David Lewis's Philosophy of Language[1]

Richard Holton

Lewis never saw philosophy of language as foundational in the way that many have. One of the most distinctive features of his work is the robust confidence that questions in metaphysics or mind can be addressed head on, and not through the lens of language.[2]

That said, Lewis's work in the philosophy of language has had enormous impact. The Lewis-Stalnaker analysis of counterfactuals provides the standard; so too does Lewis's account of convention. His account of truth in fiction is widely accepted, and has provided the inspiration for recent fictionalist accounts of various domains.[3] Elsewhere he provides, if not the orthodoxy, then the standard alternative to orthodoxy: this is true, for instance, of his counterpart semantics for modal language, and of his descriptivist account of the semantics of singular terms.

Moreover, Lewis is a member of that small group of philosophers whose work is read and used by linguists, at least those working towards the semantic end of things. A check through the index pages of *Linguistics and Philosophy* reveals, as expected, that he is one of the handful of authors whose work is multiply cited in every volume by both linguists and philosophers. 'Scorekeeping in a Language Game' is a central piece in pragmatics. 'General Semantics' is one of the founding texts of much contemporary formal semantics. 'Adverbs of Quantification' has

---

[1] Thanks to Sam Guttenplan, Lloyd Humberstone, Rae Langton, Bob Stalnaker, Mark Steedman and Steve Yablo. I use the following abbreviations for Lewis's collections of papers: *PP* I: *Philosophical Papers* Vol. I (New York: Oxford University Press, 1983); *PP* II: *Philosophical Papers* Vol. II (New York: Oxford University Press, 1986); *PPL*: *Papers in Philosophical Logic* (Cambridge: Cambridge University Press, 1998); *PME*: *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 1999); *PESP*: *Papers in Ethics and Social Philosophy* (Cambridge: Cambridge University Press, 2000). On first significant mention, papers are given with their original date of publication; but page references are to the reprints in the collections.

[2] The 'refusal to take language as the starting point in the analysis of thought and of modality' is one of the six 'recurring themes' which Lewis cites as representative of his approach in the Introduction to *PP* I, p.!xi.

[3] Starting with Gideon Rosen 'Modal Fictionalism' *Mind* 99 (1990) pp.!327–54.

provided inspiration for Discourse Representation Theory.[4] Amongst this wealth of material, I cannot hope to provide a summary. Instead, I shall try to highlight some central themes, and to say something about Lewis's distinctive approach to them.

In discussing language, as elsewhere, Lewis's work is systematic. Indeed, the systematicity is forced by his cost-benefit methodology. If one holds that the philosophical theories to be embraced are the ones that bring least cost, then one needs to know the ramifications of a theory across the board. Typically the costs are not local. I think that this helps explain why, despite having his heart in the details, Lewis was drawn to such systematic work; it isn't simply a matter of succumbing 'to the temptation to presuppose my views on one topic when writing on another', as he self-deprecatingly puts in it.[5]

Within the philosophy of language the systematicity comes out most noticeably in the project of achieving the 'integration of formal semantics into a broader account of our use of language in social interaction'.[6] Thus whilst we find pieces concerned predominantly with formal semantics (e.g. 'General Semantics'; 'Adverbs of Quantification', *Counterfactuals*), and pieces concerned predominantly with the social and pragmatic features of language (e.g. 'Scorekeeping in a Language Game', Truth in Fiction'), these are almost always written with an eye to the connections. All are concerned with giving an account of language and language use which is truth-conditional and intensional, couched in the framework of possible worlds.[7] In other pieces the connections are to the fore ('Languages and Language', much of *Convention*.). I shall structure my discussion under the three broad headings of Semantics, Pragmatics and Connections; but throughout my emphasis too will be on the connections.

---

[4] See Irene Heim, 'E-Type Pronouns and Donkey Anaphora' *Linguistics and Philosophy* 13 (1990) at p.!145.
[5] Introduction to *PP* I, p.!ix.
[6] Another of the six unifying themes; *PP* I, p.!xi.
[7] Though see ''Tensions' (1974) *PP* I pp.!250–60 for an argument that the distinction between intensional and extensional accounts is, at the very least, elusive.

SEMANTICS

The core of the formal semantic model that Lewis endorses is outlined in 'General Semantics' (1970)[8]. A categorial grammar generates the sentences at the base level and determines their semantic interpretations, interpretations which are understood in terms of truth conditions. A sentence is assigned an intension, a function from indices to truth values, where indices include information on the context of utterance (time, place etc.) and the possible world of utterance, and an assignment of objects to variables.[9] These base level sentences are then mapped onto the surface level sentences by a transformational component, which makes the necessary changes (such as alteration of word order and introduction of new morphological elements) to give us the ordinary sentences of English. Lewis says very little about the transformational component, other than to remark that he has in mind a model along *Aspects* lines.

'General Semantics' is one of the key papers, along with Montague's, in founding the approach which sees each syntactic rule as matched by a corresponding semantic rule (the 'rule-to-rule' approach). It is also one of the papers that marks the beginning of the recent resurgence of interest in categorial grammars. If certain features of his proposal are no longer current (such as the restriction of the basic categories to sentence, name and common noun, or his treatment of long range dependencies) then that is largely testament to the vibrancy of a field that Lewis helped to shape.

Within this general framework Lewis proposes ways of treating a large number of specific linguistic phenomena. Always the project is to provide a truth-conditional account of constructions which seem to resist it. In most cases, these proposals are not inseparably tied to Lewis's framework, but have been adopted by theorists working in other intensional truth-conditional frameworks. I shall briefly outline some of the most influential:

---

[8] *PP* I pp.!189–232.

[9] Following Kamp and others, Lewis later came to hold that two indices were needed to capture all of the features of English (to understand the use of the operator 'actually' we need to specify *both* the world with respect to which the sentence is being evaluated, and the world which is actual); and, following Cresswell, that one of these indices would need to be open ended, in that one could not specify in advance what sort of information it might need to contain. See 'Index, Context and Content' (1980), *PPL* 21–44.

Modality Lewis holds that modal discourse involves quantification not just over the occupants of the actual world, but over the occupants of other possible worlds too. Modal operators can then be replaced by standard quantification over individuals. Unlike in standard quantified modal logic, no object can exist in more than one world. But objects can (and typically will) have *counterparts* in other worlds, the counterpart relation being understood purely in terms of resemblance.[10] Lewis, of course, takes the talk of possible worlds realistically; but whether that is avoidable is an issue that will not be broached here.[11]

Counterfactuals Lewis holds that a counterfactual 'If it were the case that A, it would be the case that C' is true just in case, *very* roughly, the closest world at which A is true is a world at which C is true.[12]

Indicative conditionals Lewis holds that indicative conditionals 'If A then C' have the simple truth conditional semantics of the material conditional, so that they are true in every case except where A is true and C is false. However, they have the assertion conditions of the associated conditional probability P(A/C).[13]

---

[10] 'Counterpart Theory and Quantified Modal Logic' (1968) *PP* I, pp.!26–46; for further discussion see *On the Plurality of Worlds* (Oxford: Basil Blackwell, 1986) pp.!5–20

[11] The bulk of *On the Plurality of Worlds* consists of an extended argument that the many benefits of possible world talk, of which those for philosophy of language are only a part, cannot legitimately be enjoyed by those who do not give it a realistic construal.

[12] The account sketched is actually closer to that give by Robert Stalnaker in papers published a few years before Lewis's (see especially 'A Theory of Conditionals' in N. Rescher (ed) *Studies in Logical Theory* (Oxford: Blackwell, 1968)). Lewis developed his account independently. To get an idea of how Lewis's account diverges, consider what would happen if there were no single closest world in which A were true; either because there were ties, or because there were an infinite sequence of A-worlds, such that for each one, one could always find one closer. For the full account see *Counterfactuals* (Blackwell: Oxford, 1973), and for a summary 'Counterfactuals and Comparative Possibility' (1973) *PP* II pp. 3–31; for a discussion of differences between Lewis's account and those of Stalnaker, Pollock and (within a rather different framework,) Kratzer, see 'Ordering Semantics and Premise Semantics for Counterfactuals' (1981) *PPL* pp.!77–96.

[13] 'Probabilities of Conditionals and Conditional Probabilities I' (1976) *PP* II, pp.!133–56. In the appendix to that paper Lewis abandons the details of his own positive account in favour of a slightly amended version of the account put forward by Frank Jackson. For further discussion see 'Probabilities of Conditionals and Conditional Probabilities II' (1986), *PPL* pp.!57–65.

Adverbs of quantification Lewis holds that adverbs such as 'Always', 'Usually' 'Often' 'Sometimes' 'Rarely' and 'Never' should be understood as quantifiers ranging, not over times or events, but over cases.[14] Such adverbs are 'unselective quantifiers', quantifiers that bind *all* the free variable within their scope, and not, like standard quantifiers, just the variables that they explicitly mention. Thus in 'Usually, x reminds me of y iff y reminds me of x' the adverb 'Usually' binds both x and y; the sentence is true just in case most pairs of persons (i.e. most cases) satisfy the open sentence that follows it.

Declarative Performative Sentences Against Austin, Lewis holds that declarative performative sentences (e.g. 'I command you to open the door') have truth conditions in a straightforward way.[15] 'I command you to open the door' is true just in case I command you to open the door; and since, typically, there is no more to commanding than uttering the appropriate sentence with the appropriate intent, such utterances will typically make themselves true. (But contrast 'I *legitimately* command you to open the door' which requires much more if it is to be true.)

PRAGMATICS

Far and away Lewis's most influential paper in this area is 'Scorekeeping in a Language Game' (1979).[16] Central to the paper is the idea of *accommodation*. It is a familiar idea that language is contextual: the correct interpretation of a sentence is partly determined by the context in which it is uttered. Thus, for instance, the interpretation of 'now' is determined by the time of utterance of the sentence in which it occurs. Lewis's contribution is to take this idea much further. He shows, not only that the contextual factors are more pervasive than they initially appear, but that the relevant context for an utterance is often provided by the very utterance itself. The result is that uttering a sentence can change the 'conversational score' so that that

---

[14] 'Adverbs of Quantification' PPL pp.!5–20.
[15] 'General Semantics' p.!224; 'Scorekeeping in a Language Game' *PP* I, pp.!247–8. Lewis was here drawing on the work of others, most notably Lemmon: for details see 'Scorekeeping' p.!247 n.7. Rather less convincing is Lewis's contention that *all* sentences are at the base level declarative: that 'Open the door!' is the result of the operation of a transformational rule on a base level sentence along the lines of 'I command that you open the door' ('General Semantics' pp.!222–4).
[16] *PP* I pp.!233–49.

very sentence is made acceptable or even true. This is the process of accommodation. Thought of in this way a wide variety of pragmatic phenomena can be unified; and once again this is done in a way which renders them amenable to treatment in a truth-functional framework.

To see how this works, take first Lewis's treatment of presupposition[17]. If someone says "Even George can do it", their utterance requires the presupposition that George is not generally much good. So we might expect that the sentence can only be successfully uttered in a context in which that presupposition is *already* granted. But that isn't so. Uttering the sentence itself *creates* the presupposition (unless someone blocks it): the context accommodates the presupposition by changing so that the sentence uttered is acceptable.

Sometimes the accommodation affects not just the acceptability of a sentence but its truth. When I say that we can't now get to London in time for the conference, I don't mean that it is physically impossible to do so: a speed far short of the speed of light would take us there with hours to spare. Nor need I even mean that a chartered private helicopter that took us door to door would be too slow. Rather I mean that none of the standard options, the kind of thing that we might be prepared to pay for, would get us there in time. So my modal needs to interpreted relative to a context whereby only these worlds in which these options are realized are the worlds that count. Once again, this context doesn't need to be already in place: uttering my sentence puts it in place, and so provides the context against which my utterance can be true.[18]

To accord an important place to the process of accommodation is not to say, of course, that it will *always* guarantee truth. When I say that we can't get to the conference in time, that is simply false if our normal train would get us there easily. There are constraints on which worlds can be excluded by the process of accommodation. Similarly there are constraints on which presuppositions can be introduced by accommodation. If a sentence has just been asserted and accepted, you can't introduce its negation by uttering another sentence that presupposes it. But within the constraints, accommodation does a lot of work.

---

[17] This is taken from Stalnaker; see his 'Presuppositions', *Journal of Philosophical Logic* 2 (1973) pp.!447–57.

[18] Here Lewis builds on Angelika Kratzer's work; see 'What 'Must' and 'Can' Must and Can Mean' *Linguistics and Philosophy* 1 (1977) pp.!337–5.

Lewis applies the notion of accommodation to many other cases. I mention three: the process whereby performatives such as promises make themselves true; the process whereby a domain is restricted so that definite descriptions really do come out as uniquely denoting; the process whereby standards of precision for vague terms are fixed so that sentences containing them can come out as true.

This last case is of particular interest since it provides the basis of Lewis's account of knowledge, sketched in one paragraph in 'Scorekeeping' and developed in detail nearly twenty years later in 'Elusive Knowledge' (1996).[19] The account of knowledge itself is simple: I know P just in case my evidence eliminates every world in which not P.[20] But when we say 'every world' what do we mean? Knowledge for Lewis is a vague term, and so the standards shift. Do I know that I am at work at my desk? Yes, by the ordinary standards, since the set of worlds that I need to eliminate consists of the relatively mundane ones in which I am in the kitchen, or out for lunch or whatever; and these I can eliminate by simple observation. But once the set of worlds is increased to include those in which the evil demon tricks me, my evidence no longer eliminates them all, and so my claim to knowledge fails.

How is the relevant set of worlds determined? Lewis gives a set of conditions on which worlds are 'properly ignored' i.e. are kept out of the set that need to be eliminated by my evidence. Central is the *Rule of Attention*: a rule of accommodation whereby any possibility that is mentioned automatically joins those that need to be eliminated if we are to claim knowledge. The result is that introducing sceptical possibilities in the process of doing epistemology itself destroys knowledge; once I say that I might be dreaming, this becomes a possibility that needs to be eliminated.

Hasn't this gone too far? The murderous axeman knocks at your door; you have hidden your friend upstairs. "Do you know where he is?" demands the axeman. "Interesting that you ask", you reply, "since I was just thinking that this might all be a bad dream. No, I don't know where he is." This might indeed confuse the axeman, but it scarcely gets one off the charge of lying. It seems that the parties to the conversation have to somehow *accept* the new standards as relevant; here the mere

[19] *PME* pp. 418–45.

[20] As Lloyd Humberstone has pointed out to me, this definition appears to commit Lewis to an S5 logic for knowledge, rather at odds with the general externalist tenor of the account; although whether this appearance can be sustained in the light of the shifting standards is unclear.

mention of sceptical possibilities does not bring such acceptance. Indeed, even if the axeman had briefly considered the possibility you mentioned, he could surely have returned to the original standards easily enough ("Enough of this tomfoolery: do you really not know where your friend is?"). Such worries do not discredit Lewis's pragmatic framework for the account of knowledge; but they do raise questions about the particular rules he gives.[21]

In closing the discussion of 'Scorekeeping', it is interesting to contrast the process of accommodation with the more familiar Gricean mechanisms. The Gricean mechanisms exploit a distinction between what is *said* and what is *implicated*, and do not affect the truth conditions of the latter. If I say that war is war, what I have said is a tautology; but I have implicated something more significant. In contrast, accommodation affects the truth conditions of what is said. Normally the two processes will work side by side; but we might wonder whether there are cases where they will offer rival explanations. Suppose there are two birds in the garden; I have noticed only the rosella and not the treecreeper. I say "The bird in the garden is a rosella". Have I spoken falsely, as a Gricean approach would suggest (there's no such thing as 'the bird'), even though I manage to implicate something true (that the bird I have noticed is a rosella)?[22] Or does my utterance constrain the domain of eligible objects by means of the process of accommodation, so that only the Rosella falls within it , and hence my utterance comes out true?

CONNECTIONS

Finally let us turn to the work in which Lewis most explicitly treats the connections between *languages*, thought of as formal semantic models, and *language* thought of as a social phenomenon. In 'Language and Languages', Lewis analyses conventions

[21] I doubt that this example even counts in favour of the rule Lewis offers to the fallibilist—i.e. that the propositions we know are those that eliminate every possibly except those that *should be* ignored—at least not if the fallibilist is held to be claiming that even in the epistemology classroom we retain our knowledge; see p.!436.
[22] Implication is here understood in a very liberal sense, along the lines discussed by Kripke in 'Speaker's Reference and Semantic Reference', P. French, T. Uehling and H. Wettstein (eds.) *Contemporary Perspectives in the Philosophy of Language* (Minneapolis: University of Minnesota Press, 1977) pp.!6–27. There is good reason to think that Grice himself would have been highly suspicious of such a liberalization of his notion.

as, roughly, regularities to which people conform because their belief that others conform gives them pragmatic reason to conform.[23] A given language L (understood formally) can then be said to be used by a population of speakers (a concrete group of people), just in case there is a convention of *truthfulness* and *trust* in L: *truthfulness*, in that the speakers utter only true sentences of L; and *trust* in that they believe others to be truthful in L, and so come to believe what they say. And since the formal semantics work in terms of truth conditions, it is a straightforward matter to see what it would be for a sentence to be true. The formal is connected to the social.

Does this account fully determine the language which is used by a population? There are two issues here. First, knowing that there is a convention of truthfulness and trust to use a certain language leaves open the *mechanisms* by which the interpretation of that language is fixed. Take the relation of names to their bearers. It could be, for instance, that we had a convention of using a Lagadonian language in which objects name themselves, and so are related by identity.[24] Or it could be that we use a language in which names are related to their bearers by causal chains; or one in which they are related by descriptions. We need to know how things are for English, and for the other natural languages in which we are interested. Second, there is the issue of whether, even given this information, the language is fully determined: whether there is semantic indeterminacy, perhaps of the innocuous kind in which our conventions simply do not fix whether some particular word is properly assigned one sense or another; or of a more systematic and disturbing kind, for instance as discussed by Quine, or by Goodman and by Wittgenstein, as interpreted by Kripke. I take these two issues in turn.

Despite the influence of causal accounts of reference, Lewis maintains a staunchly descriptive theory. That is, he holds that the reference of an expression is

---

[23] This is rough, since in addition Lewis holds (i) that there must be other regularities that could do the job (or else the regularity that is followed would be forced upon the people involved, rather than being followed as convention); (ii) that the people must have a general preference that others conform (or else we would have deadlocked conflict), and (iii) that all of these facts must be mutual knowledge amongst those involved. The account differs in one respect from the earlier account offered in *Convention* (Oxford: Blackwell, 1969), an account which was based on Lewis's doctoral thesis. In the earlier work the definition had been framed in terms of a conditional preference to conform if others conformed, rather than in terms of each person having reason to conform if the others conformed. The advantages of the revised account are spelled out in 'Languages and Language' pp.!169–71.

[24] *On the Plurality of Worlds* p.!145

determined by a more or less precise description, associated more or less explicitly with that expression by speakers of the language. He does not deny that there might be an important causal element; but this is subsumed within the descriptive account. The causal chains are to be understood as picked out by indexical elements, themselves understood in descriptive terms: 'the object at the end of such-and-such causal chain that led to this idea', or similar. It might indeed be the case that we will be unable to identify what stands at the end of the causal chain by looking at the description in the mind of the speaker; but then that is so for many indexical elements. If the speaker does not know the time, we will not come to know it by examining her understanding of the term 'now'. What we have here is an instance of the now familiar distinction between content and character. The description in the mind of the speaker fixes the character, even if this is not enough to fix the content.[25]

In later work Lewis identifies a number of further features that descriptions are likely to have if the descriptive theory is to be plausible. As well as containing reference to causal chains, they might contain rigidifiers ('the man who *actually* invented the zip'); they might contain egocentric indexical elements ('the stuff that plays the watery role on *this* planet'); and they might make reference to word or thought tokens ('the man called 'Moses' by us'). Moreover, it might be that a given description will not perfectly pick out any referent. In such cases we might have to go to the best deserver of the term, and then it might be indeterminate whether that one is good enough, or whether the term should be thought of as lacking reference. Equally troubling, there might be two or more equally good deservers, and then it might be indeterminate whether one, or both, or neither should be thought of as the true referent.[26] The descriptive theory will thus bring with it a kind of moderate indeterminacy.

Will the descriptive theory tell the *whole* story of how reference is fixed? Can we think of *every* word as having its reference fixed by a description? This needn't lead us to a regress. We would just have to think of all of the words having their meanings fixed together by something like a huge set of simultaneous equations. Certainly Lewis thinks that large numbers of terms can have their reference fixed together in this way. Indeed that is what typically happens in a scientific theory: large

[25] 'Languages and Language' pp. 172–3

[26] for the best summary of this see 'Putnam's Paradox' (1984) *PME*, at pp. 59–60.

numbers of terms are simultaneously defined as referring to whatever it is that best satisfies the theory.[27] But this cannot be done for the whole of language at once; that, Lewis thinks, is the true moral to be drawn from Putnam's model theoretic argument. For if it could, then, providing there was no inconsistency, it would be easy enough to ensure that *every* sentence came out true no matter what: simply assign meanings to terms in whatever way is necessary to ensure truth.[28] But that is not how language works. Contingent sentences run the risk of being false: that is part of what gives them content. So there must be some constraint on the interpretation of our words in addition to that provided by a form of global descriptivism.

What could this constraint be? An obvious candidate is a causal constraint, but as we have seen, Lewis denies this much significance. In so far as he accepts causal factors at all, they just become part of the description, and hence cannot provide an independent constraint. Instead, he thinks that the necessary constraint is provided by the primitive relation of naturalness. Some properties—those of flawless physics perhaps—are perfectly natural. Others—those defined relative to human interests for instance— are less so. But even here there are differences. 'Funny' is less natural than 'Fatal'; 'Grue' is less natural than 'Green'. The idea of primitive naturalness does a great deal of work in Lewis's later philosophy, working to characterize such diverse things as physicalism, supervenience, lawhood and intrinsicness.[29] In addition, it provides the necessary constraint on interpretation.[30] Whilst satisfying their associated descriptions, words are to be assigned meanings that are as natural as possible.[31] That is what prevents us from permuting the meanings so that every sentence comes out true.

These considerations have consequence for the second issue to be discussed here, that of indeterminacy of reference. In 'Languages and Language' Lewis took Quine's worries seriously, arguing that we should be suspicious of any account that

---

[27] 'How to Define Theoretical Terms' *PP* I pp.!78–95.
[28] If the logical constants were also to be defined in this way, then even consistency would not provide a constraint: one would simply interpret so as to avoid inconsistency.
[29] 'New Work for a Theory of Universals' (1983) *PME* 8–55; Defining "Intrinsic"' (1998), with Rae Langton, *PME* 116–32.
[30] 'Putnam's Paradox'; 'New Work for a Theory of Universals' pp.!45–55.
[31] Or perhaps we should say 'appropriately natural'. Perhaps we should require that the terms of sociology come out less natural than the terms of biology, and so that any interpretation that reversed that order would be inferior to one that did not.

attempted to tell us whether ‘gavagai’ referred to enduring rabbits or to their time slices.[32] Similarly, at that time he took seriously the indeterminacy worries raised by Goodman and later by Kripke, at least in so far as they seemed to show that it was impossible to extrapolate a general rule from a sample.[33] But once he had accepted natural properties, he came to reject the Goodmanian worries: we classify things as green rather than grue, since the former is more natural.[34] We might wonder whether Lewis could make the same reply to Quine. That depends on whether we think that there is a strict ordering of naturalness amongst the possible candidates, and if there is, on what it is. Certainly we might think that the idea of a rabbit is more natural than that of an undetached rabbit part. So some of Quine’s candidates for the meaning of ‘gavagai’ might be ruled out. But how do we adjudicate between enduring rabbits and rabbit time slices? It depends on which we think more natural. For most people, I suspect it would be the rabbit. For Lewis, that is less clear; indeed, given his views on the nature of persistence, we might expect the time slices to be more natural.[35]

Whatever should be said about Quinean indeterminacy, on Lewis’s account a great deal of piecemeal indeterminacy remains. Mention has already been made of that generated by his descriptivism; indeterminacy in the counterpart relation, and in the relation of similarity between worlds needed for counterfactuals, give rise to more.[36] I conclude with one more source. Since the late 1960s, it has become the orthodoxy that natural kind terms like ‘water’ are rigid: on other planets (actual or possible) where other substances play the role of water, those substances are not properly described by us, as ‘water’. Lewis is sceptical. He suggests that ordinary English might well be indeterminate on this issue too (an idea given some credence

---

[32] Pp.!177–8.
[33] As later characterized in ‘Meaning without Use: Reply to Hawthorn’ (1992) *PESP* pp.!145–51, at p.!150.
[34] ‘New Work’ p.!48; ‘Putnam’s Paradox’ pp.!64–8.
[35] For those views, see ‘Postscript B’ to ‘Survival and Identity’ (1976) *PP* I pp.!76–7; *On the Plurality of Worlds*, pp.!202–4; ‘Rearrangement of Particles: Reply to Lowe’ (1988) *PME* pp.!187–95. In fact though Lewis doesn’t accept that on his view rabbit time slices are any more natural than rabbits; he thinks that, if anything, they are slightly less so (‘Putnam’s Paradox’ p.66 n.14). This might have something to do with his insistence that in understanding persistence in terms of perdurance he is not involved in the project of reduction (‘Postscript B’, p.!77); but I am not clear on what is at issue here.
[36] On the former see Appendix C to ‘Counterpart Theory and Quantified Modal Logic’, *PP* I pp.!42–3; on the latter see *Counterfactuals*, pp.!91–5.

by the difficulty one has in getting undergraduates to see the point). If analytic philosophers are convinced that they are rigid, this might be because of the dialect that has become standard amongst them: a dialect whose acceptance has much to do with the prestige of Kripke and Putnam.[37]

When we try to capture indeterminate parts of language with formal models, we will need to take care, for the formal models are typically not so indeterminate. Of course there are devices that can help us—supervaluation, of which Lewis was an early and steadfast proponent, springs to mind.[38] But the main lesson is simply that we should not expect determinate answers to questions that do not have them.

---

[37] 'Reduction of Mind' (1994) *PME* pp.!291–324, at pp.!312–14.

[38] 'General Semantics' pp.!228–9. Lewis also suggests that the machinery of relevance logic could be usefully employed to prevent the explosion of logical consequences in cases where ineliminable ambiguity leads to contradiction: 'Logic For Equivocators' (1982) *PPL* pp.!97–110.