

David Lewis (1941–)

ROBERT STALNAKER

Introduction

David Lewis is a philosopher who has written about a wide range of problems in metaphysics and the philosophy of mind and language, including the metaphysics of possible worlds, the analysis of counterfactual conditionals, causation and probability, the problems of universals, of intentionality, of personal identity, the foundations of decision theory, of set theory, of semantics. A distinctive and comprehensive metaphysical theory has emerged from his discussions of philosophical problems: a theory that combines realism about possible worlds with a kind of nominalism, a materialist account of mind, and Humean skepticism about unanalyzed natural necessity. But Lewis's discussions have also yielded conceptual tools that have applications both within and outside of philosophy that are independent of the grand metaphysical scheme, for example, an analysis of common knowledge that has been influential in game theory and theoretical computer science, and work on generalized quantifiers in natural language and on the role of extra-linguistic context in the interpretation of speech that has influenced the development of linguistic semantics.

Lewis studied at Harvard with W. V. Quine and Nelson Goodman, and the influence of those two philosophers is evident in his own philosophical method, in the problems he has focused on, and in the substance of the views he defends. But Lewis developed Quinean and Goodmanian themes with a distinctive twist that takes them in unanticipated directions and that has resulted in a theory that combines features his teachers would applaud with features they would abhor. For Quine and Goodman, the rejection of the analytic/synthetic distinction motivated a holistic philosophical methodology, a method of "reflective equilibrium" that helped to make metaphysics respectable for the heirs of the positivist tradition (see GOODMAN and QUINE). (One consequence of abandoning the two dogmas of empiricism, Quine wrote, was "a blurring of the line between speculative metaphysics and natural science," Quine 1953: 20.) Lewis adopted the holistic method, and accepted the invitation to do metaphysics with a clear conscience, but he defended the analytic/synthetic distinction, and the intelligibility of truth by convention. He followed Goodman in seeking a reductive analysis of counterfactual conditionals, but rejected Goodman's demand for a reduction of the possible to the actual. He adopted Quine's standards for ontological commitment, and for philo-

sophical clarification, but used them to reach very different conclusions about what there is, arguing that Quine's "creatures of darkness" – intensions, propositions, possible worlds – *can* find a place in a world-view that meets the rigorous standards of adequacy that Quine set down. The actual world, according to the metaphysical theory Lewis defends, is much as Quine and Goodman thought. Their only mistake was to think that the actual world is the only world there is.

The emphasis in this exposition will be on the general metaphysical framework that provides the context for Lewis's many constructive philosophical analyses. I will begin with some general remarks about philosophical method and metaphysics in the next section, and after that discuss Lewis's modal realism and finally his Humean account of counterfactuals, laws, and causation.

Method and metaphysics

During the first half of the twentieth century, the word "metaphysics" had mostly a pejorative use within the analytic philosophical tradition. The logical empiricists taught that metaphysics was the result of equivocation between questions about meaning, which called for a decision about what linguistic framework to use and questions that arise within the context of an accepted framework. But Quine noted that the methods used within the scientific framework for deciding which theoretical claims were true were not very different from the methods used to make the practical decisions about what language forms to adopt. In both cases, one chose the theory or framework that did the best job of making sense of one's experience. He argued that the line between internal and external questions, and between decisions that constituted linguistic stipulation and decisions that constituted empirical judgments was arbitrary. If decisions about what general framework to theorize in are not separable from judgments about what is true, then there is room for metaphysics after all. "The quest of a simplest clearest overall pattern of canonical notation," Quine wrote, "is not to be distinguished from a quest of ultimate categories, a limning of the most general traits of reality" (Quine 1960: 161).

Lewis's account of his philosophical method follows that of Quine and Goodman closely. We begin with a collection of opinions. "Some are commonsensical, some are sophisticated; some are particular, some general. . . . A reasonable goal for a philosopher is to bring them into equilibrium," Lewis 1983b: x) And like Quine, Lewis emphasizes that the method of reflective equilibrium should not be taken to have relativist or anti-realist consequences. Philosophy may be a matter of opinion, but some opinions, even some that are in some philosopher's reflective equilibrium, may nevertheless be false. But unlike Quine and Goodman, Lewis did not tie his epistemological holism to the rejection of the analytic/synthetic distinction. His first major philosophical project responded to Quine's critique of this distinction, and of truth by convention. Lewis accepted the terms of Quine's demand for an analysis: one must break out of the tight circle of concepts (synonymy, semantic rule, meaning, etc.), and explain what it is to be an analytic truth in terms of the dispositions and behavior of language users. But he argued that this could be done with the help of a general analysis of the notion of a *convention*, and a distinction between two different notions of language: language as defined by a set of syntactic and semantic rules and language as defined by a popula-

tion of speakers. The definition of an abstract language simply stipulates that the language is constituted by certain semantic rules that determine a class of analytic truths. The work is done in explaining, in terms of an analysis of convention, what it is about the behavior, expectations, dispositions of a given population of speakers for a language defined in this way to be the language spoken by that population.

There are two ways in which Lewis's account of analyticity, even if fully adequate on its own terms, will fail to satisfy an unreconstructed Quinean. First, conventions are explained in terms of intentions, beliefs, and knowledge, and so an explanation of semantic notions such as meaning and analyticity in terms of convention would not be an explanation that solved the problem of intentionality. Quine thought that to the extent that mentalistic intentional notions such as belief and intention could be explained at all, they would be explained in terms of the intentionality of language – believing, for example, in terms of holding true – and so he would not be satisfied with an explanation of semantic notions that took belief and intention for granted. In this regard, Lewis is like H. P. Grice, separating problems about linguistic meaning from the more general problem of intentionality, and taking the intentionality of thought as more basic. But Lewis, Quine, and Grice would all agree that whichever comes first, an adequate account of linguistic and mental intentionality must ultimately explain them in materialistically acceptable terms.

A second way in which the account will disappoint a Quinean is perhaps the more significant one. As Lewis emphasized, his account of analyticity makes reference to possible worlds, and so does not provide an informative analysis of one of the notions – necessity – that Quine would have put in his tight circle of problematic concepts. But Lewis argued that the metaphysical notions of necessity and possibility do not belong in this circle, since they are not semantic notions. Analytic truths are necessary because they express propositions that are necessary. The account of the conventions of language explain why the sentences used by the members of some population express the propositions they express, but the necessity or contingency of the propositions themselves has nothing to do with convention, or with language.

Though he defends analyticity, Lewis does not assume that speakers are authoritative about the conventions of their own language, so about the analytic truths. Even if there is a sharp line between truths of meaning and truths of fact, there is no sharp line between linguistic intuition and beliefs about substantive theory. "Our 'intuitions' are simply opinions, and our philosophical theories are the same" (Lewis 1983b: x). We can draw the line between analytic and synthetic, but the decision about where we draw it, like our other decisions about what to believe, is a part of a judgment about the global theory that, all things considered, best makes sense of our experience.

Modal realism

Possible worlds have played a prominent role in Lewis's philosophical analyses from the beginning. Being a good Quinean, Lewis recognized an obligation either to admit them into his ontology, or to reduce them to something else. And if they are to be accepted, it should be clear what kind of thing they are. "We ought to believe in other possible worlds and individuals," he argues, "because systematic philosophy goes more smoothly in many ways if we do" (Lewis 1986b: 354). Lewis makes no attempt to mini-

mize the counterintuitive character of the ontological commitment he is prepared to make; possible worlds, as he uses the term, are concrete particulars: other things of the same kind as the universe of which we are a part. Merely possible highways in lands that will never be actual are made of concrete that is just as real as that used to make the actual highways on which we drive. Just as people who live at other times and places in the actual world are as real as we are, so, according to Lewis's modal realism, are the non-actual people who inhabit other possible worlds. Their nonactuality consists in the fact that they are spatiotemporally disconnected from us. Lewis grants – in fact emphasizes – that the belief in a plurality of parallel universes conflicts sharply with common opinion, and since he takes common opinion seriously, he acknowledges that this is a serious cost to be balanced against the benefits of this metaphysical theory. But while he takes the “incredulous stare” that is a common response to this theory to reflect a formidable objection, he argues that more theoretical arguments against modal realism fail, as do attempts to analyze possible worlds away, or to give a more innocent explanation of what they are. In the end, he judges that the cost of offending common opinion is outweighed by the many benefits that modal realism brings. So the strategy for defending modal realism combines an exposition of the many benefits of the possible worlds framework, responses to theoretical arguments against modal realism, and arguments against attempts to get those benefits without the counterintuitive commitment.

It is useful to divide the doctrine of modal realism into a semantic and a metaphysical component. First, there is the metaphysical thesis that there is a large plurality of parallel universes, where a single universe consists of everything that is spatiotemporal related to anything in it. Second, there are the semantic analyses that relate this plurality of worlds to the many modal, epistemic, and intentional concepts whose clarification provide the benefits of modal realism. As an example of a thesis that belongs to the second component, consider the analysis of possibility as truth in some possible world. Lewis emphasizes that the theory must be evaluated as a package, and he would agree that each component would lose all plausibility without the other. On the assumption that the metaphysical thesis is false – that common sense is right that our universe is the only one – the semantic analysis of possibility has no plausibility, since on that assumption the possible collapses into the actual. On the other hand, if we look at the metaphysical claim in isolation from the semantic analyses, it looks like an extravagant and gratuitous empirical hypothesis. Why should one believe in all these other universes? Lewis's answer – that systematic philosophy goes more smoothly if we do – has force only when the metaphysical hypothesis is combined with the semantic analyses that connect the hypothesis with the phenomena that systematic philosophy seeks to explain.

The possible worlds framework promises to clarify not only *de dicto* modal claims, such as that it is necessary that all bachelors are unmarried, but also *de re* modal claims such as that no bachelor is *essentially* unmarried. Modal realism uses *counterpart theory* to analyze claims about the modal properties of things. As with the general modal realist thesis, we can distinguish a metaphysical and a semantic component of Lewis's counterpart theory. There is the metaphysical claim that individuals exist in only one possible world, and the semantic claim that *de re* modal properties should be analyzed in something like the following way: an individual has the property of being possibly

F if and only if it has a counterpart that has the property of being *F*, where the counterpart relation is a contextually determined relation of similarity in relevant respects. As with the general thesis, Lewis would emphasize that the semantic and metaphysical parts of the package must be evaluated together.

The metaphysical doctrine has been criticized (for example, by Alvin Plantinga and Nathan Salmon) on the ground that it has the implausible consequence that all properties are essential properties; but this criticism simply assumes that Lewis's semantic analysis of what it is to have a property essentially is mistaken. The semantic thesis has been criticized (for example by Saul Kripke) on the ground that it has the consequence that when we say that Humphrey might have won the election, we are not really talking about *Humphrey*. But Lewis rightly insists that, on the counterpart analysis, it is Humphrey himself who has, in the actual world, the property of being a possible winner. It is just that he has this property in virtue of his resemblance to someone else who (in another possible world) has the property of being a winner. The counterpart semantics may be more complex and less straightforward than the standard analysis, but given the metaphysical thesis, it gives a better account of our modal beliefs. And if one accepts the general doctrine that other possible worlds are parallel universes, it seems most reasonable to think that no one can inhabit more than one of them. Whatever one's verdict about the plausibility of modal realism as a whole, it seems clear that counterpart theory belongs in the package.

Both modal realism's central metaphysical thesis and its semantic analyses of necessity, possibility, and other modal notions conflict with unreflective common opinion. It is not only that it strains credibility to hypothesize that there is a vast plurality of parallel universes, it also seems counterintuitive to many people to claim that our opinions about what might or would have happened are opinions about the existence of such parallel universes. As noted above, Lewis grants that modal realism conflicts with unreflective common opinion, and that this conflict is a strike against the theory, but he argues that the cost is outweighed by the benefits. Since he agrees that if some alternative account could provide the benefits without the cost, modal realism would not be defensible, it is an important part of its defense to criticize attempts to reconcile the explanations that the possible worlds framework provides with a more modest account of what possible worlds are.

"Ersatz modal realism" is Lewis's label for the attempt to get the benefits of modal realism without the costs by explaining possible worlds as something other than parallel universes. Most of his critical discussion of this project is devoted to attempts to reduce possible worlds to some kind of linguistic object: state descriptions, maximal consistent sets of sentences, complete novels. This is a common and seductive strategy for explaining what possible worlds are, but there is a lot wrong with it, as Lewis's criticisms bring out. The most serious problem is that this kind of explanation seems to foreclose one of the most important uses of possible worlds: to represent the contents of speech acts and propositional attitudes. If sentences, or sets of them, are to represent possible worlds adequately, they must be interpreted sentences – sentences with their truth conditions. We will have a serious circularity if we try to combine this kind of explanation of possible worlds with an explanation of the truth conditions of a sentence in terms of the possible circumstances, or possible worlds, in which the sentence would be true. There are, however, philosophical accounts of possible worlds that agree

with Lewis that possible worlds are non-linguistic things, suitable for representing truth-conditional content, while disagreeing with the thesis that possible worlds are something like other universes parallel to our own.

According to the simplest and most straightforward attempt to explain what possible worlds are in a way that is compatible with actualism (a thesis that common opinion might regard as trivially true: that what actually exists is all there is), possible worlds (or less misleadingly, possible states of the world) are a kind of property: ways the world might be, or might have been. This is obviously not a reduction of possible worlds to something else: it is intended simply as a characterization of the kind of thing that a possible world is. To call possible worlds properties is to say two things about them: first, they are things that are, or may be, instantiated. Second, they are the kind of thing that is (at least *prima facie*) independent of language and thought. A property of something (such as the property of being the first child born in the twenty-first century) is different both from the thing (if any) that has the property and from a thought or a predicate that expresses the property. The significance of this characterization is that it provides us with a way to reconcile a commitment to the existence of possible worlds, construed as non-mental entities, with the apparently contradictory thesis that the actual world is the only world there is.

Anyone who takes literally the claim that there are possible worlds has to respond to this *prima-facie* paradox: unrealized possibilities – counterfactual situations – are situations that turned out not to exist. How can there *be* situations, or worlds, that don't exist? Any response to this problem will make a distinction between a sense in which non-actual possible worlds exist, and a sense in which they do not. Lewis's strategy is to distinguish two different scopes for the quantifier. Quantifiers are often restricted to some contextually determined subdomain of all there is, and one very general restriction, according to Lewis, is to the domain of things that inhabit the actual world. When we say that there are no talking donkeys, we normally mean that there are no *actual* talking donkeys. But there is also an unrestricted quantifier, which ranges over absolutely everything there is. Common opinion may not distinguish what exists from what actually exists, but Lewis would say that the distinction is implicit in their modal discourse. The actualist response to this puzzle makes the distinction, not in terms of a difference of domain, but in terms of an ambiguity in the terms "possible world" and "actual world." Just as we can distinguish the property of being the first child born in the twenty-first century from that child, so we can distinguish the property of being a universe of a certain kind from a universe that is of that kind. According to the actualist, there are (and *actually* are) many ways the world might have been, but there is only one world that is one of those ways.

This construal of possible worlds as properties allows us many of the benefits of the possible worlds framework (for example, the formal semantic analysis of modal and epistemic notions, the clarification of counterfactuals and causal and temporal structures, the representation of probability as a measure on state spaces, the representation of mental and linguistic content, and of speech contexts) without either denying the ontological commitment to possible states of the world, or challenging pretheoretical common opinion. It does not produce incredulous stares to say that there are many ways the world might have been. To see why Lewis resists this actualist interpretation of possible worlds we need to consider another one of his metaphysical priorities that

has its root in the Quine–Goodman legacy: a penchant for nominalism (see GOODMAN and QUINE).

One of the benefits of modal realism, according to Lewis, is that it provides us with an analysis of properties: of what properties are, and what it is to have a property. Properties, according to Lewis's modal realism, are just sets, and to have a property is just to be one of its members. The domain of all possibilities provides an answer to the standard objection to the identification of properties with their extensions, an answer that those with only the impoverished domain of the actual things cannot avail themselves of. Distinct properties can have the same extension in the actual world, but they are distinguished by the difference in their extensions in other possible worlds. Even if in the actual world, all and only creatures with a kidney are creatures with a heart, the two properties are distinguished by the fact that there are possible creatures with one property, but not the other. So modal realism offers the virtues of a simple extensionalist account of properties without the defects of actualist versions of that account.

Lewis recognizes that an adequate theory of properties needs to distinguish between different kinds of properties. Some properties (sets) and relations (sets of n -tuples) are *natural* or *fundamental*. Among the fundamental relations, some are *spatiotemporal* relations. These are primitive distinctions of Lewis's theory, but he argues that they are distinctions that any plausible metaphysical theory must make. With just these primitive concepts for classifying properties and relations, Lewis suggests, we can give a full characterization of the logical space of possible worlds while continuing to maintain that properties are nothing but sets.

World-mates (inhabitants of the same possible world) are individuals that stand in spatiotemporal relations with each other. A possible world is fully characterized by specifying a set of world-mates, and by saying which fundamental properties they have, and how they are related by the fundamental relations. All the properties and relations of the things in any world *supervene* on the fundamental properties and relations of those things: possible worlds that are indiscernible from each other with respect to fundamental properties and relations are identical. This a priori supervenience claim is not substantive, since the fundamental properties are just those that are necessary to give a complete characterization of a possible world. Substantive (and contingent) metaphysical hypotheses can be stated as theses about what the fundamental properties of the actual world are. So, for example, *materialism* is explained as the thesis that only physical properties and relations are fundamental. (That is, materialism is true of possible world w if the fundamental properties and relations of things in w are all physical.) The thesis of *Humean supervenience*, which we will discuss in the next section, is the thesis that only intrinsic properties and spatiotemporal relations are fundamental.

The theory of properties as sets is a crucial part of Lewis's modal realism, and unlike many of the fruits of the possible worlds framework, this analysis cannot be reconciled with the actualist interpretation of possible worlds. Lewis's theory can, of course, make the distinction to which the actualist appeals between properties and their instances – it is just the distinction between sets and their members – but it will be no help in avoiding a commitment, not just to ways the world might be, but to worlds that are those ways. For if properties are sets, and if possible states of the world are identified with maximal properties that the world might have, then a possible state of the world is a

unit set with a world that is in that state as its member. According to the actualist metaphysics, there is only one thing to be the member of such a set, and so if possible states of the world are properties, and Lewis is right about what properties are, there is only one possible state of the world. So while actualists can avail themselves of many of the benefits of the framework of possible worlds, they will have to forego Lewis's elegant reductive account of properties.

Counterfactuals and causation

Lewis's second book undertook to give a reductive analysis of counterfactual conditionals, a project motivated by the same Humean skepticism about natural necessity that motivated Nelson Goodman to try to give such an analysis more than twenty years earlier. For Goodman, the crux of the problem was the modal character of counterfactuals: they seemed to be about unrealized possibilities. The task was to explain the possible in terms of the actual. Lewis, of course, had no problem with non-actual possibilities, and took counterfactuals at face value as statements about counterfactual possible worlds. But counterfactuals (and statements about cause and effect, dispositions and propensities, dependency and chance) are for the most part contingent statements. One has to explain how a statement about counterfactual possibilities can be contingently true or false in the actual world. For such statements to be contingent, the counterfactual worlds that are relevant to the evaluation of a conditional must be determined by their relation to the actual world.

Lewis's formal semantics gives truth conditions for conditionals in terms of a three-place *comparative similarity relation* on possible worlds (world x is more similar to world w than y is to w). The rough idea of the analysis is that a conditional, "if A , then C " is true (in a possible world w) if and only if C is true in those possible worlds in which A is true that are most similar to w . This first approximation is not quite right, since if there is an infinite sequence of ever more similar worlds in which A is true, there will be no closest such possible worlds. To allow for this case, Lewis's favored analysis is as follows: "If A , then C " is true in w if and only if some world in which $A \& C$ is true is closer to w than any world in which $A \& \sim C$ is true. This analysis provides an abstract formal semantics for counterfactual conditionals, but we don't have a reductive analysis until we have explained the relevant respects of similarity. The semantic analysis is just the first step of a larger project, a defense of the doctrine that Lewis labeled "Humean supervenience." The project is motivated by a Humean skepticism about real relations between "distinct existences."

For the Humean, spatiotemporal relations (such as contiguity) are acceptable, as are logical relations, or relations of ideas. Relations of resemblance between things are acceptable, so long as the respects of resemblance are spelled out, since they are explicable in term of the sharing of specified properties. But causal relations, and others in the same family, must be analyzed in terms of global regularities, with the help of relations of the unproblematic kind. If the respects of similarity between possible worlds that are relevant to the interpretation of counterfactuals can be specified, Lewis's analysis will yield an account of counterfactuals that should satisfy a Humean, and so an account that permits the Humean to use counterfactuals to analyze relations of causation and causal dependence and independence.

Lewis's Humean project has the following separable components: (1) an abstract semantic analysis of conditionals in terms of comparative similarity; (2) an explanation of the respects of similarity that are appropriate for interpreting of the kind of conditionals that are relevant to the analysis of causal relations. Since Lewis's response to this problem appeals to laws of nature, he needs (3) an account of laws of nature in terms of global patterns of particular fact, and finally, (4) an analysis of causation in terms of counterfactuals. This is an ambitious agenda. Some parts have been carried out in detail and with precision; in other cases, there are only sketchy suggestions about the kind of account that should be given. And some parts of the project are ongoing.

A conditional is true if the consequent is true in the possible world in which the antecedent is true that is most similar, in relevant respects, to the actual world. But what are the relevant respects? One might be tempted to appeal to an intuitive notion of overall similarity. Lewis notes that we do make and understand judgments of overall similarity between complex object such as cities, and we do have intuitions about which possible worlds are more and less alike. A general impressionistic notion of similarity would be both vague and context-dependent, but as Lewis notes, counterfactuals are both vague and context-dependent. There would, however, be at least two problems with relying on such a notion of similarity. First, an impressionistic notion of similarity would be suitable for the project of Humean reduction, since judgments of similarity between worlds might be based in part on comparison of unanalyzed facts about causal relations. But second, in any case there are counterexamples that show that overall similarity is not the right relation. It seems intuitively clear that small events can have large consequences. If Oswald had missed Kennedy in 1963, the course of American politics between then and now probably would have been quite different. But isn't a possible world in which Oswald misses, but someone else succeeds, and the course of American politics proceeds much as it actually did much more similar, overall, to the actual world? If certain conspiracy theorists are right, and there were backup assassins ready to act if Oswald failed, then it might be true that if Oswald hadn't killed Kennedy, someone else would have, but we don't want an analysis of counterfactuals to ensure that such conspiracy theories are true.

One might be tempted to build a temporal asymmetry into the account of comparative similarity that is relevant to the interpretation of counterfactuals: perhaps similarity of earlier parts of history should have much greater weight than similarity of later times. But to do this would be to explain the temporal asymmetry of causal and counterfactual dependence as a consequence of convention and not as a fact about the world. Lewis's aim was to define a temporally neutral notion of comparative similarity between possible worlds, and use it to explain how temporal asymmetries in the pattern of facts in the actual world results in a *de facto* asymmetry of counterfactual dependence.

Lewis's account of the relevant respects of comparative similarity between worlds gives highest priority to avoiding large and widespread violations of laws of nature. The second priority is to maximize *exact* agreement of particular fact. Small and local violations of laws of nature are permissible to achieve the second priority, and in a deterministic world, such "small miracles" will always be required. *Approximate* agreement of fact counts for very little: deviations from the laws, even small ones, in order to increase approximate similarity of fact are not permitted, and possible worlds that agree

exactly for a period of time are more similar than worlds that agree only approximately, but over a much longer period of time. There is no attempt to make this account of comparative similarity precise, but Lewis argues that it helps to explain some temporal asymmetries, and that it points to the kind of explanation that can vindicate a reductive account of counterfactuals.

A Humean cannot, of course, rest with an appeal to an unanalyzed notion of law of nature. Here Lewis endorses an idea of Frank Ramsey's: that the laws of nature (in a given possible world) are the factual regularities that are consequences of the simplest and strongest systematization of the truths of that world. The criteria for evaluating systems of truths remain to be explained, but Lewis argues that so long as the relevant criteria of strength and simplicity are non-contingent, this will be an account of law of nature that meets the standards of Humean supervenience.

Counterfactuals, once explained in terms of resemblance of the facts and regularities of the possible worlds, are then available to the Humean for an analysis of causation. The first step is to define counterfactual dependence: one truth *B* counterfactually depends on another *A* if *B* would not have been true if *A* had not been true. If *c* and *e* are distinct events that occur, it seems a good first approximation to say that *c* is a cause of *e* if and only if both occur, and *e* would not have occurred if *c* had not. This proposal would account for many of the features of causation that create problems for a simple regularity analysis. Cause can be distinguished from effect without making explicit appeal to temporal order, and events that are regularly connected because one causes the other can be distinguished from events that are connected because they are each effects of a common cause.

But cases of preemptive causes show that one cannot, in general, identify causation with counterfactual dependence. Suppose the hit man was successful, but if he had missed, another was waiting in the wings to do the job. The victim's death was caused by the hit man's action, but, because of the backup potential cause, was not counterfactually dependent on it. Lewis's first strategy for accommodating preemptive causation was to define the causal relation as the transitive closure of the relation of counterfactual dependence (between distinct events). The death does not depend counterfactually on the shooting, but there will be intermediate events which are dependent on the shooting, and on which the effect is dependent. This move accounts for some cases of preemption, but not for all. A second strategy for dealing with preemption cases argues that while the man would still have died if the backup assassin had done the job, he would have died a different death, and so despite the preemption, the event that was the effect was still counterfactually dependent on the actual assassin's act. But it is difficult to find and motivate an account of the modal properties of events that will explain all cases of preemption in this way without intuitively implausible consequences. Even taking the resources of counterfactual conditionals for granted, the analysis of causation has proved to be a surprisingly recalcitrant problem. This is now a lively area of ongoing research.

Conclusion

Lewis's metaphysical framework, and his philosophical method, provide a rich context for the clarification of philosophical problems, the articulation and defense of philo-

sophical theses, and the formulation of constructive conceptual analyses. He has formulated and defended a materialist theory of mind, with accounts both of intentional states such as belief, and sensory states such as pain. In the context of the defense of Humean supervenience, he explored the relation between objective and subjective probability – chance and degree of belief – and between subjective probability and conditional propositions. And he provides a foundation for causal decision theory. Building on his early work on convention, he developed a foundation for semantics and pragmatics that clarifies the relations between speech and thought, and that also makes substantive contributions to compositional semantic theories for natural languages. Modal realism’s use of set theory motivated an exploration of the foundations of set theory itself that clarifies the relation between mereology (the theory of parts and wholes) and set theory.

Even though Lewis’s general metaphysical theory has a coherence and unity that tie the different parts together, many of his constructive analyses are separable from the system that provides the context for their development. This is appropriate, given Lewis’s pragmatic cost–benefit methodology: he recognizes that others with different priorities may not be prepared to swallow his system whole. Those who reject modal realism or Humean supervenience will still find much to learn and to adopt from his philosophical work. And even those who are skeptical about this metaphysical theory can appreciate the power of a system that has generated so many clarifying philosophical analyses.

Bibliography

Works by Lewis

1969: *Convention*, Cambridge, MA: Harvard University Press.

1976: *Counterfactuals*, Oxford: Blackwell Publishers.

1983a: “New Work for a Theory of Universals,” *Australasian Journal of Philosophy* 61, pp. 343–77.

1983b: *Philosophical Papers*, vol. I, Oxford: Oxford University Press.

1986a: *On the Plurality of Worlds*, Oxford: Blackwell Publishers.

1986b: *Philosophical Papers*, vol. II, Oxford: Oxford University Press.

Works by other authors

Quine, W. V. (1953) *From a Logical Point of View*, Cambridge, MA: Harvard University Press.

— (1960) *Word and Object*, Cambridge, MA: MIT Press.