Cognitive Psychology 72 (2014) 54–107

Contents lists available at ScienceDirect

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

Independence and dependence in human causal reasoning

Bob Rehder *

Department of Psychology, New York University, New York, NY 10003, United States

ARTICLE INFO

Article history:
Accepted 11 February 2014

Keywords:
Causal reasoning
Causal inference
Causal Markov condition
Conditional independence
Screening off

ABSTRACT

Causal graphical models (CGMs) are a popular formalism used to model human causal reasoning and learning. The key property of CGMs is the *causal Markov condition*, which stipulates patterns of independence and dependence among causally related variables. Five experiments found that while adult's causal inferences exhibited aspects of veridical causal reasoning, they also exhibited a small but tenacious tendency to violate the Markov condition. They also failed to exhibit robust *discounting* in which the presence of one cause as an explanation of an effect makes the presence of another less likely. Instead, subjects often reasoned "associatively," that is, assumed that the presence of one variable implied the presence of other, causally related variables, even those that were (according to the Markov condition) conditionally independent. This tendency was unaffected by manipulations (e.g., response deadlines) known to influence fast and intuitive reasoning processes, suggesting that an associative response to a causal reasoning question is sometimes the product of careful and deliberate thinking. That about 60% of the erroneous associative inferences were made by about a quarter of the subjects suggests the presence of substantial individual differences in this tendency. There was also evidence that inferences were influenced by subjects' assumptions about factors that disable causal relations and their use of a conjunctive reasoning strategy. Theories that strive to provide high fidelity accounts of human causal reasoning will need to relax the independence constraints imposed by CGMs.

* Address: Dept. of Psychology, 6 Washington Place, New York, NY 10003, United States.
E-mail address: bob.rehder@nyu.edu

http://dx.doi.org/10.1016/j.cogpsych.2014.02.002
0010-0285/

0. Introduction

People possess numerous beliefs about the causal structure of the world. They believe that sunrises make roosters crow, that smoking causes lung cancer, and that alcohol consumption leads to traffic accidents. The value of such knowledge lies in allowing one to infer more about a situation that what can be directly observed. For example, one generates explanations by reasoning backward to ascertain the causes of the event at hand. One also reasons forward to predict what might happen in the future. On the basis of, say, a friend's inebriated state, we predict dire consequences if he were to drive and so hide his car keys.

A large number of studies have investigated how humans make causal inferences. One simple question is: When two variables, X and Y, are causally related, do people infer one from the other? Unsurprisingly, research has confirmed that they do, as X is deemed more likely in the presence of Y and vice versa (Fernbach, Darlow, & Sloman, 2010; Meder, Hagmayer, & Waldmann, 2008, 2009; Rehder & Burnett, 2005; see Rottman & Hastie, 2013, for a review). But causal inferences quickly become more complicated if just one additional variable is introduced. For example, suppose that X and Y are related to one another not directly but rather through a third variable Z. Under these conditions, the question of how one should draw an inference between X and Y will depend on the direction of the causal relations that link them via Z. Three possibilities are shown in Fig. 1. First, X and Y might both be effects of Z (Fig. 1A), a topology referred to as a *common cause* network. For example, a doctor might diagnose a disease (Z) on the basis of a particular symptom (X), and then also predict that the patient will soon exhibit another symptom of that disease (Y). Second, the variables might form a *causal chain* in which X causes Z which causes Y (Fig. 1B). For example, politicians may (X) calculate that pandering to extremists will lead to their support (Z), which in turn will galvanize members of the opposing party (Y). Finally, Z might be caused by X or Y, forming a *common effect* network (Fig. 1C). A police detective might release an individual (Y) suspected of murder (Z) upon discovering the murder weapon in possession of another suspect (X).

A formalism that specifies the permissible forms of causal inferences and that is generally accepted as normative is known as *causal graphical models*, hereafter *CGM* (Glymour, 1998; Jordan, 1999; Koller & Friedman, 2009; Pearl, 1988, 2000; Spirtes et al., 2000). CGMs are types of *Bayesian networks* (or *directed acyclic graphs*) in which variables are represented as nodes and directed edges between those variables are interpreted as causal relations. Note that a CGM need not be complete in the sense that variables may have exogenous influences (i.e., *hidden causes*) that are not part of the model; however, these influences are constrained to be uncorrelated. This property, referred to as *causal sufficiency* (Spirtes, Glymour, and Scheines, 1993, 2000), in turn has important implications for the sorts of

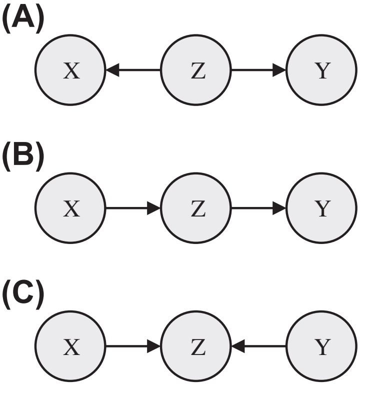

Fig. 1. Three causal networks that can be formed from three variables. (A) A common cause network. (B) A chain network. (C) A common effect network.