

Close Does Count: Evidence of a Proximity Effect in Inference from Causal Knowledge

Russell C. Burnett (rburnett@mssc.huji.ac.il)

School of Education, The Hebrew University of Jerusalem
Jerusalem 91905, Israel

Abstract

Two studies are reported in which participants drew inferences about variables in systems of causal relationships. Previous work has shown that such inferences are influenced by information about variables that, on a normative account of causal reasoning, should be irrelevant. The present studies tested two hypotheses about how relevance is assigned to these normatively irrelevant variables. Though results were mixed, they suggest that greater relevance is assigned to variables that are closer in known causal structure to the variable about which an inference is being made.

Though the learning of causal relationships from correlational evidence has received a fair amount of attention from psychologists, the use of causal knowledge to make inferences and predictions has not. This may be due to the fact that until recently most studies have focused on how learners detect the existence or strength of a causal relationship between just two dichotomous variables: a single cause and a single effect, each either present or absent. Drawing inferences or predictions from knowledge of a single causal relationship is presumably straightforward in the sense that each variable is predictive of the other: An effect is more likely present when its cause is present, and vice versa.

This paper deals with inferences about variables in complex causal systems, that is, systems of causal relationships among three or more variables. Such inferences are often less straightforward than inferences from single causal relationships, in that it is not always so clear whether or under what conditions variables are relevant to one another. Suppose, for example, that a doctor knows that virus X causes a certain enzyme deficiency, which in turn causes liver damage. That is, the doctor knows this three-variable causal chain:

virus X \rightarrow enzyme deficiency \rightarrow liver damage .

Suppose this doctor sees a patient for whom both the presence or absence of the virus and the presence or absence of the enzyme deficiency are known, and the doctor must make an inference about whether this patient is at risk for liver damage. In this case virus X and enzyme deficiency can be called *observed variables*, and liver damage an *unobserved variable*. To what extent would (or should) information about each of the two observed variables influence the doctor's inference?

Previous work suggests that such inferences are partly but not fully explained by a normative or rational theoretical framework, variously known as causal Bayesian network theory or graphical causal model theory (Pearl, 2000; Spirtes, Glymour, & Scheines, 2000), which has in recent

years been applied to the psychology of learning complex causal systems (e.g., Gopnik et al., 2004; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). The central principle in this framework is known as the *causal Markov condition*, and it can be interpreted as saying just when variables in causal systems are relevant to one another and when they are not. Formally, the causal Markov condition says that a variable is independent of all variables that are not its descendants in causal structure, conditional on its immediate cause(s). For the current example, this means that liver damage is independent of virus X (a nondescendant in causal structure) if the presence/absence of the enzyme deficiency (the immediate cause of liver damage) is known. Since it is known whether the patient has the enzyme deficiency, the presence or absence of virus X is irrelevant to an inference about liver damage.

This kind of conditional independence—dependence of two variables just when the state of a third is known—is often called “screening off.” We say, for example, that virus X is screened off from liver damage by the enzyme deficiency. Indirectly related variables in causal chains are screened off from one another by mediating variables, but this is not the only form that screening off can take. For instance, the causal Markov condition also implies that variables with a single common cause are screened off from one another by that cause. To modify the current example, if it were known that virus X causes, by separate mechanisms, both the enzyme deficiency and liver damage, then enzyme deficiency would be irrelevant to an inference about liver damage given information about the virus. This form of screening off—screening off by a common cause—was described by Reichenbach (1956).

Rehder and Burnett (2005) asked participants to draw inferences like the ones just described and found that greatest relevance was indeed assigned to variables deemed relevant by the causal Markov condition. However, inferences were also influenced by variables that, according to the causal Markov condition, should have been screened off from the variable about which inferences were being made. This was found for causal systems with several different structures, including the chain and common-cause structures. When making inferences about a variable in a chain, participants implicitly assigned relevance to indirectly related variables even when the state of a mediating variable was known. When making inferences about one of multiple effects of a single cause, participants assigned relevance to the other effects even when the state of the common cause was known (for a related finding see Walsh & Sloman, 2004). This phenomenon was termed *nonindependence*, since relevance was assigned in violation

of the independencies specified by the Markov condition.

The aim of the current study is to clarify *how* relevance, or inferential support, is assigned to normatively screened-off variables—that is, to clarify the form that nonindependence naturally takes.

Uniformity versus Proximity

Here it is proposed that nonindependence has a rational basis, and from this rational basis are derived two hypotheses about how inferential support is assigned to normatively screened-off variables.

The causal Markov condition holds only if a causal model is complete, in the sense that there is no unknown common cause of any combination of known variables (an assumption that Spirtes et al., 2000, call “causal sufficiency”) and no unknown causal path between any two known variables. If a model fails to satisfy these conditions, then variables may be relevant to one another in ways that violate the Markov condition. Consider again the doctor who knows

virus X \rightarrow enzyme deficiency \rightarrow liver damage .

If, unbeknownst to the doctor, there is a common cause of virus X and liver damage, or a relationship between virus X and liver damage that is not mediated by the enzyme deficiency, then the virus and liver damage may be relevant to one another even when the presence/absence of the enzyme deficiency is known. This is important because natural causal knowledge tends to be surprisingly incomplete (Keil, 2003; Rozenblit & Keil, 2002). Indeed, as Hausman and Woodward (1999) have noted, it is often incomplete in just the ways that invalidate the Markov condition (see also Cartwright, 1999). Nonindependence, then, can be seen as a rational way of compensating for a mismatch between an assumption of graphical causal model theory and a characteristic of natural causal knowledge. On this account, reasoners assign inferential support more liberally than predicted by the causal Markov condition so as to allow for incompleteness in their knowledge of causal systems.

One way to allow for incompleteness would be to reason as if from an augmented causal model in which a single hidden common cause underlies all of the variables in the known model. This method would assign inferential support to variables in the following way. Since all of the variables that are normatively screened off from one another in the known model are related in just the same way in the augmented model (via the hidden common cause), they should, all else equal, provide equal degrees of inferential support to one another. That is, the relevance or support that is assigned to normatively screened-off variables should be distributed uniformly over these variables; this can be called the *uniform hypothesis*. It predicts, for example, that a reasoner who knows the model in Figure 1 and makes an inference about D given information about A , B , and C will assign equal support to A and B (in addition, of course, to the support assigned to C , which is the one normatively relevant variable in this inference).

The theory that people reason as if from an augmented model with a single common cause was proposed by Rehder and Burnett (2005) in a context where a causal model represents relationships among features of a category of objects. In this context, the theory is nicely consistent with psychological essentialism, or people’s tendency to suppose that a category’s features arise from a single deep cause (Medin & Ortony, 1989).

A more precise method of allowing for incompleteness would be possible if there were some regularity in the relationship between causal knowledge and the true causal structure of the world—that is, if causal knowledge were more likely to be incomplete in some ways than in others. In this case, inferential support could be assigned to normatively screened-off variables according to their probabilities of being related to the variable in question in some unknown way. One strong possibility is that two variables are more likely to be related by an unknown common cause or an unknown path if they are closer to one another in a known causal model. If this is right, then a reasoner who knows the model in Figure 1 would do well to suppose (explicitly or implicitly) that B and D are more likely to be related in some unknown way than are A and D , and to assign greater inferential support to B than to A in an inference about D . This would constitute a proximity effect in the assignment of inferential support, and the hypothesis that inference naturally works in this way can be called the *proximity hypothesis*.

The current experiments were designed to distinguish between the uniform hypothesis and the proximity hypothesis.

Experiment 1

In Experiment 1 participants learned causal systems, developed by Rehder and Hastie (2001), with the chain structure shown in Figure 1. Consider an inference about D given knowledge of the states of A , B , and C . The uniform hypothesis predicts that A and B provide equal support to D , so that inferences are sensitive to whether neither, one, or both of them are present. This prediction is shown in Figure 2a, where the horizontal axis (disregarding the shaded region for now) represents the states of variables A and B (00 = A absent and B absent; 01 = A absent and B present; and so on). The proximity hypothesis predicts the pattern shown in Figure 2b: Both A and B provide support, but B provides greater support than A .

Causal knowledge supports inferences from effect to cause (as in medical diagnosis), as well as from cause to effect. In addition to inferences about D , participants made inferences about A given knowledge of B , C , and D . Predictions are shown by relabeling the axes in Figure 2 as shown in the shaded region. Here B is the normatively relevant variable, and the proximity hypothesis predicts that C provides greater support than D .



Figure 1: Causal chain used in the current experiments.

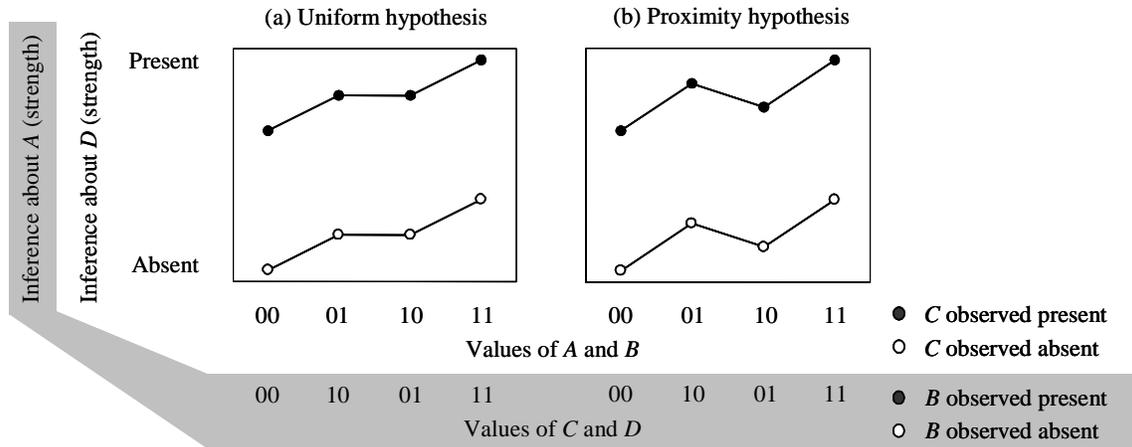


Figure 2: Predictions for Experiment 1.

Method

Participants Participants were 18 introductory psychology students at Northwestern University who received course credit.

Stimuli Stimuli were adopted from Rehder and Hastie (2001). They consisted of six causal systems, each with the chain structure shown in Figure 1. These systems were in various domains: biology, astronomy, chemistry, automobile engineering, and computer design. One of the two biological systems, for example, was said to exist in a kind of shrimp: *A* was the level of the neurotransmitter acetylcholine in a shrimp (high or normal), *B* was the duration of the shrimp's flight response (long or normal), *C* was the rate of the shrimp's sleep cycle (accelerated or normal), and *D* was the shrimp's body weight (high or normal). The non-"normal" values (e.g., high, long, accelerated) were the ones related to one another by causal mechanisms (e.g., a high level of acetylcholine causes a long flight response). We call these values *present*, and the "normal" values *absent*.

Procedure Participants were assigned at random and in equal numbers to the six causal systems. The experiment was administered by computer and involved two phases: a learning phase, in which the participant learned about the assigned causal system, and an inference phase, in which the participant made inferences about unobserved variables in a series of instances.

In the learning phase, the participant read several screens of information about the four variables and the three causal relationships. This information included the mechanisms behind the causal relationships; for example, for the shrimp system, an accelerated sleep cycle was said to cause a high body weight because shrimp feed after waking, and a shrimp that sleeps and therefore wakes more often will eat more. In addition to verbal descriptions of the causal system, the participant was presented with a graphical depiction like Figure 1 (but with values like "accelerated sleep cycle" instead of variables). In order to complete the

learning phase, the participant had to pass a 21-item multiple-choice test on the variables, relationships, and mechanisms. To correctly answer the questions about the three causal relationships, the participant had to rule out other possible relationships among the four variables. Consequently, by the conclusion of this phase, the participant had learned that the four variables were related in just the ways shown in Figure 1.

The inference phase involved a series of 32 instances in which three variables were observed and one was unobserved (e.g., a description of shrimp with a high level of neurotransmitter, a long flight response, a normal sleep cycle, and unknown body weight). On each item, the participant was asked to make an inference about the unobserved variable by positioning a slider on a rating scale. The scale was said to represent probability or confidence; one end represented certainty that the feature in question was absent (e.g., normal body weight), and the other end represented certainty that it was present (e.g., high body weight). Ratings were recorded in the range [0, 100], where 0 = absent, and 100 = present (though participants never saw these numbers). The series of instances comprised all 32 possible items in which three variables were observed (each either present or absent) and one was unobserved. Items were presented in a different random order for each participant.

There was a third phase, administered just before or just after the inference phase, in which participants judged the degree to which each of the 32 instances was a good example of the learned causal system, but this phase is irrelevant to present purposes and will not be discussed.

Results and Discussion

Inferences about *A* and *D* are shown in Figure 3. Inferences about *B* and *C* are less useful for distinguishing between the uniform and proximity hypotheses, since they involve normatively screened-off variables at only one distance from the variable in question. Consequently these won't be reported or analyzed.

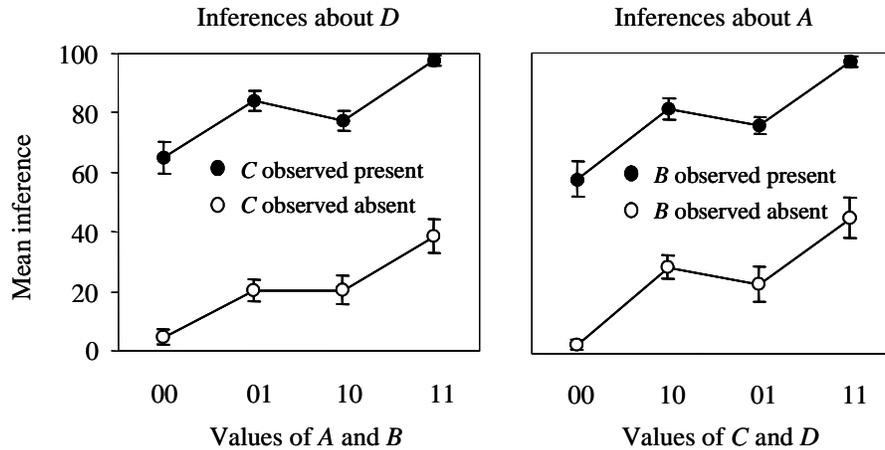


Figure 3: Results, Experiment 1.

As expected, inferences were based heavily on the normatively relevant variable (C in inferences about D , B in inferences about A). The nonindependence effect was also clearly evident: Inferences were influenced by normatively screened-off variables, as reflected in the upward trends in Figure 3. These observations were confirmed by submitting each participant's ratings to a regression analysis with three predictors representing the states of the observed variables (1 = present, or -1 = absent). In inferences about D , the mean weight assigned to C (29.9) was greater than the mean weights assigned to A (7.4), $t(17) = 5.458$, $p < .001$, and B (9.2), $t(17) = 5.573$, $p < .001$. In evidence of nonindependence, the mean weights assigned to A and B were greater than zero, $t(17) = 4.901$, $p < .001$, and $t(17) = 7.013$, $p < .001$, respectively. Likewise, in inferences about A , the mean weight on B (26.7) was greater than the mean weights on C (11.6), $t(17) = 5.078$, $p < .001$, and D (8.8), $t(17) = 4.796$, $p < .001$, and the weights on C and D were greater than zero, $t(17) = 8.160$, $p < .001$, and $t(17) = 4.689$, $p < .001$.

On distinguishing between the uniform and proximity hypotheses, the results were mixed. First consider inferences about D . In these inferences the mean weight assigned to B (9.2) was not reliably greater than the mean weight assigned to A (7.4), $t(17) = 1.204$, $p = .25$, though these means did differ in the expected direction. For finer resolution, participants were grouped according to whether they assigned greater weight to B than to A (consistent with proximity), equal weight to the two, or greater weight to A than to B . The numbers of participants who gave these three orderings of weights were 10, 3, and 5, respectively. This is suggestive of a proximity effect, though a chi-square test on these frequencies falls short of reliability, $\chi^2(2) = 4.333$, $p = .11$. In sum, these overall analyses reveal some evidence of a proximity effect, but this evidence is not statistically significant.

On the other hand, the relative degrees of support assigned to A and B seem to depend somewhat on whether C , the normatively relevant variable, was present or absent. To see this, let each stimulus be named by its values on the four variables such that, for example, 101x indicates an item

in which A is present, B is absent, C is present, and the state of D is to be inferred. When C was absent, mean inference ratings were equal (to 21) when just A was present (in 100x) and when just B was present (in 010x); that is, uniform support was assigned to A and B . But when C was present, there was some evidence of a proximity effect. Mean ratings given to 011x and 101x (84 and 77, respectively) differed in the expected direction, and this difference was marginally reliable, $t(17) = 2.097$, $p = .05$. For finer resolution, participants were grouped according to whether they gave a higher rating to 011x than to 101x (consistent with proximity), equal ratings to the two items, or a lower rating to 011x than to 101x. The numbers of participants who fell into these three groups were 11, 3, and 4, respectively, which suggests a reliable tendency for proximity-based inference, $\chi^2(2) = 6.333$, $p < .05$. In contrast, the numbers of participants whose responses fit these patterns when C was absent were 5, 7, and 6.

Inferences about A show a similar pattern. Overall, the average weight on C (11.6) was not reliably greater than the average weight on D (8.8), $t(17) = 1.603$, $p = .13$, though again the direction of the difference was consistent with a proximity effect. The numbers of participants who assigned greater weight to C , equal weights to C and D , and greater weight to D were 10, 3, and 5, respectively, which is again suggestive of a tendency toward proximity-based inference, $\chi^2(2) = 4.333$, $p = .11$. When B was present, the mean ratings given to x101 and x110 were 76 and 81, $t(17) = 1.47$, $p = .16$, and the numbers of participants whose ratings of these items were in the order predicted by proximity, equal, and in the opposite order were 8, 4, and 6, respectively, $\chi^2(2) = 1.333$, $p = .51$. When B was absent, the mean ratings given to x001 and x010 were 23 and 29, respectively, $t(17) = 0.77$, $p = .45$, and the numbers of participants who ratings of these items were in the order predicted by proximity, equal, and in the opposite order were 10, 3, and 5, $\chi^2(2) = 4.333$, $p = .11$.

In sum, all differences between means were in the direction predicted by proximity, and proximity-consistent orderings of ratings and weights were most frequent in all cases; however, most differences fell short of reliability.

It should be noted in retrospect that the power to detect a proximity effect was low. Given 18 participants and the observed variance, the power to detect a difference in weights of 2.0 was around .2. A study with greater sample size and a longer causal chain (so that normatively screened-off variables are at more than two distances from the variable in question) is planned. Meanwhile, Experiment 2 approaches the problem in a different way.

Experiment 2

Experiment 1 provided some suggestive evidence in support of the proximity hypothesis. To distinguish more powerfully between the uniform and proximity hypotheses, participants in Experiment 2 were asked to make forced choices rather than ratings on a continuous scale. In each forced-choice problem, one variable was unobserved, and two normatively screened-off variables were pitted against each other. These two variables were at different distances from the unobserved variable in question, and so the proximity hypothesis made a clear prediction on each problem. The uniform hypothesis, in contrast, predicted no preference for either choice.

The same causal systems as in Experiment 1 were used. Half of the problems involved inferences about *A* (the initial variable in the chain), and half were about *D* (the final variable). In half of the items the immediate neighbor of the variable in question (*B* in problems concerning *A*, or *C* in problems concerning *D*) was observed present, and in half it was observed absent. This was to test the possibility, raised in Experiment 1, that proximity-based inference is more likely when the screening-off variable is present than when it is absent.

Method

Participants Participants were 22 members of the Northwestern University community.

Stimuli The experiment was run as a paper-and-pencil task. Each causal system was described on a cover page, in much the same way as in the learning phase of Experiment 1 (with both verbal description and graphical depiction). Attached to this cover page were two inference problems. Each problem involved descriptions of two configurations of values on three observed variables. The state of the fourth variable was said to be unknown, and the participant was asked to indicate in which of the two configurations the unobserved variable was more likely to be present. For example, a participant might be given the two configurations 101x and 011x and the question “Which of these shrimp do you think is more likely to have high body weight [variable *D*]?” The instructions were to check one of the two options and to provide a justification.

Procedure Participants were assigned at random and in roughly equal numbers to the six causal systems. Each participant made forced choices on two different inference problems. One concerned the final variable in the chain (*D*), and in this problem *C* was either observed present (011x versus 101x) or observed absent (010x versus 100x). The other concerned the initial variable in the chain (*A*), where *B*

Table 1: Numbers of choices consistent with proximity.

Variable in question	Neighbor present	Neighbor absent
<i>D</i>	8/10	9/11
<i>A</i>	8/10	8/11

Note. Neighbor is *C* in inferences about *D*, *B* in inferences about *A*.

was either observed present (x110 versus x101) or observed absent (x010 versus x001). Whether *C* (in problems concerning *D*) or *B* (in problems concerning *A*) was observed present or observed absent was counterbalanced across participants and inference problems, and the order of the two problems was counterbalanced across participants. Whether the choice predicted by the proximity model appeared on the left or on the right varied randomly.

Results and Discussion

Choice data are presented in Table 1. (One participant is not represented in these counts because she declined to choose, saying that both choices on each problem were equally likely to have the unobserved variable present—a response consistent with the causal Markov condition.) Overall there was a strong tendency to choose in accordance with the proximity hypothesis. In inferences about *D*, 17 of 21 choices were consistent with proximity, $\chi^2(1) = 8.05$, $p < .01$. In inferences about *A*, 16 of 21 choices were consistent with proximity, $\chi^2(1) = 5.76$, $p < .05$. There was no evidence that this tendency depended either on whether the normatively relevant variable was present or absent or on whether inferential support was derived from upstream or downstream in the causal chain (i.e., whether inference was about *A* or *D*). Dividing the data on either of these dimensions yields nearly equal numbers of proximity-consistent choices.

Justifications fell into three main categories:

(1) Proximity. For example, in an inference about *A*, a participant who based his choice on *C* rather than *D* wrote that “[*D*] is a more distant emergent property of [*A*] than is [*C*].” On an analogous item another participant wrote, “[*A*] more closely linked to [*C*].”

(2) Theories about hidden causal structure. Several justifications involved explicit reasoning about hidden common causes and hidden paths. For example, one participant theorized that a shrimp’s quantity of the neurotransmitter acetylcholine (*A*) and body weight (*D*) had a common cause, the amount of choline-rich algae eaten by the shrimp. (The choline-rich algae was mentioned in the cover story.) Another participant, who had learned about a causal system involving characteristics of a certain kind of molecule, inferred a hidden path between a molecule’s structure (*B*) and its reactivity (*D*): “The pyramid structure...seems to contain more overall energy, so it is possibly more prone to react.”

(3) Consistency with known causal relationships. For example: “Since we’re seeing some obvious causal violations, I’ll use the same reasoning as before: The system with the most causal violations is more likely to show another violation.” This justification implies, for example,

that 011x is more likely to have *D* present than is 101x, because 011x shows just one violation of known causal structure, whereas 101x shows two.

Overall, 35 of the 37 interpretable justifications were in these three categories; 17 involved proximity, 11 involved theories, and 7 involved consistency with known causal relationships. Of the theories about hidden causal structure, 8 involved specific hidden paths, 2 involved specific hidden common causes, and 1 was a general appeal to the possibility of hidden relationships. Notably, these theories supported proximity-inconsistent choices (6 times) as often as proximity-consistent choices (5 times). More importantly for present purposes, the most frequent kind of justification was an explicit appeal to structural proximity.

Conclusion

Taken together, results of these experiments constitute evidence of a proximity effect in inference from complex causal models. In previous work, inferences were shown to assign relevance to variables that, on a normative account of causal reasoning, should have been screened off from the variables about which inferences were made—a phenomenon called nonindependence (Rehder & Burnett, 2005). The present results clarify the form that nonindependence naturally takes. They suggest that relevance is assigned to normatively screened-off variables as a function of proximity to the variable in question. In Experiment 1, though most differences fell short of statistical significance, the directions of differences between means favored the proximity hypothesis in all cases, as did the numbers of participants whose inferences and implied weightings fell in the orders predicted by proximity. Results of Experiment 2 were less ambiguous. The great majority of forced choices were as predicted by the proximity hypothesis, and the most frequent type of justification was an explicit appeal to proximity.

The proximity effect can be seen as a rational way of compensating for incompleteness in causal knowledge on an assumption about causal knowledge and causal truth, namely, that proximity between variables in known causal structure reflects the likelihood that these variables are related via unknown common causes or unknown paths in true causal structure.

The theory that people reason as if from an augmented causal model with a single common cause of all known variables was proposed by Rehder and Burnett (2005) in a context where causal models represent relationships among features of objects. The proximity effect can be seen as complementing or elaborating on this theory. A single common cause may often be the most salient kind of hidden causal structure to allow for. This may be especially true in reasoning about features of categories of objects, which are often thought of as having single deep hidden causes (psychological essentialism; Medin & Ortony, 1989). One interpretation of the proximity effect is that, beyond allowing for a single hidden common cause, inference also allows for shallower hidden common causes (that is, ones that underlie just subsets of the known variables) and hidden paths between known variables. Though the current work has been framed in general terms, the stimuli involved

features of objects, and it is an open question whether the proximity effect is stronger for other sorts of causal systems.

It might be argued that the tendency to choose in accordance with proximity in Experiment 2 was an artifact of the forced-choice procedure. Participants may have chosen proximity only because it was more appealing than its opposite. Justifications provide evidence against this interpretation, in that they tended to imply serious reasoning about causal principles (e.g., “*D* is a more distant emergent property of *A*”). Still, further empirical work will be informative. One possibility is that inferences made as in Experiment 1 will reveal the proximity effect more clearly when they involve longer causal chains and otherwise more elaborate causal structures.

Acknowledgments

Many thanks to Bob Rehder for helpful discussions of this work, which grew from a collaboration, and for the stimuli. Thanks also to Doug Medin, Lance Rips, Galen Bodenhausen, Serge Blok, Judith Avrahami, Yaakov Kareev, and anonymous reviewers for comments on this work at various stages.

References

- Cartwright, N. (1999). *The dappled world: A study of the boundaries of science*. Cambridge University Press.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 3–32.
- Hausman, D. M., & Woodward, J. (1999). Independence, invariance and the Causal Markov Condition. *British Journal for the Philosophy of Science*, *50*, 521–583.
- Keil, F. C. (2003). Folkscience: Coarse interpretations of a complex reality. *Trends in Cognitive Sciences*, *7*, 368–373.
- Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*. Cambridge University Press.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge University Press.
- Rehder, B., & Burnett, R. C. (2005). Feature inference and the causal structure of categories. *Cognitive Psychology*, *50*, 264–314.
- Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: The effects of causal beliefs on categorization, induction, and similarity. *Journal of Experimental Psychology: General*, *130*, 323–360.
- Reichenbach, H. (1956). *The direction of time*. University of California Press.
- Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, *26*, 521–562.
- Spirites, P., Glymour, C., & Scheines, R. (2000). *Causation, prediction, and search* (2nd ed.). MIT Press.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–489.
- Walsh, C. R., & Sloman, S. A. (2004). Revising causal beliefs. *Proceedings of the 26th annual conference of the Cognitive Science Society*.