‘Emeralds are expensive because they are rare’: Plausibility of Property Explanations

Daniel Heussen (D.Heussen@city.ac.uk)
James A. Hampton (J.A.Hampton@city.ac.uk)
Department of Psychology, Northampton Square
London, EC1V 0HB UK

Abstract

Research on explanation has primarily focused on event explanations, overlooking how we explain properties. A qualitative study revealed that when people explained a property of an entity, they regularly referred to another property to provide an explanation. Why axes are dangerous was explained by their property of being sharp. The present study looked at what affects the relative plausibility of such explanations. A set of 224 explanations of the form ’x has p because it has q’ were judged for plausibility. Measures of counterfactual relations between the two properties (i.e. likelihood of having p without q), co-occurrence and mutability of each property, as well as a measure of conceptual coherence based on network diagrams (Sloman, Love, Ahn, 1998) were used in a regression analysis to predict plausibility. Conceptual coherence followed by counterfactuals were the strongest predictors of plausibility in a model explaining almost 56% of the variance in plausibility of property explanations.

Introduction

Explanations come in a range of different forms and guises. They answer why-questions about people’s behavior, provide accounts of physical processes and shed light on the occurrence of events. The diversity of explanation has led some to suggest that explanations might not constitute a uniform phenomenon (Keil & Wilson, 2000).

Despite this diversity of explanations, theoretical and empirical research on explanation has focused on events, i.e. explanations of why a particular event occurred. For instance, Hempel’s deductive nomological model was only intended to cover singular events as remarked in a footnote in Hempel and Oppenheim’s original paper:

> “Our analysis will be restricted to the explanation of particular events, i.e. to the case where the explanandum, E, is a singular sentence.” (Hempel & Oppenheim, 1948, p.159)

Similarly Salmon’s (1984, 1998) and Lewis’s (1986) causal theories of explanation refer to the causal history of an event in terms of a succession of events that constitute its explanation.

Theories in psychology like attribution theory (Ajzen & Fishbein, 1975; Heider, 1958, Hilton & Slugoski, 1986) and its modern derivatives in the form of Bayes Nets (Glymour, 2000) and Causal Power Theory (Cheng, 1997) mirrored this emphasis on events in theorizing about explanations. This preoccupation with events has left a gap in the literature on how we understand, judge and generate explanations of properties.

Property Explanations

Properties are attributes, features and characteristics of things. In Philosophy there is an ongoing debate about the nature of properties, whether they exist without being instantiated and whether they can do any explanatory work at all (Swoyer, 1996). However, here the focus was on how people understand and judge explanations of properties. Thus we take properties to be any descriptive phrase that people consider to apply generally to an object or a natural kind. In fact, the properties used here were all sampled from a database based on asking people to list as many properties for a particular object as they could (Cree & McRae, 2003). As a consequence, these properties are mostly enduring characteristics that apply to the complete class of things in question. In addition superordinate category membership was also considered a property of the entities.

Explanations of properties are ubiquitous in everyday life. We wonder why structurally complex organisms are mostly diploid, why flat pack assembly instructions have to be so complicated and why most leading conservative politicians are so badly dressed. Questions about properties sometimes seem more fundamental than questions about events. They are about the characteristics of things that endure and sometimes define or make the entity what it is. They ask about the system of the entity, the underlying nature and the interplay of properties and therefore seem less likely to be explicable circumstantially. As the focus of research on explanation has so far been on events, it is of empirical interest to explore how we understand and judge explanations of properties.

In this paper the focus was on a particular type of explanation of properties. In an exploratory study we found that when people were asked to generate an explanation for a property of an object they regularly referred to another property of that object to provide the explanation. For instance, participants consistently refer to the rarity of emeralds when asked to explain why they are expensive. Property explanations commonly took the following form:

“Xs have p because they have q!”

where ‘p’ and ‘q’ are two properties of the concept ‘X’. The natural question that arose was what determines whether ‘q’ is a good explanans for ‘p’ in any given X. That is the question that we address in this paper.

Study

The overall aim of our study was to investigate the determinants of plausibility in property explanations. Regression analysis was used to establish the relative contribution of a range of predictors to the plausibility of property explanations. The predictors consisted of measures of causal, statistical and local coherence relations between properties and a global measure of feature centrality for individual properties. Pairs of properties were selected and each property pair was tested in both directions for each measure. This enabled us to test whether property explanations exhibit symmetry; i.e. where two properties are equally able to explain one another. For instance, do people judge 'Emeralds are expensive because they are rare' as equally plausible as 'Emeralds are rare because they are expensive'. Similarly the causal, statistical and dependence relations were measured in both directions.

Method

Materials A sample of 28 natural kind and 28 artifact concepts was randomly drawn from Cree and McRae's (2003) database. Using the same database two properties for each concept were drawn at random and two additional properties were selected so that roughly equal numbers of property types (see Wu & Barsalou's, 2002 classification) were present in each of the two domains.

The four properties per concept were paired up to create two pairs such that at least one of the resulting explanations seemed plausible to the experimenter. For instance the four properties 'having a siren', 'being a vehicle', 'being large' and 'being used for emergency' for the concept 'Ambulance' could be paired up in three different ways. The relation between 'having a siren' and 'being used for emergency' seemed most likely to produce a plausible explanation, and consequently this pairing was adopted.

Each pair of properties was used to generate two explanations (e.g. "Axes are dangerous because they are sharp." vs. "Axes are sharp because they are dangerous.") resulting in four explanations per concept. The exact same principal applied to co-occurrence judgments (e.g. "Of all man-made things that are dangerous, what percentage is also sharp?" vs. "Of all man-made things that are sharp, what percentage is also dangerous?") and counterfactuals (e.g. "If axes were not dangerous, would they be sharp?" vs. "If axes were not sharp, would they be dangerous?").

Mutability was judged for each property individually (e.g. "How difficult is it to imagine axes that are not dangerous?" vs. "How difficult is it to imagine axes that are not sharp?"). Dependence was measured by network diagrams consisting of the concept with all four properties displayed. Participants saw all 56 concepts and were asked to draw arrows between the properties of a concept indicating the strength of the dependence by using different colors. Based on pilot work, participants were asked to consider carefully the direction of the dependence, and only to draw bidirectional arrows if they were convinced that they applied. (see Figure 1; a solid line represents the strongest dependence relation, a dotted line represents weakest dependence). Arrows are drawn from one property to a property that it depends on. The full set of measures is illustrated in Figure 2.

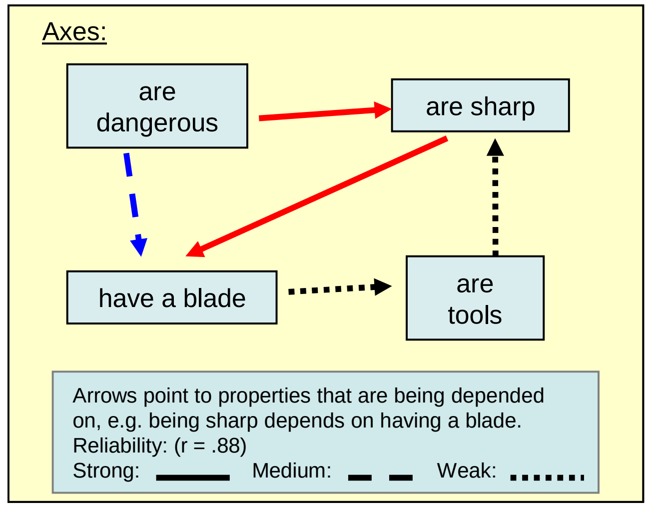

Figure 1. Dependence Net Diagrams

Participants A total of 386 participants were recruited for the five different measures. The three paper-based measures were collected at City University London with a total of 80 participants each for co-occurrence and plausibility judgments and 21 for the dependence measure. The sample size was 109 and 96 for counterfactuals and mutability respectively. The whole sample consisted of University undergraduate and postgraduate students with an average age of 22 years. The majority (88%) were native English speakers with the remainder having a high command of English as their second language.

Procedure & Design The complete set of 224 items was divided into four subsets, each of which contained one of the four explanations per concept. Each questionnaire therefore contained 56 target items with an additional four warm-up items at the beginning of each questionnaire. Two different random orders of items were used for each questionnaire. Each subset and each measure was completed by different groups of participants. Counterfactuals and mutability were collected online using a website and recruitment took place via email. The paper-based questionnaires were collected in classrooms. For each measure, participants had the chance to win a £20 voucher from Amazon as an incentive to participate.

Results

All measures had a Cronbach Alpha between .88 and .94 (see figure 2). The data were averaged across participants for each of the measures and the analyses were carried out on the individual items. All variables apart from the dependence measure were roughly normally distributed. Dependence was consequently transformed by taking the square root of its values. In the analyses that follow the transformed version of both dependence measures were used.

Concept: ***Axes***

Property p: ***dangerous*** **Property q:** ***sharp***

Plausibility Judgments (p > q): X has p because it has q

How plausible is the explanation? {r = .88 - .92}*

Axes are dangerous because they are sharp.

Dependence (p>q) p depends on q {r = .88}

For axes being dangerous depends on being sharp.

Dependence (q>p) q depends on p

For axes being sharp depends on being dangerous.

Counterfactuals (¬ p): If not p then q? {r = .90 - .95}

If axes were not dangerous would they be sharp?

Counterfactuals (¬ q): If not q then p?

If axes were not sharp would they be dangerous?

Co-occurrence (%q in p): Percentage of q in p {r = .90 - .93}

Of all man-made things that are dangerous, what percentage is sharp?

Co-occurrence (%p in q): Percentage of p in q

Of all man-made things that are sharp, what percentage is dangerous?

Mutability (mu - p): Mutability of p {r = .93 - .94}

How difficult is it to imagine axes that are not dangerous?

Mutability (mu - q): Mutability of q

How difficult is it to imagine axes that are not sharp?

(*Numbers in brackets represent reliability range for each measure)

Figure 2. Summary of all Measures

First we compared the average scores across domains (artifact vs. natural kind) on each of the measures using a one-way multivariate analysis of variance. Results of evaluation of assumptions of normality, homogeneity of variance-covariance matrices, linearity and multi-collinearity were satisfactory.

The combined dependent measures were significantly affected by domain, $F(5, 218) = 6.63$, $p < .001$. Although significant, the results only showed a small association between domain and the combined dependent variables, partial $\eta^2 = .13$, reflecting the small effective difference between the two domains across the different measures.

In univariate ANOVAs of each measure only plausibility ($F(1, 222) = 20.03$, $p < .001$) and counterfactuals ($F(1, 222) = 12.75$, $p < .001$) showed significant differences between the two domains. Pairs of artifact properties produced both more plausible explanations and were more counterfactually dependent on each other than those of natural kinds (see Table 1).

Property Explanations across Property Types. Table 2 provides frequencies of plausible explanations (items that were rated above 5 on a 10-point scale for plausibility) broken down by the different property types for the two positions in the explanation and the two domains. Starting with natural kinds we can see that out of the 24 plausible explanations 10 (42%) had a superordinate property in the 'q'-position. A further 6 explanations consisted of a superordinate property in the 'p'-position resulting in two thirds of all plausible property explanations for natural kinds containing a superordinate in either the 'p' or the 'q'-position.

For artifacts, the 52 plausible explanations were more evenly distributed across the different property types with components (44.2%) and functions (25%) being the most strongly represented property types in 'q'-position with a similar distribution for the 'p'-position.

However, taking the base rate of the different property types into account revealed a different picture. Figure 3 represents the proportion of plausible explanations over the total number of explanations for each property type. Of all explanations that contained a function in 'q'-position 86% were plausible, compared with only 47% of property explanations containing a component in q position. Thus despite the range of property types in plausible explanations for artifacts, taking the base rate of the different property types into account there was a clear dominance of functional properties as explanations for artifacts.

Correlation and Regression The pattern of correlations was virtually the same in each domain in terms of rank order and absolute strengths. Regression models including all predictors for the two domains were compared by calculating the difference in regression coefficients across domain. Comparing all significant predictors in either domain with their counterpart in the other domain showed no differences greater than two standard errors. In addition, a principle components analysis showed the same pattern of loadings for the two domains. Both domains produced a two factor solution with artifacts showing smaller cross-loadings than natural kinds. Thus, in what follows the items were collapsed across domains.

Table 1: Domain Differences with Mean and (Standard Deviation) for each Measure.

	Natural Kind		Artifact	
Plausibility	4.11	(1.40)	5.09	(1.80)
Dependence	0.64	(0.64)	0.74	(0.65)
Counterfactuals	2.17	(0.75)	2.58	(0.94)
Mutability	5.76	(1.80)	5.58	(2.00)
Co-occurrence	48.61	(18.0)	49.93	(21.0)

Table 2: Distribution of property type pairing across domains and position.

Natural Kind

	'q' – position				
'p' – position	Component	Function	Superordinate	Other	Total 'p'
Component	0	0	6	2	33%
Functional	3	0	0	2	21%
Superordinate	3	0	0	3	25%
Other	1	0	4	0	21%
Total 'q':	29%	0%	42%	29%	24

Artifact

	'q' – position				
'p' - position	Component	Function	Superordinate	Other	Total 'p'
Component	10	7	3	4	46%
Functional	8	0	1	3	23%
Superordinate	4	1	0	0	10%
Other	1	5	3	2	21%
Total 'q':	44%	25%	13%	17%	52

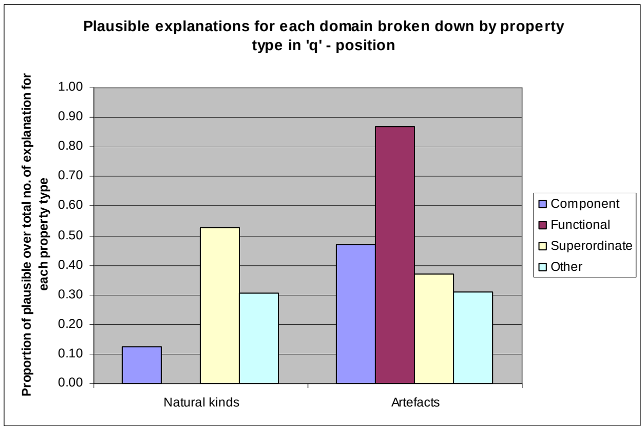

Figure 3. Plausible explanation over total no. of explanations by domain and property type

Table 3 shows the correlations between plausibility and the significant predictors. In line with our predictions both dependence (p>q) (r = .625) and counterfactual (¬q) (r = .567) were strongly correlated with plausibility (p>q). Mutability (mu-q) (r = .319) and co-occurrence (%p in q) (r = .305) were only moderately correlated with plausibility. Interestingly the other co-occurrence measure (%q in p) (r = .372) had just as strong a correlation with plausibility as its converse. For example, “axes are dangerous because they are sharp” is rendered more plausible as an explanation if a high percentage of dangerous things tend to be sharp (%q in p), as well as if a high percentage of sharp things tend to be dangerous (%p in q)

In the instructions for the dependence diagrams participants were asked to choose and only draw the stronger direction of the two dependence relations. As a result the two complement dependence measures showed a small but significant negative correlation (r = -.132, p < .05). For all the other predictors there were no significant correlations between the two complements. The only other correlation among complements that was significant was for plausibility (r = .441, p < .001) indicating that many property explanations work in both directions.

Suprisingly we found that a number of explanations were equally plausible in both directions. For instance, ‘Whistles are used for alerting because they are loud’ was as plausible as ‘Whistles are loud because they are used for alerting.’ This symmetry was more prominent for artifacts. Nevertheless some natural kind explanations also showed symmetry; (e.g.) ‘Carrots are roots because they are found underground’ and ‘Carrots are found underground because they are roots.’

The overall aim of the study was to establish a model to predict plausibility judgments of these property explanations. A standard regression with plausibility (p>q) as dependent variable and each of the two variables for counterfactuals, co-occurrence, mutability and dependence as predictors was carried out. As mentioned above dependence was the only measure to show a deviation from normality in its distribution. A square root transformation of dependence was used to normalize the distribution.

Dependence (p>q) turned out to be the strongest predictor of plausibility (p>q) followed by the equally strong counterfactuals (¬q) and (¬p) and then mutability-q. Table 3 provides the results for the final regression model. Overall the model predicted 56% of variance in plausibility judgments. Neither of the two co-occurrence measures entered the model despite their moderate positive correlations with plausibility. Interestingly despite the weak correlations with plausibility both complement measures of dependence (q>p) and counterfactuals (¬p) entered the model. Dependence (p>q) turned out to be the strongest predictor of plausibility (p>q) followed by the equally strong counterfactuals (¬q) and (¬p) and then mutability-q. Table 3 provides the results for the final regression model. Overall the model predicted 56% of variance in plausibility judgments. Neither of the two co-occurrence measures entered the model despite their moderate positive correlations with plausibility. Interestingly despite the weak correlations with plausibility both complement measures of dependence (q>p) and counterfactuals (¬p) entered the model.

Table 3: Regression Table for the final model with plausibility as dependent variable.

	Plausibility (p>q)	Dependence (p>q)	Counterfactual (¬p)	Counterfactual (¬q)	Dependence (q>p)	Mutability -q	B	Beta
Dependence (p>q)	0.625**						2.010	.478
Counterfactual (¬p)	0.299**	-0.004					.474	.241
Counterfactual (¬q)	0.567**	0.546**	0.094				.441	.224
Dependence (q>p)	0.182**	-0.132*	0.546**	-0.004			.522	.124
Mutability-q	-0.319**	-0.281**	0.194**	-0.396**	0.065		-.133	-.151
Mean	4.59	.69	2.37	2.37	.69	5.67	R^2	**.569**
StD.	1.70	.64	.41	.41	.64	1.94	Adj. R^2	**.559**

**. Correlation is significant at the 0.01 level (2-tailed).

*. Correlation is significant at the 0.05 level (2-tailed).

A further point to note was that, despite dependence having the lowest Cronbach alpha with .88, it turned out to have the greatest influence on plausibility. In theory two equally strong predictors would show differential influence on the dependent variable as a result of their reliability with the variable with higher reliability showing a stronger influence. Thus, despite the slight underestimation of the influence of dependence (p>q) on plausibility, it turned out to be the strongest predictor in this model. However the correlation between dependence (p>q) and plausibility (p>q) ($r = .625$) was not significantly different from the correlation between counterfactual (¬q) and plausibility (p>q) ($r = .567$), which means that counterfactual (¬q) could have equally turned out to be the strongest predictor. Thus no strong claims can be made here about the dominance of one of the two measures over the other.

Discussion

Property explanations were found to be more plausible for artifacts than for natural kinds. For natural kinds, plausible property explanations contained mostly superordinates in either the 'p' or the 'q'-position of the explanation. In contrast for artifacts the dominant property type in plausible explanations was function. Plausibility judgments of some property explanations in both domains showed symmetry, i.e. 'p' and 'q' were equally capable of explaining one another. In the regression, dependence (p>q) followed by both counterfactual measures turned out to be the strongest predictor of plausibility. However, as the correlations of dependence (p>q) with plausibility did not significantly differ from the correlation of counterfactuals (¬q) with plausibility, no strong claims can be made about the superiority of dependence over counterfactuals as predictors of plausibility.

The domain difference for plausibility and for counterfactuals between living natural kinds and man-made artifacts was significant. Despite the effect being rather small, artifacts had twice as many property explanations with high plausibility ratings as did natural kinds. One attempt to explain the domain difference might be by reference to the predominant property types in plausible explanations. Most of the plausible explanations for natural kinds contained a superordinate in either the 'p' or the 'q'-position of the explanation, whereas for artifacts a range of property types were explanatory. Thus, if natural kind property explanations were constrained by having to contain a superordinate in order to be plausible and only a fifth of all the explanations in our item pool contained a superordinate then the proportion of plausible explanations for natural kinds is inevitably smaller than that for artifacts.

This raises the question of why superordinates are the only explanatory property for natural kinds. Intuitively one might say that the reason why living natural kinds have the properties they have, is because of some evolutionary processes that brought them about.[1] Only superordinates might be considered to be associated with evolutionary processes and thus have explanatory value. Apart from that, properties in and of themselves don't refer to or instantiate evolutionary processes that could be explanatory, therefore property explanations for natural kinds were less plausible overall.

Furthermore, comparing the two domains, functions play a large explanatory role for artifacts, whereas for whole natural kinds, they are not explanatory at all (McLaughlin, 2000). The only functional properties generated for natural kinds in Cree and McRae's (2003) database were functions that the entity had for humans, e.g. 'horses are used for pulling things.' But these kinds of functional properties do not readily explain why for instance the horse has certain features. Thus in the present study, functions of natural kinds were not explanatory.

One might object that a human heart is a natural kind and that it clearly has a function that has explanatory power for its features. However this objection only holds if the natural kind is part of a self-reproducing system, where the function provides an evolutionary advantage and thereby explains its existence (McLaughlin, 2000). In the present case all natural kinds were complete entities, for which participants did not generate functions as their properties (Cree & McRae, 2003).

Artifacts conversely do have functions. Their existence seems to be based on the functions they are or were meant to perform. Most of their properties can be explained by reference to their intended function. Thus, for artifacts, having both superordinates and functions as explanatory properties resulted in the overall observed domain difference.

But why are functions so explanatory for artifacts and superordinates so explanatory for natural kinds? Psychological essentialism (Medin & Ortony, 1989) may provide an account for both. Psychological essentialism is the view that we represent things in the world as having essences which bring about their properties and make the entities what they are. For natural kinds, superordinate properties are the most likely candidates to stand in for essential properties and therefore are most able to explain surface features of natural kinds.

One view argues that for artifacts that intended functions constitute their essences (Bloom, 1996). The essential features bring about the non-essential properties in an entity and as a result are explanatory for them.

Another view might simply be that functions in artifacts and superordinates in natural kinds constitute a common cause structure for their properties, which doesn't rely on the idea of essences. The present study though was not equipped to distinguish between these two possibilities. Whether or not we hold the stronger view of psychological essentialism, what we found was that some kind of common cause structure underlies plausibility judgments for property explanations. The most likely candidates for these common causes were superordinates for natural kinds and functions for artifacts.

A surprising finding made in this study was the symmetry that some of the explanations exhibited. If property explanations are predominantly based on causal relations as is thought to be the case for event explanations, then these explanations should not exhibit symmetry because causal

[1] Our sample of items only contained living natural kinds.

relations are asymmetrical (Salmon, 1998). One view might be that property explanations are not based on causal relations between the two properties and therefore are able to exhibit symmetry, which would undermine any account that mainly relies on causal relations (Sloman, 2005). However correlations between the symmetry of plausibility and the symmetry of counterfactuals were not significant, suggesting that for some items that were equally plausible in both directions, counterfactual judgments were only high in one direction.

Another account of the symmetry might be that one of the directions is plausible in a diagnostic rather than explanatory way. "Flamingos are birds because they have wings" might be understood as 'I can tell that flamingos are birds because they have wings.' This may be a plausible account for natural kinds involving superordinates, however it would not explain the symmetry of artifact items like: "Sofas have cushions because they are used for relaxing." Future studies will have to address this distinctive feature of property explanations.

The regression suggested that the best way to model plausibility of property explanations was in the form of dependence net diagrams that capture the local aspect of conceptual coherence. Mutability in contrast captures the more global aspect of coherence by measuring the centrality of a property to its concept. The combined finding of the strong influence of dependence and the lack of influence of mutability on plausibility suggests that the local aspect of conceptual coherence is more important for plausibility judgments than the global.

The present results are mostly in line with Thagard and Verbeurgt's (1998) view of coherence as a constraint satisfaction model. In their view concepts are coherent sets of properties. Coherence depends on the relations between these properties. Certain relations might then be seen to provide stronger constraints than others. As in Quine's (1960) idea of a coherent net of knowledge, certain properties are more difficult to remove from our conception of a concept because a number of other properties stand in strong constraint relations to them. They constitute local coherence constraints. Other properties of the same concept might not depend on any of these properties and may be part of their own local coherence net. Coherence under this view is a net of pair-wise constraint relations with some parts that are more tightly connected and others more loosely.

This view provides an explanation for why conceptual coherence was so much more influential when measured as pair-wise dependence rather than as mutability. According to Sloman et al. (1998) dependence is a very general and basic notion. "Every directional, semantic relation between features can be treated as a generic dependency relation." (Sloman et al., 1998, p. 204). Thus, when judging dependence, people are able to form a general impression that encompasses causal, logical or any other type of relation between the properties. With this view of dependence in place, we can see that judgments of dependence would incorporate the other measure of counterfactuals, co-occurrence and mutability, so that the dominance of dependence as a predictor of plausibility would follow directly. Future study will have to address both the symmetry and the domain difference in plausibility of property explanations.

References

Ajzen, I., & Fishbein, M. (1975). A bayesian analysis of attribution processes. *Psychological Bulletin, 82*, 261-277.

Bloom, P. (1996). Intention, history, and artifact concepts. *Cognition*, *60*, 1–29.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104 (2)*, 367–405.

Cree, G., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General, 132*, 163–201.

Glymour, C. (2000). Bayes nets as psychological models. In F. Keil & R. Wilson (Eds.), *Explanation and Cognition*. MIT Press.

Heider, F. (1958, reprinted 1983). *The psychology of interpersonal relations*. NJ, US: Lawrence Erlbaum Associates.

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

Hilton, D. J., & Slugoski, B. R. (1986). Knowledge-based causal attribution - The abnormal conditions focus model. *Psychological Review, 93 (1)*, 75–88.

Keil, F., & Wilson, R. (2000). Explaining explanation. In F. Keil & R. Wilson (Eds.), *Explanation and Cognition*. MIT Press.

Lewis, D. (1986a). Causal explanation. In D.-H. Ruben (Ed.), *Explanation*. Oxford, UK: Oxford University Press.

McLaughlin, P. (2000). *What functions explain: Functional explanations and self-reproducing systems*. Cambridge University Press.

Medin, D., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou and A. Ortony (Eds.), *Similarity and analogical reasoning.* Cambridge University Press.

Quine, W. (1960). *Word and object*. MIT Press.

Salmon, W. C. (1998). *Causality and explanation*. Oxford University Press.

Sloman, S. A. (2005). *Causal Models*. MIT Press.

Sloman, S. A., Love, B. C., & Ahn, W.-k. (1998). Feature centrality and conceptual coherence. *Cognitive Science, 22 (2)*, 189–228.

Swoyer, C. (1996). Theories of Properties: From Plenitude to Paucity, *Philosophical Perspectives*, 10; 243-264.

Thagard, P., & Verbeurgt, K. (1998). Coherence as constraint satisfaction. *Cognitive Science, 22 (1)*, 1–24.

Wu, L. L., & Barsalou, L. W. (2002). Grounding concepts in perceptual simulation: Evidence from property generation. Manuscript submitted for publication.