In *Workshop on Conditional Independence Structures and Graphical Models*, Toronto, Canada, September, 1999.

The Observability of Independence under Monetary-Based Elicitation

David M. Pennock and Michael P. Wellman

University of Michigan AI Lab, Ann Arbor, U.S.A.

Graphical models [1] have proved invaluable as a language for encoding probability distributions, due primarily to their facility in expressing *independencies* among variables. Sufficient independencies yield not only parsimonious models, but insights into the underlying system being represented; conversely, the lack of any independence structure renders a graphical description, per se, essentially devoid of content.

Consider a modeler attempting to faithfully reproduce the dependency structure of a single expert.[1] If the expert's beliefs were somehow directly observable, (e.g., his self-reports were reliably accurate and honest) then the problem would be simply one of transcribing his beliefs. Realistically, the modeler has access only to the *actions* of the expert, for example his betting behavior or verbal reports, which are inextricably linked to utility as well as belief. Various elicitation aids have been designed to extract beliefs from actions, including lotteries, scoring rules [6], and promissory notes [2]. This task is underconstrained, though, and any one of a continuous family of belief–utility pairs offers an equally valid explanation for a particular behavior. Elicitation devices based on monetary incentives actually reveal *risk-neutral* probabilities, or the probabilities at which a risk-neutral agent would choose the same actions, which may differ significantly from true probabilities [3,4].

The modeler seeks to describe the joint distribution across a set of S binary, uncertain events, $Z = \{A_1, A_2, \ldots, A_S\}$. Let $\Omega = \{\omega_1, \omega_2, \ldots, \omega_{2^S}\}$ be the set of all 2^S possible joint outcomes of the events, also called the *atomic states*. The expert has a subjective probability distribution Pr over Ω and utility $u(y)$ for y dollars, which may in general depend on the state. Let $\mathrm{CI}(A_j, W, X)$ be shorthand for the conditional independence $\Pr(A_j|WX) = \Pr(A_j|W)$.

Risk-neutral probabilities are defined as $\Pr^{\mathrm{RN}}(\omega) \propto \Pr(\omega)u'(Y_\omega)$, where Y_ω is the agent's monetary wealth in state ω, and u' is the derivative of utility [4]. The expert's behavior, manifested as actions chosen to maximize expected utility, is indistinguishable from that of a hypothetical agent with transformed probabilities $\Pr^{\mathrm{RN}}(\omega)$ and reciprocally transformed, risk-neutral utility $u^{\mathrm{RN}}(y) \equiv u(y)/u'(Y_\omega)$. Any test employed to monitor the expert's choices (e.g., querying the prices at which he would buy or sell various lottery tickets) essentially reveals $\Pr^{\mathrm{RN}}$, and *not* Pr. Define $\mathrm{CI}^{\mathrm{RN}}$ analogously to CI.

We examine first the case where $S = 2$. It can be shown that

$$\mathrm{CI}(A_1, \emptyset, A_2) \,\&\, \left(\frac{u'(Y_{\bar{A}_1 A_2})}{u'(Y_{A_1 A_2})} = \frac{u'(Y_{\bar{A}_1 \bar{A}_2})}{u'(Y_{A_1 \bar{A}_2})}\right) \Rightarrow \mathrm{CI}^{\mathrm{RN}}(A_1, \emptyset, A_2). \tag{1}$$

That the expert considers the events to be independent is *not* sufficient for the modeler to ascertain it. The second term of the precondition of (1) is true (and indeed $\Pr^{\mathrm{RN}} = \Pr$) if the expert's marginal utility is constant across states. This holds if the expert is risk neutral, and holds approximately if utility is state-independent and $Y_{\omega_j} \approx Y_{\omega_k}$. The modeler can enforce this approximation by choosing small enough inducements, if those constitute the

[1] The difficulties raised here are compounded when the modeler consults more than one expert, or includes his own assessments in the mix [5].

only stakes in the events available to the expert. In general, however, the expert may hold nonnegligible outside stakes beyond the modeler's control [3].

One might hope that, despite any quantitative difference between $\Pr^{\mathrm{RN}}$ and Pr, the underlying qualitative independence might be preserved. This turns out to be true only in very limited circumstances. Assume the expert has *constant absolute risk aversion* (CARA), which corresponds to negative exponential utility $u(y) = -e^{-cy}$. Further assume that $Y_{A_1A_2} = Y_{A_1} + Y_{A_2}$, implying that the expert may own stakes in the events, but not in arbitrary joint outcomes of the events. Now we find that

$$\frac{u'(Y_{\bar{A}_1} + Y_{A_2})}{u'(Y_{A_1} + Y_{A_2})} = \frac{ce^{-cY_{\bar{A}_1}}e^{-cY_{A_2}}}{ce^{-cY_{A_1}}e^{-cY_{A_2}}} = \frac{e^{-cY_{\bar{A}_1}}}{e^{-cY_{A_1}}} = \frac{ce^{-cY_{\bar{A}_1}}e^{-cY_{\bar{A}_2}}}{ce^{-cY_{A_1}}e^{-cY_{\bar{A}_2}}} = \frac{u'(Y_{\bar{A}_1} + Y_{\bar{A}_2})}{u'(Y_{A_1} + Y_{\bar{A}_2})}.$$

Thus the precondition of (1) is satisfied, and the independence is observable.

Now consider $S > 2$. In this case, even if the expert has CARA and owns stakes only in the events, not *all* of his independencies are invariant under the transformation to $\Pr^{\mathrm{RN}}$. In particular, suppose that $S = 3$, with A_1 and A_2 independent, and A_3 dependent on both. It can be shown by example that $\mathrm{CI}^{\mathrm{RN}}(A_1, \emptyset, A_2)$ does *not* hold generically. On the positive side, a subclass of independencies, called *Markov independencies are* invariant. An event A_j is Markov independent of a set of events X, given another set W, if $\mathrm{CI}(A_j, W, X)$ and $A_j \cup W \cup X = Z$. In this setting,

$$\mathrm{CI}(A_j, W, X) \Rightarrow \mathrm{CI}^{\mathrm{RN}}(A_j, W, X)$$

holds whenever the expert has CARA and $Y_{A_jWX} = Y_{A_j} + Y_{WX}$. It follows that a *Markov network* [1] description of $\Pr^{\mathrm{RN}}$ has the same topology as one of Pr. The structure of a *Bayesian network* (BN) encoding of $\Pr^{\mathrm{RN}}$ is a moralized and triangulated version of the BN structure of Pr.

References

[1] J. N. Darroch, S. L. Lauritzen, and T. P. Speed. Markov fields and log-linear interaction models for contingency tables. *Annals of Statistics*, 8(3):522–539, 1980.

[2] B. de Finetti. *Theory of Probability.* J. Wiley, New York, 1974.

[3] J. B. Kadane and R. L. Winkler. Separating probability elicitation from utilities. *Journal of the American Statistical Association*, 83(402):357–363, 1988.

[4] R. F. Nau. Coherent decision analysis with inseparable probabilities and utilities. *Journal of Risk and Uncertainty*, 10(1):71–91, 1995.

[5] D. M. Pennock and M. P. Wellman. Graphical representations of consensus belief. In *Fifteenth Conference on Uncertainty in Artificial Intelligence*, 1999.

[6] R. L. Winkler and A. H. Murphy. Good probability assessors. *Journal of Applied Meteorology*, 7:751–758, 1968.

[University of Michigan AI Lab, 1101 Beal Avenue, Ann Arbor, MI 48109-2110 USA; emails: {dpennock,wellman}@umich.edu]