[redacted]

[redacted]

ArtiŢcial Intelligence 169 (2005) 174–179

[redacted]

www.elsevier.com/locate/artint

Book review

Judea Pearl, Causality, Cambridge University Press, 2000.

Henry E. Kyburg Jr.

University of Rochester, Institute for Human and Machine Cognition, USA

Available online 28 October 2005

This is a remarkable volume. Winner of the Lakatos award, given biennially for the book in the philosophy of science most highly regarded by an international committee, it is also crammed with formulas that will be of practical importance, as well as of interest, to epidemiologists, lawyers, economists, and other down-to-earth folk. This is not to say that it will be easy to read, for anyone, or that it is altogether correct. I shall Ţrst offer a review of the contents of the book, and then carp (minimally) about the viewpoint. The review of the contents will be highly schematic, since the book is extremely rich.

Many of the results reported in the Ţrst six chapters here parallel results also achieved by researchers at CMU [4]. The Ţrst two chapters present background, and introduce the terminology of probabilistic networks. Chapter One introduces some probability, with a one page nod to subjectivism; conditional probability is glossed as "... given that I know A" [p. 5]. Subsequently the references are to the Ţrst person plural, which suggests a relatively objective conception of probability. Most often we encounter relations among probabilities (such as $P(y \mid x, z) = P(y \mid z)$) which make perfectly good sense as objective frequencies. The issue of subjectivism is one to which we will return later.

Among the basic ideas introduced in the Ţrst chapter are these: Directed Acyclic Graphs (DAG's); conditional independence; Bayesian networks; the Markov property, and d-separation. The very important operator $do(X = x)$ that introduces a way of treating *actions* is introduced on p. 23. This operator sets a subset X of variables to constants x, yielding an *interventional distribution.* Let P_* be the set of all interventional distributions $P_x = P(v \mid do(X = x))$ (including the empty intervention). A *DAG G* is a *causal Bayesian network* compatible with P_* if and only if the following three conditions hold for every P_x in P_*:

1. $P_x(v)$ is Markov relative to G;
2. $P_x(v_i) = 1$ for all $V_i \in X$ whenever v_i is consistent with $X = x$;

E-mail address: hkyburg@ihmc.us (H.E. Kyburg).

doi:10.1016/j.artint.2005.10.003