Explanatory Conditionals

Holger Andreas*

This article aims to complement causal model approaches to causal explanation by James Woodward, Joseph Halpern and Judea Pearl, and Michael Strevens. It centers on a strengthened Ramsey Test of conditionals: $\alpha \gg \gamma$ if and only if, after suspending judgment about α and γ, an agent can infer γ from the supposition of α (in the context of further beliefs in the background). Holger Andreas and Mario Günther used this conditional as the starting point of an analysis of 'because' in natural language. I refine this analysis so as to yield a fully fledged account of (deterministic) causal explanation.

1. Introduction. This article aims to complement causal model approaches to causal explanation by Woodward (2003), Halpern and Pearl (2005b), and Strevens (2008). It does so by carrying on a conditional analysis of the word 'because' in natural language by Andreas and Günther (2018). This analysis centers on a strengthened Ramsey Test of conditionals:

> $\alpha \gg \gamma$ if and only if (iff), after suspending judgment about α and γ, an agent can infer γ from the supposition of α (in the context of further beliefs in the background).

Using this conditional, we can give a logical analysis of because:

$$\textit{Because } \alpha, \quad \gamma(\text{relative to } K) \quad \text{iff} \quad \alpha \gg \gamma \in K \quad \text{and} \quad \alpha, \gamma \in K,$$

where K designates the belief set of the agent. Andreas and Günther (2019) refined this analysis by further conditions, which results in a broadly inferential account of deterministic, actual causation. This article builds on this account so as to yield an analysis of deterministic causal explanations. The logical foundations of the belief changes that define the conditional $\gg$ are

*To contact the author, please write to: Department of Economics, Philosophy and Political Science, University of British Columbia Okanagan, 1147 Research Road, Kelowna, BC V1V 1V7, Canada; e-mail: holger.andreas@ubc.ca.

Philosophy of Science, 86 (December 2019) pp. 993–1004. 0031-8248/2019/8605-0015$10.00
Copyright 2019 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/705447 Published online by Cambridge University Press

explicated using Alchourrón, Gärdenfors, and Makinson (AGM) style belief revision theory (Gärdenfors 1988).

Why do we think that causal model approaches to causal explanation are incomplete? Woodward (2003) and Halpern and Pearl (2005b) offer the most prominent elaborations of a causal model approach in contemporary philosophy of science. Halpern and Pearl (2005a) have devised a precise semantics of causal models that centers on structural equations. Such an equation represents causal dependencies between variables in a causal model:

$$X := \phi(Y_1, \ldots, Y_n).$$

In this schema of a structural equation, the variable X causally depends on the variables $Y_1, \ldots, Y_n$. It proved highly useful to represent the causal dependencies of a causal model by a causal graph. Woodward's (2003) account of causal explanations heavily relies on such graphs.

Halpern and Pearl (2005a, 847) take the structural equations of a causal model to represent laws or mechanisms in the world. However, they do not address the problem of distinguishing between causal and noncausal dependencies of laws of nature. There is no further explanation of what it is for a variable to causally depend directly on certain other variables. The definition of causation by Halpern and Pearl (2005a) merely defines complex causal relations in terms of elementary causal dependencies, just as truth-conditional semantics defines the semantic values of complex sentences in terms of a truth-value assignment to the atomic formulas. And the corresponding account of causal explanation in Halpern and Pearl (2005b) inherits the reliance on elementary causal dependencies from the analysis of causation.

Woodward (2003) explains the notion of a direct cause in terms of interventions, but the notion of an intervention is always relative to a causal graph, so that some knowledge about elementary causal dependencies must be antecedently known. An account of causal explanation in terms of elementary causal dependencies is certainly valuable and insightful, but it is not the full story. At least, we should not give up on a more comprehensive account of causation and causal explanation too quickly (cf. Spohn 2006; Paul and Hall 2013).

The kairetic account of explanation by Strevens (2004, 2008) makes essential use of causal models as well, but it works with a more liberal notion of such a model. In this account, a set of propositions entail an explanandum E in a causal model only if this entailment corresponds to a real causal process by which E is causally produced (Strevens 2004, 165). Causal models are assumed to be founded in physical facts about causal influence, and it is assumed that these facts can be read off the true theory of everything (165).

The kairetic account is conceptually incomplete in a manner akin to the approaches by Woodward (2003) and Halpern and Pearl (2005b). This account

https://doi.org/10.1086/705447 Published online by Cambridge University Press

leaves open how we can discriminate between causal relations of logical entailment and noncausal ones in a true theory about the world, be it complete or incomplete. Relatedly, the account comes with a number of placeholders that remain unspecified. For example, causal models are assumed to define a relation of logical entailment that represents actual causal relations. What is the language of such a causal model? How is the mathematics underlying some causal processes represented? What are the distinctive properties of causal relations of logical entailment? In what follows, we make an attempt at answering these questions, focusing in particular on a characterization of logical entailment with a causal meaning. For this characterization, we define an explanatory conditional $\gg$ but impose also nonlogical conditions on the explanans and the explanandum.

Our final account of logical entailment with a causal meaning takes as input a nonmodal, first-order representation of our theories and beliefs about the world. Modality and explanatory conditionals come into play through a refined and strengthened variant of the Ramsey Test. The final analysis is thus relative to an epistemic state that represents certain theories and beliefs about the world. This yields a smooth epistemology of causal explanations insofar as we seem to know what beliefs and theories we have about the world. The analysis is thus very much in the spirit of related accounts of explanation and causation by Gärdenfors (1988, chaps. 8 and 9) and Spohn (2006). Note, however, that we are not committed to the view that causation and explanation are epistemic concepts. If we are given an epistemic state that contains the true theory of everything, our analysis explains how we can read the causal structure off such a theory or a portion thereof.

2. Belief Changes and the Ramsey Test

2.1. Belief Changes: Basic Ideas. AGM-style belief revision theory provides us with a precise semantics of belief changes for the Ramsey Test. Let us therefore make ourselves familiar with the basic ideas of this theory. In the AGM framework, one distinguishes three types of belief change of a belief set K by a formula α:

1. Expansions $K + \alpha$
2. Revisions $K * \alpha$
3. Contractions $K \div \alpha$.

An expansion of K by α consists in the addition of a new belief α to the belief system K. This operation is not constrained by any considerations as to whether the new epistemic input α is consistent with the set K of present beliefs. Hence, none of the present beliefs is retracted by an expansion. The expanded belief set is designated by $K + \alpha$.

https://doi.org/10.1086/705447 Published online by Cambridge University Press

A revision of K by α, by contrast, can be described as the *consistent integration* of a new epistemic input α into a belief system K. If α is consistent with K, it holds that $K + \alpha = K * \alpha$; that is, the revision by α is equivalent with the expansion by α. If, however, α is not consistent with K, some of the present beliefs are to be retracted, as a consequence of adopting the new epistemic input. The revised system of beliefs is designated by $K * \alpha$.

A contraction of K by α, finally, consists in retracting a certain formula α from the currently accepted system of beliefs. This operation will be used to define the *suspension of judgment about* α in our strengthened version of the Ramsey Test. The belief set after the retraction of α is designated by $K \div \alpha$.

In some contexts, it is helpful to distinguish between the belief system K and the epistemic state S that underlies it. Henceforth, we make this distinction and write $K(S)$ for the belief system K of the epistemic state S. Further, we assume that epistemic states are represented by *belief bases*. In symbols, $S = H$. A belief base H is a set of formulas that represent the explicit beliefs of an agent.

Note, finally, that we can define revisions in terms of contractions and expansions (Levi identity):

$$K(S) * \alpha = (K(S) \div \neg\alpha) + \alpha.$$

Once we have retracted α, we obtain a belief set $K(S')$ that is consistent with α. Hence, we have $K(S') * \alpha = K(S') + \alpha$. Such are the basic ideas about belief changes that will be used in our strengthened Ramsey Test and the subsequent analysis of explanation.[1]

2.2. The Ramsey Test. The Ramsey Test is an epistemic approach to conditionals. Its core idea has been expressed most clearly by Stalnaker (1968, 102): "First, add the antecedent (hypothetically) to your stock of beliefs; second, make whatever adjustments are required to maintain consistency (without modifying the hypothetical belief in the antecedent), finally, consider whether or not the consequent is then true." It was then Gärdenfors (1988) who translated this test into the language of belief changes and who insisted more forcefully than Stalnaker (1968) on an epistemic understanding of conditionals. Using the AGM framework, he was able to explicitly define a semantics of conditionals in terms of belief changes:

$$\alpha > \gamma \in K(S) \text{ iff } \gamma \in K(S) * \alpha,$$

where $>$ designates the conditional connective. Recall that $K(S) * \alpha$ designates the revision of the beliefs of an epistemic state S with the formula α. So the Ramsey Test defines that a conditional $\alpha > \gamma$ is to be accepted in a

1. For lack of space, we go no further into details of the definition of contractions and revisions. For this, the reader is referred to Gärdenfors (1988), Hansson (1999), and Andreas and Günther (2018).

https://doi.org/10.1086/705447 Published online by Cambridge University Press

belief system $K(S)$ iff it is true of $K(S)$ that the consequent γ is in $K(S)$ when revised by the antecedent α.

2.3. Strengthening the Ramsey Test. Inspired by the work of Rott (1986) on the logical analysis of 'because', Andreas and Günther (2018) defined a conditional $\gg$ with the following intuitive meaning: $\alpha \gg \gamma$ iff, after suspending any beliefs in $K(S)$ as to whether α and γ are true or false, it holds that $\gamma \in K(S) * \alpha$. The evaluation of $\alpha \gg \gamma$, thus, consists of two steps: (i) contracting $K(S)$ in such a manner that we become indeterminate about α and γ and (ii) testing whether γ is in $K(S) * \alpha$. In more formal terms:

Definition 1: Belief function $B(\alpha)$. Let T be some arbitrary classical tautology and α a formula.

$$B(\alpha) = \begin{cases} \alpha & \text{if } \alpha \in K(S) \\ \neg\alpha & \text{if } \neg\alpha \in K(S) \\ \neg T & \text{otherwise.} \end{cases}$$

$$\text{(SRT)}\quad \alpha \gg \gamma \in K(S) \text{ iff } \alpha \in (K(S) \div B(\alpha) \vee B(\gamma)) * \alpha.$$

Contracting $K(S)$ by $B(\alpha) \vee B(\gamma)$ results in a belief system $K(S')$ that does not contain α or $\neg\alpha$ or γ or $\neg\gamma$, provided that α and γ are contingent. It thus represents the operation of suspending judgment about the truth and falsity of α, γ. Hence, $\alpha \gg \gamma$ iff $\gamma \in K(S') * \alpha$.

Using $\neg\alpha \notin K(S) \div B(\alpha) \vee B(\gamma)$ and the Levi identity, we obtain

$$\alpha \gg \gamma \text{ iff } (K(S) \div B(\alpha) \vee B(\gamma)), \alpha \vdash \gamma,$$

where $\vdash$ stands for the relation of logical consequence in classical logic. Thus, $\alpha \gg \gamma$ means that the consequent γ is inferable from the antecedent α, together with other explicit beliefs, after judgment has been suspended about γ and α.

2.4. Ramsey Test Explanations. Having strengthened the Ramsey Test, we are in a position to give a preliminary account of explanation:

Definition 2: Explanation (preliminary account). Let S be an epistemic state that is represented by a belief base. The set A of antecedent conditions and the set G of generalizations explain the fact F—relative to S—iff

(E1) For all $\alpha \in A$, all $\gamma \in G$, and all $\beta \in F$: $\alpha, \gamma, \beta \in K(S)$.
(E2) For all nonempty $A' \subseteq A$, $\bigwedge A' \gg \bigwedge F \in K(S)$.

https://doi.org/10.1086/705447 Published online by Cambridge University Press

Here, $\bigwedge C$ is shorthand for $\bigwedge_{\phi \in C} \phi$ and, hence, designates a conjunction of all members of a set C. We do not merely require that $\bigwedge A \gg \bigwedge F$ to ensure that all propositions of A are relevant for the explanandum; F is a set because the description of the explanandum may be complex.

This preliminary analysis bears substantial commonalities with Strevens (2008) and accounts in the logical empiricist tradition. All of these accounts require that the explanandum be entailed by a set of antecedent conditions, together with a set of generalizations or laws. In line with Strevens (2008), we are aiming to capture the intuition that the explanandum is *causally produced* by the antecedent conditions of the explanation. For this to be achieved, we need to impose further constraints on the Ramsey Test, the explanans, and explanandum.

3. Causation

3.1. Time. For C to be a cause of E, it must hold that

(C1) $C \gg E \in K(S)$.

This condition allows us to capture a large range of causal relations. It is too liberal, however. If a theory T is deterministic and time symmetric, we are not only able to infer the effect from the cause, but also the other way around. Classical mechanics is such a theory.

Now, the idea of production seems to imply a temporal asymmetry between the producing event and the effect: the cause must precede its effect. Hence,

(C2) $t(C) < t(E)$,

where $t(C)$ is a function that yields the time at which the event C occurs. This condition expresses an old Humean dictum, which is also central to Spohn's (2006) ranking-theoretic analysis. It helps to not mistake effects for causes in deterministic, time-symmetric theories. If A, B, or A and B, are temporally extended events, we take $t(A) < t(B)$ to mean that A comes into being before B, while A and B may well overlap. In the next section, we see that condition C2 is too restrictive and so needs to be liberalized—not only because of ideas about backward causation.

3.2. Levels of Theoreticity. The temporal asymmetry between cause and effect appears to hold consistently in everyday contexts. It does not always hold in scientific contexts. For example, we say that forces produce accelerations and that electromagnetic fields produce forces upon charged particles, without there being a temporal delay between, respectively, forces and accelerations, and electromagnetic fields and electromagnetic forces.

https://doi.org/10.1086/705447 Published online by Cambridge University Press

Is this accidental? If not, what is the distinctive relation that allows us to recognize a causal direction in the latter cases?

We conjecture that an ordering of *theoreticity* plays an important role here. Forces are higher up in this ordering than accelerations because Newtonian mechanics determines the meaning of force and mass on the basis of the concepts of space and time but not vice versa.[2] The notion of an ordering of theoreticity has been made precise recently by Schurz (2014). The core idea of this definition is that a concept C_1 is higher up in this ordering than a concept C_2 iff there is a chain of theories that define measurement procedures that lead from measured values of C_2 to measured values of C_1. The ordering of theoreticity thus defined may well be partial. For lack of space, we cannot go further into the details.

Note that there is also a philosophical motivation for recognizing a connection between theoreticity and causation. A major motivation for introducing theoretical concepts is to give a unified account of certain phenomena. In light of unificationist ideas about explanation by Kitcher (1989), this implies that theoretical concepts are—for conceptual reasons—supposed to help devise explanations. If we then further assume a close connection between causation and explanation, we can infer that theoretical concepts are supposed to provide causal explanations.

Using our conjecture concerning theoreticity and causation, we are able to liberalize condition C2 as follows:

(C2*) $t(C) < t(E)$ or $c(E) <_t c(C)$.

In this notation, $c(A)$ gives us the set of concepts in terms of which the sentence A is expressed. And $c(E) <_t c(C)$ is intended to mean that there is at least one concept in $c(C)$ that is higher up in the theoreticity ordering than all the concepts of $c(E)$, while there is no concept in $c(C)$ that is below a concept in $c(E)$ in the sense of this ordering.

3.3. Joint Effects. Joint effects of a common cause can pose a problem for an inferential approach to causation and explanation. Take the neuron diagram shown in figure 1: C fires and, thereby, sends signals to D and B so that D and B are excited. E is excited in the course of receiving a signal from D. Intuitively, the excitation of B is not a cause of the excitation of E. However, the excitation of B strongly conditionally implies—in the sense of SRT—the excitation of E.

Counterfactual approaches solve the problem by excluding *backtracking* counterfactuals (Paul and Hall 2013, 71–72). Andreas and Günther (2019) adopt a similar strategy: the counterintuitive result about joint effects is

2. This is so on all formal semantics of theoretical concept.

https://doi.org/10.1086/705447 Published online by Cambridge University Press

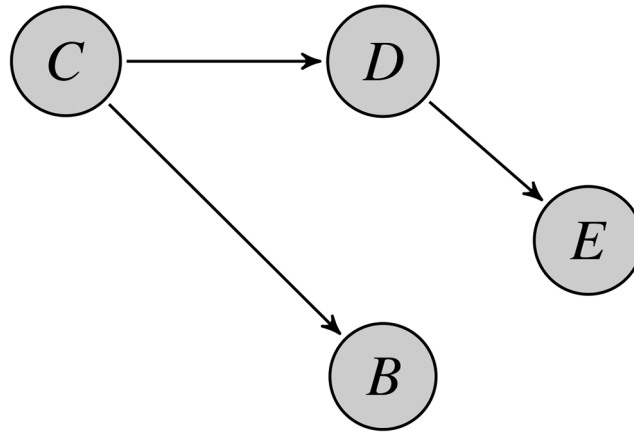


Figure 1. Neuron diagram. (Source: Paul and Hall 2013, 71.)

avoided if all inferences that lead from the presumed cause C to the putative effect E are nonbacktracking. They make this idea precise in proof-theoretic terms using the notion of a *forward-directed proof*. Such a proof conforms to the temporal order of events in the following sense:

> **Definition 3:** $H \vdash_F C$. Let H be a set of formulas and C be a formula. Only literals and conjunctions of literals are taken to assert the occurrence of an event. We say there is a *forward-directed natural deduction proof* of C from H—in symbols $H \vdash_F C$—iff there is a natural deduction proof of C from H such that (i) for all inferential steps P/I (of the main proof and any subproof), if I asserts the occurrence of an event, then this event does not precede any event that is asserted by a premise in P or by a premise in a subproof that is a member of P, and (ii) the assumption of any subproof is consistent with H.

Using this notion of a forward-directed proof, we can impose a temporal constraint on our Ramsey Test:

(SRT_F) $A \gg_F C \in K(S)$ iff $(K(S) \div B(A) \vee B(C)), A \vdash_F C$.

That is, C is a forward-directed strong conditional implication of A iff there is a forward-directed proof of C from A and the explicit beliefs of the epistemic state S, after suspending judgment on the antecedent A and the consequent C. Now, we require there to be a forward-directed proof of the putative effect from the presumed cause:

(C1*) $C \gg_F E \in K(S)$.

https://doi.org/10.1086/705447 Published online by Cambridge University Press

This condition solves the problem of joint effects. For, the inference from the firing of B to the firing of C is backward directed. Hence, there is no forward-directed proof of E from A. Note that a forward-directed proof merely excludes backward-directed inferences. Such a proof may still involve inferences where the events asserted in the premises and the conclusion are simultaneous.

3.4. Spurious Causation. Unfortunately, there is another problem with the causal scenario of figure 1. As there is a stable correlation between the firing of B and E, it is reasonable to have the implication $B \rightarrow E$ as a generalization in the belief base. Hence, $B \gg E$. If the firing of B precedes that of E, the inference is forward directed, and so $B \gg_F E$ holds as well. So we would have to consider the firing of B as a cause of the firing of E, which is counterintuitive. This is the problem of spurious causation.

The paradigm of a spurious cause is the drop of the barometer that does not count as a genuine cause of any storm. Since, however, the correlation between the drop of the barometer and the storm is commonly considered to be probabilistic, this example is not well suited for the current analysis of deterministic causation. Here is a better example: a causal analysis of a thunderstorm should distinguish at least three events: (i) the electrical discharge between a cloud and the ground, (ii) the flash of the lightning, and (iii) the thunder. Physics tells us that the electrical discharge between a cloud and the ground is—via the rapid production of heat within the region of the air where electricity is conducted—the common cause of the flash and the thunder. Physics would deny that the flash is a genuine cause of the thunder. So we have to conclude that the flash is a mere spurious cause of the thunder. But the event of the flash satisfies conditions C1* and C2* with respect to the thunder. So we must further refine our analysis.

Why is the electrical discharge rather than the flash a genuine cause of the thunder that temporally follows? Let us take a closer look at the various inferential paths in this causal scenario. The flash and the electrical discharge of the lightning are alike in that we can infer from their occurrence the occurrence of the thunder. The inferential paths, however, differ from one another. We can infer the thunder from the electrical discharge, in a forward-directed manner, using only generalizations that are fundamental, or nonredundant, *laws of nature*. In essence, it is the laws of electrodynamics, atomic theory, and acoustics that are used in this path. By contrast, there is no forward-directed inferential path from the flash to the thunder such that all generalizations thereof are fundamental laws of nature. The inference from the flash of the lightning to the thunder is either not forward directed (when it goes via the electrical discharge) or not based on fundamental laws of nature (when the thunder is directly inferred from the flash, without going through the common cause of the flash and the thunder). Figure 2 may help distinguish the two inferential paths.

Let us be more precise about the notion of a fundamental, or nonredundant, law of nature. To make our approach as parsimonious as possible and to avoid

https://doi.org/10.1086/705447 Published online by Cambridge University Press

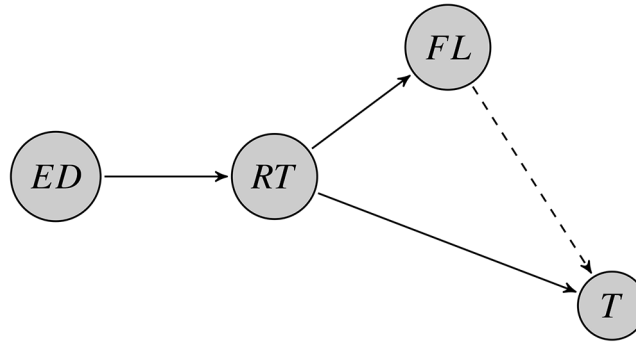


Figure 2. Inferential path. ED = electrical discharge, RT = rise of temperature, FL = flash of lightning, and T = thunder.

conflict with alternative accounts of laws of nature, we prefer the notion of a *nonredundant generalization* over the notion of a law of nature. Yet, we explicate the former notion in terms of a *best system account* of laws of nature. In fact, a very weak understanding of what a best system is suffices to make the crucial distinction:

> **Explanation 1:** Nonredundant generalization. A generalization is nonredundant iff it is a member of a deductive system that does a good job at balancing strength and simplicity and that is not clearly inferior to an alternative system.

This formulation acknowledges that the optimization of strength and simplicity cannot be accomplished in a straightforward manner.[3] However, it is strong enough to see that the generalization asserting a regular connection between a flash at the sky and a thunder is redundant. For, (i) it is uncontroversial that the laws of electrodynamics, atomic theory, acoustics, and so on (i.e., the laws that are used in the inferential path from the flash to the thunder via the electrical discharge), are members of at least some deductive systems that do a good job at maximizing strength and simplicity and that are not clearly inferior to an alternative system. Further, (ii) our generalization concerning flash and thunder can be derived from the laws described in i. More precisely, it can be derived that the generalization holds universally for a system that contains clouds above the ground and an atmosphere of a specific

3. For an accessible summary, sympathetic reconsideration, and refinement of the best system account, the reader is referred to Cohen and Callender (2009). Notably, their account is prepared to face a variety of best systems, as explanation 1 is. It is worth noting, moreover, that Cohen and Callender (2009) are concerned with deductive systematizations of our knowledge of whatever domains.

https://doi.org/10.1086/705447 Published online by Cambridge University Press

composition. Otherwise, there would not be an inferential path from the flash to the thunder via the electrical discharge. (iii) Deductive systems that lack the laws of electrodynamics, atomic theory, acoustics, and so on, are clearly inferior to at least some deductive systems that contain these laws. Conditions i, ii, and iii imply that the critical generalization concerning flash and thunder is redundant in the sense of explanation 1.

In a similar vein, it can be shown that the assertion of a regular connection between the firing of B and E (in the above neuron diagram) is redundant in the sense just explained. In light of these observations, an inferential characterization of nonspurious causes almost falls into place: for any nonspurious cause, there must be a forward-directed inferential path to the effect such that all generalizations of this path are nonredundant (in the sense of explanation 1). Let us express this condition in terms of Ramsey Test conditionals:

> **Definition 4:** $H \vdash_{FN} C$. $H \vdash_{FN} C$ iff there is a forward-directed natural deduction proof of C from H such that all generalizations of this proof are nonredundant.

This gives rise to another conditional:

(SRT_{FN}) $A \gg_{FN} C \in K(S)$ iff $(K(S) \div B(A) \vee B(C)), A \vdash_{FN} C$.

Thus, we obtain

(C1**) $C \gg_{FN} E \in K(S)$.

This condition says that there must be a forward-directed inferential path between cause and effect such that all generalizations of this path are nonredundant.

4. The Final Account. We can now merge C1** and C2* into our preliminary account of causal explanation from section 2.4:

> **Definition 5:** Causal explanation. Let S be an epistemic state that is represented by a belief base. The set A of antecedent conditions and the set G of generalizations causally explain the fact F—relative to S—iff
>
> (E1) For all $\alpha \in A$, all $\gamma \in G$, and all $\beta \in F$, $\alpha, \gamma, \beta \in K(S)$.
> (E2) For all nonempty $A' \subseteq A$, $\wedge A' \gg_{FN} \wedge F \in K(S)$.
> (E3) For all $\alpha \in A$ and all $\beta \in F$, $t(\alpha) < t(\beta)$ or $c(\alpha) <_t c(\beta)$.
> (E4) For any $\gamma \in G$, E2 fails to hold for $K(S) \div \gamma$.

If the sets A, G, and F satisfy these conditions, we say that $A \cup G$ stands to F in the relation of logical entailment with a causal meaning.

https://doi.org/10.1086/705447 Published online by Cambridge University Press

REFERENCES

Andreas, H., and M. Günther. 2018. "On the Ramsey Test Analysis of 'Because.'" *Erkenntnis*. https://doi.org/10.1007/s10670-018-0006-8.
———. 2019. "Causation in Terms of Production." *Philosophical Studies*. https://doi.org/10.1007/s11098-019-01275-3.
Cohen, J., and C. Callender. 2009. "A Better Best System Account of Lawhood." *Philosophical Studies* 145 (1): 1–34.
Gärdenfors, P. 1988. *Knowledge in Flux*. Cambridge, MA: MIT Press.
Halpern, J. Y., and J. Pearl. 2005a. "Causes and Explanations: A Structural-Model Approach." Pt. 1, "Causes." *British Journal for the Philosophy of Science* 56 (4): 843–87.
———. 2005b. "Causes and Explanations: A Structural-Model Approach." Pt. 2, "Explanations." *British Journal for the Philosophy of Science* 56 (4): 889–911.
Hansson, S. O. 1999. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Dordrecht: Kluwer.
Kitcher, P. 1989. "Explanatory Unification and the Causal Structure of the World." In *Scientific Explanation*, ed. P. Kitcher and W. Salmon, 410–505. Minneapolis: University of Minnesota Press.
Paul, L. A., and N. Hall. 2013. *Causation: A User's Guide*. Oxford: Oxford University Press.
Rott, H. 1986. "Ifs, Though, and Because." *Erkenntnis* 25 (3): 345–70.
Schurz, G. 2014. "Criteria of Theoreticity: Bridging Statement and Non-statement View." *Erkenntnis* 79 (S8): 1–25.
Spohn, W. 2006. "Causation: An alternative." *British Journal for the Philosophy of Science* 57 (1): 93–119.
Stalnaker, R. 1968. "A Theory of Conditionals." In *Studies in Logical Theory*, ed. N. Rescher, 98–112. American Philosophical Quarterly Monograph Series 2. Oxford: Blackwell.
Strevens, M. 2004. "The Causal and Unification Approaches to Explanation Unified—Causally." *Noûs* 38 (1): 154–76.
———. 2008. *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.
Woodward, J. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

https://doi.org/10.1086/705447 Published online by Cambridge University Press