Précis of *Depth*

Michael Strevens

Draft of January 2012

1. Explanation and Causation

Many scientific explanations have the form of deductive derivations from laws or other generalizations. Hempel and Oppenheim's (1948) equation of explanation with law-involving deduction is stymied, however, by cases such as the flagpole and the shadow: two deductions, formally almost identical, have quite different explanatory properties. The length of the flagpole's shadow can be explained by deriving it from the height of the flagpole and the position of the sun, but the height of the flagpole is certainly not explained by deriving it from the length of its shadow. The former derivation owes its explanatoriness, apparently, to something that goes beyond mere logical deduction.

The modern reaction to this case (not so different from that insinuated by Aristotle in *Posterior Analytics* 1.13) is that the deduction is only a representational instrument: the meat of the explanation inheres in what the deduction stands for, which is almost universally taken to be a causal relation. A deduction of the length from the height, if done right, represents the causal process by which the flagpole and sun together cause the shadow. The reverse deduction fails to represent a causal process, because there is no such process running from shadow to pole. What explains, in short, are the causal facts—the facts that constitute the process causing the explanandum (the thing to be explained). A law-involving deduction is explanatory only because it represents such facts.

If that is the right approach to understanding scientific explanation, then three items must appear on the agenda of the philosophical study of explanation:

1. What are causal relations?

2. How do deductive derivations and other semantic apparatus in science represent causal relations?

3. Besides the need to represent the causal production of the explanandum, what other norms govern the construction of scientific explanations?

The great preponderance of work in this area since Hempel has struggled with the first question. *Depth* attempts to look beyond that question, without resolving it, to the other two, most especially the third.

As to causality: when Wesley Salmon (1984) and other philosophers first attempted to reintroduce causality to the philosophy of explanation, their primary task was to convince logical empiricists and post-empiricist anti-realist types alike that causal relations provided an ontologically serious, more or less observer-independent basis on which to found explanatory practice. These days, causation proliferates in the philosophy of science. So do philosophical accounts of causation: there are successors to Salmon's account (Dowe 2000); counterfactual accounts (Lewis 1973, 2000); interventionist variations on counterfactual accounts (Woodward 2003); neo-regularity accounts (Baumgartner 2008); and more. On the one hand, the philosopher of explanation can rest assured that the status of causality in the philosophy of science is secure. On the other hand, if an account of scientific explanation is to be built on causality, then there are as many rival approaches to explanation as there are to causation.

What is an explanationist to do? He might leave causation to the experts—but that would be to place it back in the hands of devilish metaphysicians, whose expulsion from the garden by Carnap, Reichenbach, Hempel and others constitutes the heroic founding myth of modern philosophy of science. Or he might attempt the metaphysics of causation himself, in the hope that one more blow of a head well enough endowed will suffice to break through the wall.

Or . . . might it be possible to say illuminating things about explanation while bracketing the whole question of causation? It is in that spirit that *Depth* is written.

Suppose, then, that the universe is held together by some kind of causal connective tissue—call it *causal influence*. Something has to be assumed about causal influence, if not its deep, underlying nature. I assume that it is the kind of causal relation that we find in (or perhaps project onto—but then I said I would bracket these questions) fundamental physics. Newtonian force is my paradigm; I help myself to the assumption that scientists see something similar in modern physics.

I then show how a deductive argument, or similar formal structure, can represent causal influence. The successful execution of this task will depend on what kind of thing causal influence is, so you will find only a schema and not a full account in *Depth*. The idea is taken straight from the case of the flagpole and the shadow: each step in a deduction can, in the right circumstances, represent a step in a causal process. Such a step will, roughly speaking, use modus ponens to deduce, from a causal law and initial conditions, a causal consequence of those conditions. When a deductive argument is of the right sort to represent a causal process that has as a consequence the state of affairs asserted to hold by the argument's conclusion, I say that the argument *causally entails* the conclusion. A derivation that causally entails its conclusion, then, is one that can be used to represent (whether truly or not) a causal process that has the concluding state of affairs as its consequence.

Two remarks. First, causal entailment is a property of an argument or derivation, thus Hall's objection to the notion of causal entailment that there is “no such thing as *the* derivation of some claim from some other claims” is misplaced—the relevant derivation is simply the derivation to which the property of being a causal entailment is ascribed (or from which it is withheld). Second, the property of being a causal entailment is partially extrinsic. It depends, in particular, on the nature of causal laws and so on the nature of causality.[1]

2. Causal Difference-Making

The story so far: the world is one big bundle of causal influence. For any event or state of affairs of explanatory interest, we can use a deductive argument that causally entails its conclusion, or some other formal device, to represent various causal influences that together sufficed to bring about the explanandum. Is that all there is? To explanation, I mean?

It is only the beginning. What comes next is to develop an account of which elements of the causal web make a difference to the explanandum’s occurring or holding, and which do not. Only the difference-makers are explanatory, I maintain, following a grand tradition in which Garfinkel (1981), Lewis, and Salmon are only a few of the names of note.

How can something be a causal influence but not a difference-maker? Let me use Salmon’s (1997) example. A ball goes astray during a local baseball game; it hits a window and the window breaks. Among the causal influences on the window are both the ball and the shouts of the players, the latter because sounds cause the window to vibrate. But of these two influences, only the ball makes a difference to the breaking. Why? Both affect what happens to the window as it breaks. The influence of the shouting is subtle; it perhaps affects

1. Hall complains also, in a footnote, about my schematic characterization of a causal law. With some justice: I allowed an extraneous thought about grounding to invade my formulation in the passage he quotes. He might have noted, however, that my official schema for causal lawhood is spelled out immediately following that passage, and is not subject to his objection (first full paragraph of p. 77 of *Depth*).

the precise pattern of the shattering, hence the exact shapes of the shards and their exact trajectories. But these influences do not (except under very unusual circumstances) steer the trajectory of the window so as to determine whether it breaks or not. They make a difference to the *way* that the window breaks, but not to the fact that it does break. Thus they do not explain the breaking.

Crucial to this distinction is a certain feature of the explanandum: it is not a concrete event but what I call a high-level event. A concrete event is an event individuated by all of its intrinsic properties. The concrete event of the window's breaking, for example, would not have occurred if the breaking had proceeded in even a very slightly different way (at a slightly later time, or with a slightly different shattering pattern, for example). As Hempel noted long ago, we do not typically seek to explain concrete events. We explain high-level events or states of affairs: in the case of the window, we explain the fact that the window shattered within a certain time interval, or something with similarly liberal individuation conditions, hence something that would have happened even if the shattering had gone a little differently. This something is what I call a high-level event.

When explaining a concrete event, all causal influences are difference-makers, since it takes the concerted efforts of every causal influence to steer the relevant elements of the world to that precise state that constitutes the concrete event's occurrence. But when explaining a high-level event, some influences (or rather, some facts about some influences) are decisive in steering things into the state where the event occurs and some are not. The non-decisive influences' hands are on the wheel, but they make a difference not to whether the world's trajectory enters the high-level state, but only to what part of that state is entered—they make a difference to how the event occurs but not to whether it occurs or not.

What criterion should be used to distinguish difference-makers, then? To what test can we subject each of the causal influences on an event, in order to determine which are difference-makers and which are not? An old suggestion (Salmon's) is to impose a probabilistic test: an influence makes a difference only if the probability of the occurrence of the event conditional on the influence is higher than the event's unconditional probability. A suggestion tempting to many contemporary philosophers is a counterfactual test: an influence makes a difference to a high-level event if, had the influence not been present, the event would not have occurred.

Though they are not without problems, both criteria have much to be said for them in certain limited spheres of application, and in particular when applied to the explanation of discrete, local occurrences. They are less successful when applied to, say, standing states of affairs. Suppose (to use an example developed in *Depth*) you want to explain why sodium reacts so readily with a wide variety of substances. Relevant to the explanation is sodium's loosely bound outer electron; not so relevant is its exact number of neutrons. The former makes a difference to sodium's reactivity, while the latter does not.

How to use the probabilistic or counterfactual tests to diagnose these facts about difference-making? To apply the probabilistic test, you must ask yourself: what would be the probability of sodium's being reactive if it did not have a loosely bound outer electron? It is doubtful that we have a probability distribution capable of answering this question, that is, a probability distribution that ranges over various alternative chemistries in which sodium has a rather different form. Yet we answer the question easily; we cannot, then, be answering it probabilistically.

To apply the counterfactual test, you must ask yourself: in a nearby possible world where sodium lacked a loosely bound electron, would it react readily? To answer this question, it seems, you must determine what form sodium would have in such a world and how that form would interact with the rest of the world's chemistry. (You must also determine whether, in the light of this radical reconstruction of the chemical facts, the world in question is similar enough to the actual world to qualify as one of the *closest* possible worlds in which sodium lacks a loosely bound outer electron.) Expertise in these matters is perhaps just barely within our human grasp, but it seems doubtful that the philosophically innocent chemists who see so clearly the relevance of the electron are performing such advanced metaphysical calculations.

For this reason, and several others, I propose a new test for difference-making, a test that captures the judgments that are so ably provided by probabilistic and counterfactual tests in the case of spatiotemporally discrete events, but that is also easily applied to explananda such as sodium's reactivity by investigators whose expertise is scientific rather than philosophical.

This *kairetic test*, as I call it, is straightforward (and well described by the other contributors to this symposium): Begin with a deductive derivation correctly representing some set of causal influences that causally entails the explanandum (here I will ignore stochastic processes, though they are treated in *Depth*). Make this causal model as abstract as you can, by replacing its highly specific descriptions of the influences in question with more abstract descriptions of the same. When you have proceeded as far as you can—when any further abstraction would invalidate the entailment of the explanandum—you have a model containing only difference-makers for the explanandum. To be a difference-maker, then, is to appear in such a model.

What remains in a fully abstracted model are typically not specifications of individual causal influences but rather specifications of aggregative properties of large sets of influences. In the case of the broken window, for example, though the ball is a difference-maker, the abstracted model will not specify the positions of the individual particles making up the ball, or for that matter the exact number of such particles. This is because a precise specification of the positions or the number can be made more approximate without invalidating the entailment of the breaking. The end product of the abstraction process, then, will specify that the window is struck by a ball made up a large number of particles arranged so as to form a single concentrated mass, or something similar. These are the facts about the ball's causal influence on the window that made a difference to the window's breaking.

Some further remarks. First, a further constraint must be placed on the abstraction process: it must not go so far as to destroy the *cohesion* of the causal model (that is, of the deductive argument). The urgent reason for this constraint is to prevent abstraction by arbitrary disjunction, for example, to prevent the model for the window's collision-induced breaking being made more abstract by disjoining it with a model for the process in which a window is broken by a sonic boom (to use an example from Hall's contribution). More generally, the motivation for the cohesion requirement is to ensure that the model, though abstract, describes a single kind of causal mechanism or process, rather than presenting a list or disjunction of such processes. The definition of cohesion, then, is in effect a criterion for individuating causal processes. As such, it is not a distinctive ingredient of the kairetic account but rather a necessary component of any causal theory of explanation, since any theory that identifies scientific explanations with causal mechanisms must provide some criterion for telling mechanisms apart. Weatherson's contribution to this symposium makes this point admirably clear; indeed, since cohesion is discussed at considerable and rewarding length by all three contributors and also, naturally, in my replies, I will not treat it further—or even define it—here.

Second, when the abstraction process is done, you have something more than a list of difference-makers. The difference-makers will be arranged, like the causal influences from which they were derived, in the form of a causal model. I call such a model—a model containing only difference-makers—a *standalone explanation* for the explanandum, and I propose that the aim of explanatory inquiry in science is to construct standalone explanations. An event will have many standalone explanations, each taking up the causal story at a different point or points in time and space—some closer to the explanandum, some further away. All standalone explanations are good, and convey understanding of the phenomenon that is explained, but some are better than others. These questions of relative explanatory power are investigated further in chapter 4 of *Depth*.

Third, many readers will have noticed that the various difference-making criteria considered in the discussion above also appear in the literature on causality. (The kairetic criterion itself corresponds loosely to Mackie's (1974) INUS criterion for causation.) What is going on?

My view is that "was a cause of" locutions, the analysis of which is the aim of much philosophical work on causation, are a species of explanatory difference-making claim. More exactly, to say of two events *c* and *e* that *c* was a cause of *e* is to say that the occurrence of *c* is a part of the explanation of *e*. If I am right then this genre in the causation literature is in fact a part of the philosophy of explanation. But only a part: the causation literature tends to confine itself to event causation, where the counterfactual and probabilistic criteria work well, and so to ignore the kind of causal difference-making that accounts for sodium's reactivity and other such things.

Another notable difference between the story told by proponents of, say, the counterfactual approach to causation and the kairetic account of explanation is that the kairetic account sees causal difference-making relations, hence the relations identified by "was a cause of" claims, as constructed from a fundamental level relation of causal influence, while counterfactualists tend to regard the "was a cause of" relation as something not to be reduced to a more primitive causal relation, but directly to non-causal (i.e., counterfactual) facts. Whereas in my view, there are two kinds of (token) causal relation—a more basic relation of causal influence and a derived relation of causal difference-making—the causation literature for the most part operates as if there is a single causal relation, identical to that asserted by "was a cause of" claims. (A notable exception is Lewis's (2000) final theory of causation, which like the kairetic approach reduces the "was a cause of" relation to a more basic influence relation that is very close to my own causal influence relation.)

If "was a cause of" claims are at bottom explanatory claims, then a theory of explanation ought to be capable of resolving many of the questions that have been posed in preceding decades about causation: How to handle tricky preemption cases? How can absences or "omissions" be causes? (That is, how can absences appear in the first position of "was a cause of" claims?) What is the causal status of prevention? Is "was a cause of" transitive? Chapter 6 of *Depth* tackles each of these questions, arguing for example that all well known arguments against transitivity have subtle flaws.

3. The Uses of Causal Difference-Making

A causal approach to explanation seems irresistible in the light of the flagpole and other counterexamples to Hempel's purely deductive approach. Yet for all that the explanation of the flagpole's shadow obediently follows the direction of causation, countless other scientific explanations seem quite deaf to the causal order. Whether from insouciance or ignorance, they disregard or distort important causal facts, or talk the language of stochasticity where causal reality is strictly deterministic.

An appreciation for the extreme explanatory abstraction that so often accompanies the search for causal difference-makers makes sense of much of this behavior without abandoning the fundamental thesis that scientific explanations do nothing but describe causal processes—that explanations are solely concerned with telling causal stories.

Let me give you three examples, each developed at length in *Depth*.

Equilibrium Explanation Elliott Sober (1983) has memorably argued that equilibrium explanations are not causal, because they do not identify the causal trajectory that a system takes from its initial state to the state that is to be explained. An evolutionary biologist, for example, might explain the roughly one-to-one sex ratio in a certain species by showing that any population with a different ratio will be pulled by natural selection toward the even ratio, thus that the even ratio is a unique evolutionarily stable equilibrium. In explaining the population's even sex ratio, then, they do not specify the ratio at any earlier time, or the path by which the earlier ratio (if non-even) became an even ratio. Apparently, they describe no causal process at all.

A number of commentators have noted that the equilibrium explanation is not bereft of causal content: in showing that a certain state is an equilibrium, you describe the causal dynamics that make it so. *Depth* goes further and shows that the equilibrium model is precisely what you will obtain if you abstract from a more specific model in accordance with the kairetic criterion for difference-making. Take a causal model that specifies the population's sex ratio at some earlier time, then, and identifies the trajectory taken to the ratio's current even state. According to the kairetic account, this additional information—everything over and above what is described by the equilibrium model—is causally irrelevant, because it makes no difference, to the explanandum. Whereas Sober supposes that without explicit information about trajectories, the model cannot be causal, I maintain that the canons of causal explanation require that such information be removed. The equilibrium explanation is the one explanation of the even ratio that is causally good and true.

In his contribution to this symposium, Weatherson objects that many legitimate equilibrium explanations fail to satisfy the kairetic account's cohesion requirement. I will deal with this complaint in my reply.

Idealization Scientific explanations almost always get the causal facts wrong, and deliberately so. Economic explanations assume that decisions are made by actors who are perfectly well informed or, in a certain sense, perfectly rational. Biological explanations assume that populations undergoing natural selection are infinitely large, or that genes assort independently. Explanations in the physical sciences might assume zero friction, zero perturbation, zero exogenous noise.

One way to account for these miserable errors is as a reflection of the fallen state of humankind: our weak intellects cannot bear reality in all of its radiant, complex profusion, so we water down our explanatory models with simplifying falsehoods to achieve something inferior but more cognitively palatable. On this view, an idealized model is explanatorily worse than what you might call its veridical counterpart—a corresponding model in which the idealized model's distortions have been rectified—but is preferred on practical grounds.

In *Depth* I argue, by contrast, that the idealized model offers a better explanation than its veridical counterpart. How can it be, if there is nothing to explanation but telling a true causal story?

Two claims defended in *Depth* supply the explanation. First, the things that are distorted or omitted in idealizing explanations are not difference-makers. Getting them right, as an idealized model's veridical counterpart does, therefore fails to provide additional explanatory information.

Second, an idealized model's distortions, when interpreted correctly, contain useful explanatory information: the distortions tell you that the distorted factors, or forces or tendencies attributable to the distorted factors, are not difference-makers. When the explanation of Boyle's law, for example, falsely assumes that gas molecules do not collide, it is telling you that the collisions in a real gas make no difference to the fact that it conforms (approximately) to Boyle's law. Likewise, when biologists assume that populations are infinite, and thus that there is no genetic drift, they are telling you that drift makes no difference to their explanandum (often, the fixation of some character in the population). And when economists assume perfect knowledge, they are saying that, for the purposes of the salient explanation, ignorance is not a difference-maker.

An idealizing explanation does not incorrectly represent certain fictional factors as making a difference to the explanandum, then, but rather correctly represents certain actual factors as not making a difference. For this reason, such an explanation is superior to its veridical counterpart, which by correctly describing non-difference-makers, falsely implies that they make a difference.

There is a certain code to idealization, I suggest: non-difference-makers are distorted by replacing their true descriptions with descriptions of some corresponding default state. Often this is zero, as when a small, non-difference-making force is represented as not existing at all. Sometimes it is an infinite value, as with the evolutionary biologists' populations. Or it might be something more complex, such as probabilistic independence (in genetics), perfect information or rationality (in economics), or the kind of perfect ignorance that is represented by a uniform probability distribution (an example that I owe to Weatherson's contribution to this symposium).

Scientists know the code; thus, they know the right way to read idealized models. Properly interpreted, such models have nothing but true things to say about the relevant causal process.

Probabilistic Explanation in Deterministic Systems It is often convenient to use stochastic elements in models of systems that are known, or at least suspected, to have a deterministic dynamics. Social scientists and many others do it when they introduce an "error term" into their causal models.

In some special cases, these probabilities are not merely useful but explanatory. Classical statistical mechanics provides one example: even if the underlying physics is deterministic, the statistical mechanical probabilities seem to give genuine insight into the reason why, say, the temperature of two adjacent objects tends to equilibrate. Another example is evolutionary biology, where probabilities are used to explain the random extinction in small populations of perfectly good alleles through genetic drift.

The explananda in these examples are produced by deterministic causal processes. Surely, then, a causal account of explanation should direct us to explain such phenomena by telling a deterministic causal story. How can it be helpful to refer to probabilities that play no role in fundamental causal reality?

As always, there is a practical account: feeble minds, limited resources, satisficing, and so on. As usual, I am interested in something more principled.

Depth argues that, as in the case of equilibrium explanation, the kairetic criterion for difference-making *forces* us to go stochastic. Start with a deterministic characterization of the process, then, and abstract away everything you can; what you are left with is a model that appeals to probability-like entities. I say "probability-like" for two reasons. First, in *Depth* I take no position as to whether these are "real" probabilities or not; that's one for the metaphysicians (which might be me on a different day, as in Strevens (2011)). Second, the probability-like descriptions that remain at the end of abstraction may not be genuinely indeterministic: they may use the mathematical apparatus of probability theory, but without allowing any genuine latitude in outcomes. What you might have, for example, is a set of claims about probabilities and stochastic independence that, when fully analyzed, entail rather than merely probabilifying the fact that a certain outcome occurs with a certain frequency (this latter state of affairs being the explanandum).

The previous comment notwithstanding, some optimal explanatory models of deterministically produced outcomes are genuinely non-deterministic; to understand why, however, requires a notion that will be introduced only in the next section.

* * *

A concluding bonbon: *Depth* proves, in a brief detour through political philosophy, that even in an election won by a sizable majority, your vote does make a difference (p. 200). (For similar reasons, your flight to the APA this year makes a difference to global warming, which perhaps provides additional reason to exercise your vote.)

4. Beyond Causal Difference-Making

Laws and Entanglement On a causal approach to explanation, to explain the law *In conditions Z, all Fs are G*, you should specify the causal mechanism by which F-ness, in conditions Z, produces G-ness. *Depth* endorses this view, which makes life easy because the relevant notion of a causal mechanism is the one that has been developed already for event explanation.

There is one complication: in almost any high-level scientific law or generalization, the antecedent property F plays no causal role in the mechanism. All ravens are black, but ravenhood itself (which is a partly historical property) is not in on the causal action. Pre-1977 Ford Pintos have a tendency to explode when struck from behind, but their "Pinto-hood" is not the cause. In both cases, what is doing the causing is rather some physical properties of, respectively, ravens and pre-1977 Pintos.[2]

To explain why all ravens are black, then (or at least why, in normal conditions, most ravens are black), it had better not be necessary to supply a causal mechanism in which ravenhood together with normal conditions causes blackness, because there is no such mechanism, which would mean that raven blackness is inexplicable.

A possible alternative: let P denote the physical properties that cause ravens to be black. Can raven blackness by explained by exhibiting the causal mechanism by which P, in normal conditions, causes blackness? Not quite. A model of this mechanism would explain why things with P are black, but that is a different explanandum: possession of P may be neither necessary nor sufficient for ravenhood. (Perhaps carrion crows as well as ravens have P. Perhaps ravens could have evolved so as to lack P.) The mechanism is a part of the explanation for raven blackness, but to complete the story, another fact must be added to the mix: that most ravens have P.

This is what I call a *basing generalization*. It need not be a causal claim;

2. Ingenious philosophers can no doubt find ways to dispute the causal neutrality of ravenhood and Pinto-hood, but as you will see, I have a better solution.

indeed, in the case at hand, it is not causal, because ravenhood no more causes possession of P than it causes blackness. To deal with the explanation of laws and other generalizations in the high-level sciences, then, we need to supplement a causal account of explanation with an account of basing generalizations that allows them to play an explanatory role despite their occasional non-causality.

Such an account forms a large part of chapter 7 of *Depth*. Let me share two of my conclusions. First, to play its explanatory role, the property F (in the example, ravenhood) connected to a causally relevant property P by a basing generalization must bear toward P a two-way modal relation I call *entanglement*. What F's entanglement with P amounts to, very roughly, is: (a) things with F tend to have P, in the sense that in the actual and most nearby possible worlds, most things with F have P, and (b) if a thing loses its F-ness, it will tend to lose its P-hood. (It is (a) that accounts for the robustness of raven blackness.)

Second, ravenhood is relevant to the explanation of blackness in virtue of the non-causal relation of entanglement. Not all explanatory relevance, then, is causal, though this does not imply that some explanation is non-causal, because basing generalizations enter into explanations only at the front end of causal mechanisms.

Note that clause (a) of the definition of entanglement, with its double "most", introduces an element of non-determinism to explanation that is quite orthogonal to the determinism or otherwise of explanatory causal processes. It is through entanglement that genuinely indeterministic explanatory models can explain deterministically produced outcomes, as promised in the previous section.

Explanation Outside Science I classify the kairetic account of explanation as a "two-factor" account because it has the following two parts:

1. It posits a fundamental dependence structure between various elements

of reality, namely, the web of causal influence.

2. It supplies a criterion for picking out aspects of the web of dependence that make a difference to, and so help to explain, any given high-level state of affairs.

The two-part structure is feasible because the difference-making criterion is somewhat independent of the nature of the fundamental dependence relation; this is what allows me to offer an account of causal explanation without resolving the question of the nature of fundamental causality—that is, the nature of causal influence.

Let me finish by making two suggestions about explanation in areas outside of science, such as mathematics and morality. First, I speculate that explanation in these domains may have the same two-part structure: a web of dependence and a difference-making criterion. In the case of mathematical explanation, for example, there would be a dependence structure reflecting which mathematical facts were more fundamental, and a difference-making criterion determining which aspects of the fundamental facts make a difference to any particular high-level mathematical state of affairs.

Second, I suggest that the difference-making criterion might be the same (except perhaps for the details of cohesion) in every explanatory realm. The kairetic difference-making criterion developed for causal explanation in *Depth* might, in other words, be precisely the criterion that determines difference-making in mathematical explanation, in moral explanation (dictating what aspects of the fundamental moral laws make a difference to the wrongness of a particular act), in aesthetic explanation, and elsewhere. In the kairetic account, then, lies the germ of a general theory of explanation.

References

Baumgartner, M. (2008). Regularity theories reassessed. *Philosophia* 36:327–354.

Dowe, P. (2000). *Physical Causation*. Cambridge University Press, Cambridge.

Garfinkel, A. (1981). *Forms of Explanation*. Yale University Press, New Haven, CT.

Hempel, C. G. and P. Oppenheim. (1948). Studies in the logic of explanation. *Philosophy of Science* 15:135–175.

Lewis, D. (1973). Causation. *Journal of Philosophy* 70:556–567.

———. (2000). Causation as influence. *Journal of Philosophy* 97:182–197.

Mackie, J. (1974). *The Cement of the Universe*. Oxford University Press, Oxford.

Salmon, W. C. (1984). *Explanation and the Causal Structure of the World*. Princeton University Press, Princeton, NJ.

———. (1997). Causality and explanation: A reply to two critiques. *Philosophy of Science* 64:461–477.

Sober, E. (1983). Equilibrium explanation. *Philosophical Studies* 43:201–210.

Strevens, M. (2011). Probability out of determinism. In C. Beisbart and S. Hartmann (eds.), *Probabilities In Physics*. Oxford University Press, Oxford.

Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, Oxford.