

Arguments Whose Strength Depends on Continuous Variation

JAMES FRANKLIN

*School of Mathematics and Statistics
University of New South Wales
Sydney 2052
Australia
j.franklin@unsw.edu.au*

Abstract: Both the traditional Aristotelian and modern symbolic approaches to logic have seen logic in terms of discrete symbol processing. Yet there are several kinds of argument whose validity depends on some topological notion of continuous variation, which is not well captured by discrete symbols. Examples include extrapolation and slippery slope arguments, sorites, fuzzy logic, and those involving closeness of possible worlds. It is argued that the natural first attempts to analyze these notions and explain their relation to reasoning fail, so that ignorance of their nature is profound.

Résumé: Les approches dans les logiques Aristotélicienne traditionnelle et symbolique décrivent la logique en termes de traitement des symboles discrets. Pourtant, il existe plusieurs types d'arguments dont la validité dépend d'une certaine notion topologique de la variation continue, ce qui n'est pas bien saisie par les symboles discrets. On inclut comme exemples, des extrapolations, des arguments de la pente glissante, des sorites, des arguments fondés sur la logique floue, et ceux impliquant la proximité des mondes possibles. On fait valoir que les premières tentatives naturelles à analyser ces notions et à expliquer leur relation au raisonnement échouent, et que l'ignorance de cet échec est profonde.

Keywords: discrete and continuous, extrapolation arguments, fuzzy logic, open texture, possible worlds, slippery slope arguments, sorites, vagueness

1. Introduction

Both traditional Aristotelian and modern symbolic logic regard logic as a *discrete* affair: discrete symbols are manipulated according to formal rules. Even in the wider field of “informal logic,” it is regarded as an advance when there is movement towards formalization in discrete symbols and rules such as the syllogism or modus ponens, or discrete diagrams of argument structure. Topological notions like continuous variation or closeness are generally presumed to have no essential place in

logic, in the way they do in the study of space, time and motion. Yet there seems no *a priori* reason why this should be so, nor are there arguments to this conclusion beyond vague assertions that concepts that do not have sharp limits are “meaningless from the logical point of view” (Frege 1980/1896: 115) or that formal manipulation of discrete atomic symbols is “the only scheme we have for capturing normative patterns of reasoning,” and that anything else would be unintelligible (Pylyshyn 1984: 198, 51). Earlier times had no difficulty with the intelligibility of Euclidean geometry, long the paradigm of logical inference, whose reasonings involved essential reference to diagrams in continuous space (Greaves 2001: Ch. 3). If logic studies either the laws that thought should follow, or the relations between abstract entities such as propositions, no reason is immediately apparent why the processing of discrete symbols should be expected to capture all the relevant structure.

Further, there are many kinds of arguments which do depend critically on some notion of continuous variation. Let us take a simple example. The transitivity of identity is capturable in discrete symbols. For example:

Crimson is the RGB color (RGB: 211, 0, 63).
 (RGB: 211, 0, 63) is the same as (Hex: #D3003F).
 Therefore crimson is (Hex: #D3003F).

However if we weaken identity to similarity,¹ continuity enters the picture:

Crimson is very like scarlet.
 Scarlet is very like vermilion.
 Therefore, crimson is somewhat like vermilion.

The validity of the argument plainly depends on “likeness,” unlike identity, being subject to continuous variation and coming in (not necessarily numerical) degrees or gradations. Likeness thus works like closeness in space: if x is very close to y and y is very close to z , x is not far from z . That is inherent in the notions of likeness and closeness.²

¹ The debate on whether similarity is partial identity (surveyed in Morganti 2011) need not be resolved here; it is sufficient that similarity resembles identity.

² It may be said that the structure of the concepts ought to be laid out explicitly, and if they are then the original arguments are enthymemes. That is so but is no more than can be said of many other arguments, such as “ x is before y and y is before z , therefore x is before z ” (depending on the transitivity of “before”) or “Tweety is a canary and birds fly, therefore Tweety flies” (de-

It could be doubted whether similarity is “really logic” and thus whether arguments involving similarity should be counted as truly logical. That would be a sterile debate in the absence of any agreement as to what the limits of logic are (Haack 1978: Chs 1 and 9; Sher 1991: Chs 1 and 3) or even agreement as to whether logic with identity is truly logic. It is at least true that similarity, like identity, is a very general and “topic neutral” notion, crucial to argumentation across an indefinitely broad range of subject matters. The same is true of the various notions examined below.

We survey a number of fields where continuity is crucial to argumentation, with brief commentary. It is not maintained that the same analysis should apply to all of them, but since the topic is unfamiliar, it is desirable that as varied a selection of examples as possible should be on the table before analysis begins. There will then be less danger of painting oneself into a corner by becoming fixated on the special features of a single example. Naturally that necessitates some sacrifice of depth for breadth. That is inevitable in setting the scene and looking for commonalities across a range of examples.

The examples will be taken from slippery slope and extrapolation arguments, possible worlds and counterfactuals, sorites, fuzzy logic and classification, and probability and inductive logic. In all of these cases, reasoning essentially relies on some notion of gradual or continuous variation, some notion of closeness among abstract or logical entities. In each case the reader should ask, “In what space, exactly, is this variation?”

2. Slippery slope and extrapolation arguments

The first example comes from slippery slope arguments—best known in applied ethics but not confined to that area. Walton’s *Slippery Slope Arguments* distinguishes three types. First, there is the “thin edge of the wedge” argument: that if some new and small step is taken, it will create a precedent and “all hell will break loose.” Insiders of academic politics will recognize the analysis in Cornford’s *Microcosmographia Academica*:

The *Principle of the Wedge* is that you should not act justly now for fear of raising expectations that you may act still more justly in the future—expectations which you are afraid you will not have the courage to satisfy ...
The *Principle of the Dangerous Precedent* is that you

pending on the classification of canaries). The issue is that the structure of the concepts in the example does involve continuous variation.

should not now do an admittedly right action for fear you, or your equally timid successors, should not have the courage to do right in some future case, which, *ex hypothesi*, is essentially different, but superficially resembles the present one. Every public action which is not customary, either is wrong, or, if it is right, is a dangerous precedent. It follows that nothing should ever be done for the first time. (Cornford 1953: 15)

The second type of argument distinguished by Walton is close to the sorites, of which more later. It argues that “There is no cutoff point,” no natural boundary for the application of a vague term. It has been argued, for example, that there is no natural point in time between conception and birth at which a baby becomes human, so that a right to life should be extended to it immediately upon conception.

A third type of argument argues that a contemplated action would trigger a cascading series of effects that lead to disaster. The domino theory of the advance of Communism is an example, and forward defence the conclusion generally advocated. It is possible to combine all three kinds, as is usual in the euthanasia debate. There is the precedent problem, the fuzzy boundaries of notions like “human,” “alive” and “voluntary,” and the Nazi analogy of an actual causal series of events that led from euthanasia to disaster (Walton 1992: 2-6; similar in Lamb 1998: Ch. 1; Govier 1982; Burgess 1993; Lode 1999; Schubert 1994; a negative view in Spielthener 2010; doubts on the Nazi analogy in Hanaukeabel 1996).

Slippery slope arguments have sometimes been classified as fallacies. That is not in general correct, as is widely agreed in the applied ethics literature. The foundation of their effectiveness is that there really is a closeness relation between precedent and precedent, embryo-stage and embryo-stage, domino and domino, leading to the reasonableness of the inference that whatever applies to one precedent (domino, event) applies equally well (or almost equally well, or to some degree as well) to a nearby one. There is no agreement, however, as to how much they are worth. Obviously they may be overcome by evidence, such as reasons for thinking that a Nazi-like series of events “couldn’t happen here.” But defeasibility by further evidence is a feature of any non-deductive inference. The question remains, how strong is a given slippery slope argument in the absence of countervailing evidence, and how much evidence would be needed to overcome it?

If such arguments were used only in applied ethics, there might not be any strong motivation to examine them strictly as

logic. But very similar *extrapolation* arguments appear in philosophy and in science, and important conclusions sometimes hang on them.³ An example is the argument from microscopes for realism about the sub-microscopic. We tend to believe in the real existence of new things we see through magnifying glasses, because of the continuity with what we see with the naked eye. D.M. Armstrong argues:

We would trust these new deliverances because we had been able to check the native eye to some extent. We could then substitute a more powerful glass, checking *its* reliability by reference back to the original glass; and so proceed by easy stages to the most powerful microscope. (Armstrong 1961: 159-60; a similar argument of Quinton's defended in Chibeni 2006)

There has been little discussion of this argument for realism, though there is sometimes noted the weaker argument that if one wishes to give a different status to “the observed” and “the unobserved,” it is hard to draw the line between them. Another philosophical example is the argument that science has gradually purged the mystical from explanations of lightning and other physical phenomena, then of life, and can therefore be expected to conquer the last bastion of occultist explanation, the mind (Dennett 1994).

But perhaps the best reason for taking such arguments seriously as logic is that convincing examples occur in core scientific inference. Armstrong's argument concerns the extension of scientific as much as philosophical knowledge, and in modern Big Science there is often a question of cantilevering a series of instruments or methods of measurement out into the unknown. Red shift methods for measuring astronomical distances are checked against parallax methods for close stars and brightness-based methods for close galaxies. The agreement in the regions of overlap promotes confidence in the methods for more distant objects. Unfortunately the overlap of the methods is small, and there was concern about a “gap” between nearby galaxies, whose distances are pinned down by several methods, and the more distant galaxies, where relative distances are easy to find but absolute ones have been doubtful. It is claimed that recent work with the Hubble telescope has succeeded in plugging the gap (Di Benedetto 2002; Mould and Sakai 2008).

³ This is not the same issue as extrapolation in the sense of transferring the results of medical experiments on animals to humans, or other such examples of transfer of causal mechanisms, as discussed in Steel 2007.

Similar considerations apply to methods of dating ancient objects, such as calibrating carbon-14 dating against uranium-thorium results and those from counting tree-rings (the “German absolute oak chronology”) (van der Pflicht 2004). It is obvious that it is not simply a matter of deciding on the reliability of a new method in a yes-or-no fashion. However well a new method agrees with old ones in the region where they both apply, doubt accrues to the new method as it extrapolates further out into the unknown. But that doubt may be lessened by increasing the area of overlap, and by bringing more methods, even if unreliable ones, to bear.

A similar argument is crucial to Darwin’s theory at one of its weakest points. While there is plenty of evidence for the long development of species and their common descent, the evidence that chance variation with natural selection is the sole or main driving force is much thinner. There is no doubt that this mechanism can cause some changes in a population, but can it, and does it, cause speciation? Darwin argued that observed artificial selection over short time intervals can be scaled up to explain changes as large as those from one species to another. That is not an observed result, nor is it an argument from analogy; it is an argument from continuity, an extrapolation. Darwin understood how crucial the extrapolation was to his argument, which is why the first two chapters of *The Origin of Species* are entirely devoted to it. Before natural selection or homologies are even mentioned, there are many pages of mind-numbing detail about the descent of the varieties of domestic pigeon, designed to demonstrate that “domestic races of the same species differ from each other in the same manner as, only in most cases in a lesser degree than, do closely-allied species of the same genus in a state of nature” (Darwin 1859: 78). In recent times, Stephen Jay Gould has again recognized this as a crucial but relatively weak link in Darwinism, and has doubted it. He writes, “The modern synthesis drew most of its direct conclusions from studies of local populations and their immediate adaptations. It then extrapolated the postulated mechanism of these adaptations—gradual, allelic substitutions—to encompass all larger scale events. The synthesis is now breaking down ...” (Gould 1980: 121). Other extrapolation arguments relied on widely in biology include those inferring the danger of low doses of carcinogens from observations on high doses (Bolt et al 2009).

Extrapolation arguments have a resemblance to the poor cousin of Mill’s Methods, the Method of Concomitant Variations (Mill 1872: 460-70). It has been the least discussed, although it is obviously the strongest. A classic use by Galileo shows its strength and weaknesses, as well as the essential role

played in it by continuous variation. He is arguing that the Copernican system has spheres moving more slowly the farther they are from the sun, whereas the Ptolemaic system has to break that pattern suddenly by having the most distant sphere, that of the fixed stars, rotate once a day.

The improbability is shown for a third time in the relative disruption of the order which we surely see existing among those heavenly bodies whose circulation is not doubtful, but most certain. The order is such that the greater orbits complete their revolutions in longer times, and the lesser in shorter: thus, Saturn, describing a greater circle than the other planets, completes it in 30 years; Jupiter revolves in its smaller one in 12 years, Mars in 2; the moon covers its much smaller circle in a single month. And we see no less sensibly that of the satellites of Jupiter the closest one to that planet makes its revolution in a very short time, that is in about 42 hours; the next, in three and a half days; the third in 7 days and the most distant in 16. And this very harmonious trend will not be a bit altered if the earth is made to move on itself in twenty-four hours. But if the earth is desired to remain motionless, it is necessary, after passing from the brief period of the moon to the consecutively larger ones, and ultimately to that of Mars in 2 years, and the greater one of Jupiter in 12, and from this to the still larger one of Saturn whose period is 30 years – it is necessary, I say, to pass on beyond to another incomparably larger sphere, and make this one finish an entire revolution in twenty-four hours. (Galileo 1967/1632: 118-9)

Modern statistics has rechristened some of the simpler extrapolation methods “extreme value theory,” and used them with some success to predict, for example, flood peaks outside the range of data so far observed (Reiss and Thomas 2001). It is still true that the results become more unreliable as one moves further beyond the data; for example, given fifty years of data, a prediction of a once-in-a-thousand-year flood is much less secure than the prediction of a once-in-a-hundred-year flood.

In these scientific examples, it is clear what is the space in which variation occurs—in chronology it is time (and closeness of tree-rings which is presumed to reflect closeness in time), in Darwin’s extrapolation it is the closeness of species in “feature space” or number of important shared characteristics, and so on. That does not cast much light on how the strength of the arguments relates to variation in those spaces. By and large, greater closeness in the space gives a stronger argument, but it remains unclear how much, or why.

3. Conditionals and closeness of possible worlds

The next example of continuous variation in logic comes from the theory of possible worlds. A popular analysis of counterfactual conditionals like “If kangaroos had no tails, they would topple over,” is, “In all possible worlds close to the present one, in which the antecedent holds, the consequent also holds.” The question is, how much weight rests on the notion of “closeness” of worlds, how does one analyze it, and how does one know about it? David Lewis explains why closeness of worlds is crucial to the analysis:

‘If kangaroos had no tails, they would topple over’ is true (or false, as the case may be) at our world, quite without regard to those possible worlds where kangaroos walk around on crutches, and stay upright that way. Those worlds are too far away from ours. What is meant by the counterfactual is that, things being pretty much as they are—the scarcity of crutches for kangaroos being pretty much as it actually is, the kangaroos’ inability to use crutches being pretty much as it actually is, and so on—if kangaroos had no tails they would topple over.

We might think it best to confine our attention to worlds where kangaroos have no tails and *everything* else is as it actually is; but there are no such worlds. Are we to suppose that kangaroos have no tails but that their tracks in the sand are still as they actually are? Then we shall have to suppose that these tracks are produced in a way quite different to the actual way. Are we to suppose that kangaroos have no tails but that their genetic makeup is as it actually is? Then we shall have to suppose that genes control growth in a way quite different from the actual way (or else that there is something, unlike anything there actually is, that removes the tails). And so it goes; respects of similarity and difference trade off. If we try too hard for exact similarity to the actual world in one respect, we will get excessive differences in some other respect (Lewis 1973: 8-9).

Lewis did not initially give much analysis of the “similarity metric” involved in the closeness of worlds, but in response to objections later developed some suggestions on what features are important. He agreed that similarity might sometimes be relative to context, but maintained that in general, for example, a large-scale violation of laws of nature resulted in much more dissimilarity than any small change in a particular fact (Lewis 1979).

It has also been plausibly argued that a closeness-of-possible-worlds analysis is appropriate also for indicative condi-

tionals—for example, that we distinguish the indicative conditional “If Oswald did not kill Kennedy, someone else did,” from the counterfactual “If Oswald had not killed Kennedy then someone else would have” merely by the set of possible worlds we hold constant (Nolan 2003). Since nothing is more fundamental to logic than ‘if,’ a topological notion has here penetrated the heart of logic.

Notions of closeness also appear in discussion of the “truthlikeness” or “approach to truth” of scientific theories (Niiniluoto 1987: Ch. 10.3; Kuipers 1992) and in certain treatments of global philosophical scepticism, especially Nozick’s. He requires that knowledge of a proposition p should track p , but only in “close” worlds, not in distant ones like vat worlds (Nozick 1981: 240-3). He does not provide any analysis of closeness.

4. Sorites, vagueness and fuzzy logic

The first irruption of the continuous into logic occurred in the sorites “paradox” of the ancients. It has long stood as the standard and prominent example of the difficulty that discrete logic has in dealing with continuous variation. A heap of, say, grains of sand is still a heap if one grain is removed. But if one applies that rule many times, one ends up with a falsity, since having removed a sufficient number of grains, one will no longer have a heap. The problem is now studied in philosophy under the rubric of “vagueness” (Black 1970; Williamson 1994; Sorensen 2012/1997: section 3, with many references), in law under “open texture” and in other disciplines as “fuzzy logic.” The idea in each case is that many words of natural language, like “heap,” “tall,” “reasonable” and so on admit of borderline cases, and even definite cases of them can be more or less central. So removing one grain from a heap cannot make it a non-heap, but it can make it (slightly) less centrally a heap, that is, move it towards the borderline (itself imprecise) of heapness.

There is a range of philosophical approaches to the question (Hyde 2011/1997). But the main ones agree that the problem arises from a mismatch between the discreteness of language (“a heap” versus “not a heap”) and the near-continuous nature of what it describes (masses of grains of many different numbers). The definiteness of “heap” and “not-heap” does not match anything in the real-world referent. Wherever a continuum without natural boundaries is to be cut up into a discrete spectrum, a mismatch and an arbitrariness in the boundaries is inevitable. That is particularly clear in the formalization of the sorites that asks: if one grain is not a heap and very many grains

are a heap, for what number n exactly is it true that n grains are a not a heap and $n+1$ grains are? (Rescher 2008). Philosophical disagreement is mainly over whether the discrete-continuous mismatch should be thought of primarily as semantic, epistemological or ontological.

The sorites paradox illustrates the perils of concentrating on the theoretical aspects of a single example. Even if the philosophical issues were resolved, it would not help with formalizing inference with vague or fuzzy concepts. More detail is needed in cognitive science and artificial intelligence, where the aim is to *represent* fuzzy or vague predicates, for the purpose of performing inference with them in such areas as the natural language querying of databases. If a computer system is to respond appropriately to a question like “Give me the employees with fairly high salary,” it will need to represent somehow the difference between central and borderline cases of “high,” said of salaries, and how a modifier like “fairly” acts. The most common approach is the fuzzy logic one of assigning a “degree of membership” function which takes the value 1 on central members of the concept, 0 on central non-members, and values between 0 and 1 to describe how well a borderline case matches the concept (McNeill and Freiburger 1993). Questions arise such as where one is to get the actual values from, whether the values are themselves precise or fuzzy, and so on.

As remarked above, there are important examples in applied ethics. In another area of application, the question “Where do you draw the line?” which is often used rhetorically, is taken seriously, and answered, in law. It is worth surveying briefly the legal situation, since the law is required to reach solutions to real problems as they arise, and cannot afford the professional defeatism of philosophy. Since decisions must be reached, and must be reached in a manner that is reasonably consistent across cases, “drawing the line” is a matter of everyday occurrence in the courts (typical examples in Eveleth 1988; Gostin 1993; Spellman 1987), though the principles by which it is done are a matter of controversy. Legal decision-making, it is agreed, deals with “open-textured” concepts (legal jargon for fuzzy or vague). H.L.A. Hart’s *Concept of Law* explains the “fringe of vagueness” as the reason why expressing laws in rules cannot make them determinate in all cases:

In all fields of experience, not only that of rules, there is a limit, inherent in the nature of language, to the guidance which general language can provide. There will indeed be plain cases constantly recurring in similar contexts to which general expressions are clearly applicable (‘If any-

thing is a vehicle a motor-car is one') but there will also be cases where it is not clear whether they apply or not ('Does "vehicle" used here include bicycles, airplanes, roller skates?') The latter are fact-situations, continually thrown up by nature or human invention, which possess only some of the features of the plain case but others which they lack ... Faced with the question whether the rule prohibiting the use of vehicles in the park is applicable in some combination of circumstances in which it appears indeterminate, all that the person called upon to answer can do is to consider (as does one who makes use of a precedent) whether the present case resembles the plain case 'sufficiently' in 'relevant' respects. The discretion left to him by language may be very wide; so that if he applies the rule, the conclusion, even though it may not be arbitrary or irrational, is in effect a choice. He chooses to add to a line of cases a new case because of resemblances which can reasonably be defended as both legally relevant and sufficiently close. (Hart 1961: 120-4; also Bix 1991; Endicott 2000; Margalit 1979)

Problems of this kind have been with the law for a long time, and one of the oldest problems is one of the most illuminating, because it involves a concept that seems paradigmatically natural and as precise as one could hope: that of "animal". Since ancient times, it has been necessary legally to distinguish animals into "tame" (or "domestic") and "wild" (*ferae naturae*), the former being the responsibility of their owners if they damage something. The question arises in one of A.P. Herbert's misleading cases, when someone throws snails over the neighbour's fence, are snails *ferae naturae*? That is a joke, but the real problems that have arisen under the heading "What is an animal?" are hardly less bizarre. Bees are *ferae naturae*; when hived they become the qualified property of the person who hives them, but become *ferae naturae* again when they swarm. Parrots may become, but young unacclimatized parrots are not, "domestic animals." A performing bear is not a domestic animal, nor is a caged lion or a tame seagull used in a photographer's business. The phrase "bird, beast or other animal, ordinarily kept in a state of confinement" includes a ferret (James 1986: articles 'Animal', 'Domestic animal', 'Ferae naturae'). The problem is a serious obstruction to any attempt at a true formalization of legal thought, as would be necessary to achieve legal reasoning by computer (Dayal and Moles 1993; Franklin 2012).

Traditional logic textbooks often began with a discussion of "genus and species." There was some reason for that, since the terms that are the elements of propositions are the result of a

prior logical operation of classification. Whether classification of individuals into kinds is a discrete or continuous matter has important consequences. (Can a foetus become a human gradually, for example?) Aristotle has been criticized for delaying theories of evolution by imposing a doctrine of discrete, fixed and immovable species on the living world. That is the opposite of what Aristotle actually said, since he admitted continuous variation between species, even between (simple) plants and animals (Franklin 1986). Nevertheless, the historical mistake itself is instructive, a symptom of endless tension over whether classification should be discrete via trees (like Porphyry's tree, or Linnaeus'), with the tree dividing on one attribute at a time, or via a mix of features with various weights. The latter approach inevitably leads to problems with borderline cases, and one is led to picture the space of possible species as a multidimensional continuous space of features, rather than as something discrete like a tree. Psychological experiments on which style of classification humans use have yielded ambiguous results, but with at least a substantial element of a mix of features and a continuous spread of categories from exemplars (Estes 1994). A logic based on an assumed discrete classification structure of terms will miss something essential about human reasoning.

5. Logical probability with quantitative predicates

The need for topological considerations in logical spaces is nicely illustrated by an example which is common in the folklore on inductive logic (Franklin 2001; similar examples in Swinburne 1971: 326-7; Howson and Urbach 1993: 129; an ancient example in Diodorus Siculus: Bk. 3 Chs 36-7). Normally, instances of a generalization are taken to confirm it; even among inductive sceptics, they are not usually taken to disconfirm it. Consider, however, the generalization, "All humans are less than 5 meters tall." This is confirmed by present observations of people, all of whom have been observed to be less than 5 meters tall. Suppose that an expedition returns from a previously unexplored jungle and credibly reports that they have observed a human 4.99 meters tall. On the total evidence after this discovery, the probability that all people are less than 5 metres tall is nearly zero, although the generalization still has only positive instances, and in fact has more positive instances than it did before.

The lesson of the example is surely not that there is something wrong with confirmation by instances, but that informa-

tion relevant to the probability is hidden in the structure of the concepts being used. Length, as everyone who uses the concept knows, is something that admits of continuous variation, which means that the existence of something 4.99 meters tall makes it probable that there is a similar thing at least 5 meters tall. There is no justification for pretending not to know that fact, on the grounds that logic ought to be formal (that is, that instance confirmation ought to apply in the same way for all concepts irrespective of their logical structure). The dogma that logic is formal has enough difficulties even in deductive logic (for example, one cannot substitute in inference schemas such as *modus ponens* concepts that have deductive-logical oddities, such as being inconsistent or self-referential, Stove 1986: ch. 10). Nor can one apply logical probability blindly, or formally, just because it is logic, and treat concepts that have probabilistic-logical complexity as if they were simple. “Green,” “grue” and “less than 5 meters tall” have different logical structures, which can result in their behaving differently with respect to logical probability. If logic is to analyze real arguments, it must analyze the concepts used in them in enough detail to capture the inference.

Of course, the whole of non-deductive logic involves continuous variation, namely the variation in the strength with which one proposition can support another. It is hard to escape the conclusion that this is one reason why symbol-fixated logicians have tried so hard to avoid admitting that non-deductive logic really is logic, despite the strength of the reasons for doing so (for example, that it works with the strength of conjectures in pure mathematics, which is true in all possible worlds, Franklin 1987; general defence of non-deductive logic in Stove 1970). There is even a continuity argument as to why non-deductive inference should be regarded as strictly logic: if the relation between “All men are mortal and Socrates is a man” and “Socrates is mortal” is a matter of logic, the same ought to be true of the relation between “99% of men are mortal and Socrates is a man” and “Socrates is mortal.”

A final example in non-deductive inference is the tangled matter of Ockham’s Razor and inference to the best explanation. Simplicity is a property of theories in which there can presumably be continuous variation, even if measuring it is notoriously difficult. And it has been remarked that a paranoid fantasy is very like an inference to the best explanation, but with too much “explained” (Lipton 1991: 62). That suggests that there is some kind of continuous parameter measuring how eager one is to find patterns in events. Paranoiacs have it set too high, so that they see everything as part of a conspiracy. It must be possible

to have it set too low, causing one to see history as “just one damn thing after another.” If one’s parameter is properly tuned, one will be ready to fit things into an explanatory pattern, but not more ready than is reasonable. An approximately correct tuning will be essential to good reasoning.

That completes the survey of examples where the strength of inferences depends essentially on some notion of continuous variation.

6. Understanding continuous variation in logical spaces

How well do we understand such examples? How well do we need to understand them? What are the right questions to ask about continuous variation in logical spaces?

At one extreme, if one’s aim is to produce a working computer-based system to, say, perform legal reasoning, or interpret natural language for the purpose of querying a database, then one needs to understand examples like the above extremely well and in fine detail. It will be necessary to consider how to elicit from humans their methods of performing their balancing acts with competing considerations, how to represent in the computer system the continuous variation required, and how to reason with it. It is a tall order. Things are not much better from a cognitive science viewpoint. There, one must explain how humans actually do manage to reason in these areas. It is very difficult to believe that they could do so by processing discrete uninterpreted symbols, according to the programme of “symbolic AI” (McCorduck 2004: chs 6 and 11). There are of course ways of translating reasoning in continuous spaces into symbol transformations—that can be accomplished by co-ordinatizing the real line with infinite decimals, and it can be made to work in computer graphics. Still, it is not done easily, and the usual result is a very large number of symbols to manipulate. Symbolic AI *may* be able to explain plausibly how it could be done by the brain, but it will take a lot of work, and there is little evidence of a reasonable plan of attack yet. There is more hope in the approach to reasoning through “mental models” (Johnson-Laird and Byrne 1991). A mental model at least could be something that admits of continuous variation, if the mind has some kind of mental visualization facility, which allows the drawing of pictures in which spatial relations can be represented. Such a facility, called the “imagination”, was presumed to exist in older philosophy and cognitive theory (Franklin 2000). It suffered temporary eclipse in the mid-twentieth century, with Ryle, for instance, pronouncing “There are no such things as mental pic-

tures” (Ryle 1949: 254). But it has reappeared in psychology, and has had some exposure in philosophy in the “imagery debate” (Tye 1991; Blachowicz 1997; Burnett 2004).

At the other extreme (“extreme” in what space?), one may ask for a philosophical or logical in-principle understanding of what is happening in these kinds of argument. It is tempting to retreat to talk of “pragmatics” and “social practices”, as if talk about continuous variation is purely a matter of socially-constructed metaphor used for dialectical moves in “practical reasoning”. That is the view taken in Walton’s *Slippery Slope Arguments*. He argues that such arguments are “pragmadiialectical” moves designed to shift the burden of presumption, in “interactive argumentation governed by collaborative rules of politeness in speech-act conversational exchanges” (Walton 1992: 17, 20). A similar approach is taken by Stalnaker to the closeness of possible worlds. Deciding which changes to a world count as small, he says, depends on “vague conditions which are largely dependent on pragmatic considerations for their application” (Stalnaker 1981: 46).

If the pragmatic approach were essentially correct, there might be little to say about the subject from a strictly logical point of view.

Now, it can hardly be denied that the way we use counterfactuals, for example, is dependent to some extent on human interests. When we consider a world in which “things are a little different,” it does not matter to us whether there are a few more or less galaxies in distant space, although those changes may be in some absolute sense large. But to stop there is to avoid the interesting questions. Compare the analysis of the word ‘tool.’ Undoubtedly, what is classified as a tool is relative to human interests, but to say that and no more avoids the interesting questions about the objective properties of a given material that render an object made from that material useful as a tool for one purpose but not for another. The same applies to slippery slope arguments. What is it about the arguments themselves that makes them *apt* for use in dialectical or practical reasoning? Doubtless who is called “tall” and who “short” is to some extent relative to our interests, but it is an objective fact that height of people is a property that varies continuously (in one dimension), and there is no sharp divide in the population. With possible worlds, we may warp the underlying space to give prominence to changes that are of special concern to us, but what is the pre-existing structure of the underlying space being warped? It is an objective fact that a world in which kangaroos have no tails is closer to ours than a world in which all animals have no tails; no amount of warping in pursuit of special interests will make it

otherwise. More seriously, talking about “praxis” and “interests” gives us no insight into the extrapolation arguments in science, like the argument to realism from microscopes. These are *prima facie* logically similar to slippery slope arguments, and hence should be presumed to require the same analysis unless there is reason to think otherwise. The microscope argument supports realism about the sub-microscopic to the degree it does, irrespective of anyone’s interests, and we would like to know what that degree is and why. Similarly, we wish to know the objective reliability of our extrapolations in cosmic distances and ancient chronologies. The question remains, do we understand the underlying objective variation in logical spaces, and how it bears on the logical force of arguments?

If “understand” implies an ability to formalize, it seems that we know very little. Neither of the two standard formalizations of continuous variation provided by mathematics, metric spaces and topological spaces, seems adequate. A space is a *metric space* if there exists a function d from pairs of points in the space to the real numbers (which one thinks of as the distance between the points), which satisfies:

$$\begin{aligned}d(x,y) &= d(y,x) \\d(x,z) &\leq d(x,y) + d(y,z) \\d(x,y) &= 0 \text{ if and only if } x = y.\end{aligned}$$

This formalism works very well for “really spatial” spaces, like the 3-dimensional space we live in, and the “space” of cosmic distances and ancient dates. It may be adequate for such non-spatial “spaces” as the spaces of perceived colours or of probabilities. But it is too strong for many of the cases at hand. One has no idea whether it is possible to find an exact distance between one possible world and the next (Lewis’s talk of a “similarity metric” is not intended to indicate a real quantitative metric), much less one that is guaranteed to satisfy the triangle inequality (the second axiom above). If it does exist, there seems no non-arbitrary way of measuring it. And even if it did exist, it seems unnecessary to explain the possibility of the loose continuous reasoning about similarity of worlds that is actually performed.

On the other hand, the axioms for a topological space appear to be too weak. A set is said to be a *topological space* if there exists a set of subsets of it (the “open sets”), which contains the whole set, the empty set, and is closed under taking unions and under taking finite intersections. Topological spaces that are intended to describe some notion of continuity are also normally Hausdorff spaces, satisfying the extra axiom that for

every pair of points, there exist open sets containing each, which do not intersect (Dugundji 1966: 62, 137). One may well feel that one can recognize an open set in the space of possible worlds: probably the set of possible worlds which differ minimally from the present world except for making Alice smaller and smaller (without vanishing) is open. Unfortunately, that is very little help in deciding on the closeness of worlds. Since worlds can differ in many respects, it is necessary to ask which dimensions of variation are important, so as to be able to compare a large variation in one dimension with a small change in another, possibly more important, dimension. It is impossible even to ask these questions within the framework of topological spaces, since they do not support a notion of a “small” versus “large” change, or any comparison of dimensions.

One would like, then, some notion stronger than a topological space, but weaker than a metric space. There is some non-standard mathematical machinery available, such as fuzzy metric spaces, but it has yet to prove itself in applications (Kaleva and Seikkala 1984). One problem that would appear to be difficult is comparing, in possible worlds, the variation of an object in some attribute, with the object’s vanishing.

Another approach that could be taken is to construct possible worlds sufficiently simple that the answers are obvious. In Carnapian worlds with a finite number of individuals, each with a finite set of attributes chosen from a finite number of categories, it is easy to say what possible worlds are. A possible world is a point in the grid of possible choices of the attributes for all individuals. The worlds closest to a given world will be the neighbouring points on the grid, those in which one attribute of one individual is changed. That looks simple enough, at first glance, but there are still several problems which suggest one is not as well informed as one first thought. Are some dimensions of variation more important than others? Does the grid have an edge, or does it come round on itself like a grid on a torus? And how should one compare changes in attributes with the going out of existence of an individual? More important, however, is the problem that the idealization involved in Carnapian worlds has removed by fiat those properties whose nature is to vary continuously, like length and shape. As Hart says, “This would be a world fit for ‘mechanical’ jurisprudence. Plainly, it is not our world.” (Hart 1961: 125) Variation in those respects cannot be represented in the discrete on-off style of the Carnapian worlds—and surely it is variation in those respects that is important to the topological structure of the possible worlds we wish to consider.

Then, even if the formal structure of the relation of closeness were known, it would still be a further project to establish how to find it, that is, how to actually measure the closeness of worlds. And that would still leave undecided the question of how finely divisible any of these logical spaces is. The real numbers are infinitely divisible, physical space may be as well, and one at least knows how to approach deciding whether it is or not; but for logical spaces there seems little hint as to where to start.

Even if the problem of defining and measuring distances in the underlying space was solved (as it is solved in cases of extrapolation arguments based on distance or time, such as with cosmic distances and ancient chronology calibration), there remains the problem of evaluating how the strength of the argument depends on closeness in the space. That problem is very much unsolved.

The simplest case should be extrapolation arguments in one spatial dimension. One would expect that simplifying the problem sufficiently would lead to a problem solvable by standard statistical methods, but that is not the case. There is relevant mathematical machinery, but it does not solve the original questions. The natural first problem to consider is extrapolation of a function fitted to data points. Interpolation (predicting values within the range of the data points) is reasonably well understood via methods such as data smoothing, but extrapolation is much harder. John Stuart Mill justly remarked that many functional forms would give indistinguishable answers on data points (and on interpolated values), but would give wildly differing predictions outside that range (Mill 1872: 469). The problem remains unsolved. In general, without special knowledge of how the data is generated, extrapolation is regarded as unsafe. Press's authoritative *Numerical Recipes* warns, "[except when solving differential equations] the dangers of extrapolation cannot be overemphasized: An interpolating function, which is perforce an extrapolating function, will typically go berserk when the argument x is outside the range of tabulated values by more than the typical spacing of tabulated points" (Press 1992: 107). Such across-the-board pessimism is not welcomed in the ever-optimistic business world, where self-proclaimed experts have been ever ready to sell their services in "trend extrapolation" to predict the direction and scale of technological innovation (Bright and Shoeman 1973; Martino 1983, especially ch. 5; on the results see Schnaars 1989). A total pessimism is not possible in areas like nuclear power plant technology and greenhouse gas climate models, where there is no choice but to attempt inference from models calibrated on a system operating in a safe

range to the behavior of the system in a quite different, possibly unsafe regime (D'Auria et al 1995; official view on climate projection in IPCC 2007, especially FAQ 8.1). There is a good deal of engineering expertise in specific problems, but the upshot of those investigations is that extrapolation depends on knowledge of the particular problem. Tracking problems (to predict the paths of aircraft from their radar tracks, for example) require a smooth projection out from the existing track (Chui and Chen 1991, Chs 1, 8), and the brain needs to implement some similar algorithm to predict the motion of things that its organism must avoid or catch (Nijhawan 1994; Mehta and Schaal 2002). On the other hand, predictions of signals or time series of economic data normally depend on a quite different idea, the discovery of periodicities in the observed portion (Wiener 1949: Ch. 2; Chatfield 1994: Ch. 1 and section 2.6). Estimating the biodiversity of a large region from counting the species in a small subregion is a different problem again, made special by the fact that the data are distributed in space (Palmer et al. 2002; generally Haining 2003). Yet another methodology applies to purely mathematical cases like predicting a limit from a series of successively better approximations to a result (Brezinski and Redivo Zaglia 1991). All in all, existing mathematics and statistics can give only the vaguest advice on extrapolation, and most of that advice consists in warnings.

7. Conclusion

At the present stage, therefore, it must be concluded that even the simplest cases of arguments whose strength depends on some notion of continuous variation are very poorly understood. One must rest content with the purely Socratic pleasures of the realization that when it comes to any reasoning based on the continuous, there are more unsolved questions than previously thought.

References

- Armstrong, D.M. (1961). *Perception and the Physical World*. London: Routledge.
- Bix, B. (1991). H.L.A. Hart and the 'open texture' of legal language. *Law and Philosophy*, 10: 51-72.
- Blachowicz, J. (1997). Analog representation beyond mental imagery. *Journal of Philosophy*, 94: 55-84.
- Black, M. (1970). *Margins of Precision*. Ithaca, NY: Cornell University Press.
- Bolt, H.M., R. Marchan and J.G. Hengstler (2009). Low-dose extrapolation in toxicology: an old controversy revisited. *Archives of Toxicology*, 83: 197-198.
- Brezinski C. and M. Redivo Zaglia (1991). *Extrapolation Methods: Theory and Practice*. Amsterdam: North-Holland, summary in *Applied Numerical Mathematics* 15 (1991), 123-31.
- Bright, J.R. and M.E.F. Schoeman, eds (1973). *A Guide to Practical Technological Forecasting*. Englewood Cliffs, NJ: Prentice-Hall.
- Burgess, J. (1993). The great slippery-slope argument. *Journal of Medical Ethics*, 19: 169-74.
- Burnett, R. (2004). *How Images Think*. Cambridge, Mass: MIT Press.
- Chatfield, C. (2004). *The Analysis of Time Series: An Introduction*, 6th ed. Boca Raton: Chapman and Hall/CRC Press.
- Chibeni, S.S. (2006). Quinton's neglected argument for scientific realism. *Journal for General Philosophy of Science*, 36: 393-400.
- Chui, C.K. and G. Chen (1991). *Kalman Filtering With Real-Time Applications*, 2nd ed Berlin: Springer.
- Cornford, F.M. (1953). *Microcosmographia Academica*, 5th ed. Cambridge: Bowes.
- Darwin, C. (1859). *On the Origin of Species*, 1st ed. London: Murray.
- D'Auria, F., N. Debrecin and G.M. Galassi (1995). Outline of the uncertainty methodology based on accuracy extrapolation. *Nuclear Technology*, 109: 21-38.
- Dayal, S. and R.N. Moles (1993). The open texture of language: handling semantic analysis in legal decision support systems. *Journal of Law and Information Science*, 4: 330-47.
- Dennett, D. (1994). The practical requirements for making a conscious robot. *Philosophical Transactions of the Royal Society of London, series A*, 349: 71-85.
- Di Benedetto, G.P. (2002). On the absolute calibration of the Cepheid distance scale using Hipparcos parallaxes. *Astronomical Journal*, 124: 1213-20.

- Diodorus Siculus. (1st century BC). *History (Bibliotheca Historica)*.
- Dugundji, J. (1966). *Topology*. Boston: Allyn & Bacon.
- Endicott, T.A.O. (2000). *Vagueness in Law*. Oxford: Oxford University Press.
- Estes, W.K. (1994). *Classification and Cognition*. Oxford: Clarendon.
- Eveleth, J.S. (1988). Freedom or confidentiality: where do you draw the line? *Maryland Bar Journal*, 21 (Sept/Oct): 13-15.
- Franklin, J. (1986). Aristotle on species variation. *Philosophy*, 61: 245-52.
- Franklin, J. (1987). Non-deductive logic in mathematics. *British Journal for the Philosophy of Science*, 38: 1-18.
- Franklin, J. (2000). Diagrammatic reasoning and modelling in the imagination: the secret weapons of the Scientific Revolution. In G. Freeland & A. Coronas (Eds). *1543 and All That: Image and Word, Change and Continuity in the Proto-Scientific Revolution*, pp. 53-115. Dordrecht: Kluwer.
- Franklin, J. (2001). Resurrecting logical probability. *Erkenntnis*, 55: 277-305.
- Franklin, J. (2012). How much of legal and commonsense reasoning is formalizable? A review of conceptual obstacles. *Law, Probability and Risk*, 11: 225-245.
- Frege, G. (1980/1896). Frege to Peano, 29.9.1896. In *Philosophical and Mathematical Correspondence*, ed. G. Gabriel et al. Oxford: Blackwell.
- Galileo. (1967/1632). *Dialogue Concerning the Two Chief World Systems*, trans. S. Drake. 2nd ed. Berkeley: University of California Press.
- Gostin, L.O. (1993). Drawing the line between killing and letting die. *Journal of Law, Medicine and Ethics*, 21: 94-101.
- Gould, S.J. (1980). Is a new and general theory of evolution emerging? *Paleobiology*, 6: 119-30.
- Govier, T. (1982). What's wrong with slippery slope arguments? *Canadian Journal of Philosophy*, 12: 303-16.
- Greaves, M. (2002). *The Philosophical Status of Diagrams*. Stanford: CSLI Publications.
- Haack, S. (1978). *Philosophy of Logics*. Cambridge: Cambridge University Press.
- Haining, R. (2003). *Spatial Data Analysis: Theory and Practice*. New York: Cambridge University Press.
- Hanauskeabel, H.M. (1996). Not a slippery slope or sudden subversion: German medicine and National Socialism in 1933. *British Medical Journal*, 313: 1453-63.
- Hart, H.L.A. (1961). *The Concept of Law*. Oxford: Oxford University Press.

- Howson, C. and P. Urbach (1993). *Scientific Reasoning: The Bayesian Approach*, 2nd ed. Chicago: Open Court.
- Hyde, D. (2011, original version 1997). Sorites Paradox. *The Stanford Encyclopedia of Philosophy (Winter 2011 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2011/entries/sorites-paradox/>.
- IPCC (Intergovernmental Panel on Climate Change) (2007). Fourth Assessment Report: Working Group I Report, The Physical Science Basis. http://www.ipcc.ch/publications_and_data/ar4/wg1/en/contents.html
- James, J.S. ed. (1986). *Stroud's Judicial Dictionary of Words and Phrases*, 5th ed. London: Sweet & Maxwell.
- Johnson-Laird R.N. and R. Byrne (1991). *Deduction*. Hove: Erlbaum.
- Kaleva, O. and S. Seikkala (1984). On fuzzy metric spaces. *Fuzzy Sets and Systems*, 12: 215-29.
- Kuipers, T. (1992). Naive and refined truth approximation. *Synthese*, 93: 299-342.
- Kramosil, O. and J. Michalek (1975). Fuzzy metric and statistical metric spaces. *Kybernetika*, 11: 326-34.
- Lamb, D. (1988). *Down the Slippery Slope*. London: Croom Helm.
- Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. (1979). Counterfactual dependence and time's arrow. *Noûs*, 13: 455-476.
- Lipton, P. (1991). *Inference to the Best Explanation*. London: Routledge.
- Lode, E. (1999). Slippery slope arguments and legal reasoning. *California Law Review*, 87: 1469-1543.
- Margalit, A. (1979). Open texture. In A. Margalit (Ed), *Meaning and Use*, pp. 141-52. Dordrecht: Reidel.
- Martino, J.P. (1983). *Technological Forecasting for Decision Making*, 2nd ed, New York: North-Holland.
- McCorduck, P. (2004). *Machines Who Think*, 2nd ed. Natick, Mass: A.K. Peters.
- McNeill D. and P. Freiburger (1993). *Fuzzy Logic*. Melbourne: Bookman.
- Mehta B. and S. Schaal (2002). Forward models in visuomotor control. *Journal of Neurophysiology*, 88: 942-53.
- Mill, J.S. (1872). *System of Logic*, 8th ed. London: Longmans, Green, Reader and Dyer.
- Morganti, M. (2011). The partial identity account of partial similarity revisited. *Philosophia*, 39: 527-546.

- Mould J. and S. Sakai (2008). The extragalactic distance scale without Cepheids. *Astrophysical Journal*, 686: L75-L78.
- Niiniluoto, I. (1987). *Truthlikeness*. Dordrecht: Kluwer.
- Nijhawan, R. (1994). Motion extrapolation in catching. *Nature*, 370 (28 July): 256-7.
- Nolan, D. (2003). Defending a possible-worlds account of indicative conditionals. *Philosophical Studies*, 116: 215–69.
- Nozick, R. (1981). *Philosophical Explanations*. Oxford: Clarendon.
- Palmer, M.W., P.G. Earls, B.W. Hoagland, P.S. White and T. Wohlgemuth (2002). Quantitative tools for perfecting species lists. *Environmetrics*, 13: 121-37.
- Press W.H. et al. (1992). *Numerical Recipes in C*, 2nd ed. Cambridge: Cambridge University Press.
- Pylyshyn, Z.W. (1984). *Computation and Cognition*. Cambridge, Mass: MIT Press.
- Reiss R.D. and M. Thomas (2001). *Statistical Analysis of Extreme Values: With Applications to Insurance, Finance, Hydrology and Other Fields*. Basel: Birkhäuser.
- Rescher, N. (2008). Vagueness: a variant approach. *Informal Logic*, 28: 282-294.
- Ryle, G. (1949). *The Concept of Mind*. London: Hutchinson's.
- Schnaars, S.P. (1989). *Megamistakes*. New York: Free Press.
- Schubert, L. (2004). Ethical implications of pharmacogenetics: do slippery slope arguments matter? *Bioethics*, 18: 361-78.
- Sher, G. (1991). *The Bounds of Logic: A Generalized Viewpoint*. Cambridge, Mass: MIT Press.
- Sorensen, R. (2012, original version 1997). Vagueness. *The Stanford Encyclopedia of Philosophy (Summer 2012 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2012/entries/vagueness/>.
- Spellman, R.L. (1987). Fact or opinion: where to draw the line? *Communications and the Law*, 9 (Dec): 45-61.
- Spielthener, G. (2010). A logical analysis of slippery slope arguments. *Health Care Analysis*, 18: 148-163.
- Stalnaker, R. (1981). A theory of conditionals. In W. L. Harper, R. Stalnaker and G. Pearce (Eds.), *Ifs*, pp. 41-55. Dordrecht: Reidel.
- Steel, D. (2007). *Across the Boundaries: Extrapolation in Biology and Social Science*. Oxford: Oxford University Press.
- Stove, D.C. (1970). Deductivism. *Australasian Journal of Philosophy*, 48: 76-98.
- Stove, D.C. (1986). *The Rationality of Induction*. Oxford: Clarendon.

- Swinburne, R.G. (1971). The paradoxes of confirmation – a survey. *American Philosophical Quarterly*, 8: 318-29.
- Tye, M. (1991). *The Imagery Debate*. Cambridge, Mass: MIT Press.
- Van der Pflcht, J. (2004). Radiocarbon calibration – past, present and future. *Nuclear Instruments and Methods in Physics Research B*, 223-4: 353-8.
- Walton, D.N. (1992). *Slippery Slope Arguments*. Oxford: Clarendon.
- Wiener, N. (1949). *Extrapolation, Interpolation and Smoothing of Stationary Time Series*. New York: MIT Press.
- Williamson, T. (1994). *Vagueness*. London: Routledge.