ROBERT C. STALNAKER

LETTER TO DAVID LEWIS

May 21, 1972

It *does* seem to me worth noting that if P is a probability distribution, and if for any A and B, $P_B(A) = P(B > A)$, then P_B is a probability distribution too (excepting the absurd case). What it is good for, I would like to suggest, is deliberation – the calculation of expected utilities.

Let $S_1, \ldots, S_n$ be an exhaustive set of mutually exclusive propositions characterizing the alternative possible outcomes of some contemplated action. Let A be the proposition that I perform the action. My suggestion is that expected utility should be defined as follows:

$$u(A) = P(A > S_1) \times u(S_1) + \cdots + P(A > S_n) \times u(S_n).$$

Why $P(A > S_i)$ rather than $P(S_i/A)$? Because what is relevant to deliberation is a comparison of what will happen if I perform some action with what *would* happen if I instead did something else. A difference between $P(S/A)$ and $P(S)$ represents a belief that A is *evidentially relevant* to the truth of S, but not necessarily a belief that the action has any causal influence on the outcome. That a person performs a certain kind of action can be evidence that makes some state subjectively more probable, even when the action in no way contributes to the state. Suppose that this is true for some action A and desirable state S. Then $P(S/A) > P(S)$, but only an ostrich would count this as any sort of reason inclining one to bring it about that A. To do so would be to act so as to change the evidence, knowing full well that one is in no way changing the facts for which the evidence is evidence.

I am thinking of Nozick's puzzle ("Newcomb's problem", in the Hempel festschrift), which I just discovered, but which I assume you know. My intuitive reaction to this puzzle was the following: there is only one rational choice (assuming there is no backwards causation in the case), and that is to choose the dominating action. But this seems to conflict with the principle of maximizing expected utility. But from my suggested version of the principle, the rational choice follows. The principle of expected utility may be held to be universally applicable.

Since quotient conditionalization is the way to revise your beliefs, it is also rational in the Newcomb problem to bet, after having made the rational choice, that you will fail to get the million dollars. Had you made the other

W. L. Harper, R. Stalnaker, and G. Pearce (eds.), Ifs, 151–152.

choice, it would have been rational to bet that you would succeed in getting the million dollars. But *this* is no reason to wish that you had chosen differently, since you could have changed only the fair betting odds, not the facts, by acting differently.

The suggested version of the expected utility principle makes it possible for a single principle to account for various mixed cases: the probabilistic dependence may have two components, one causal and one non-causal. The components may reinforce each other, or counteract each other. They might cancel out, leaving the evidence irrelevant, even though there is a believed causal dependence. Also, it may be unknown whether the probabilistic dependence is causal or not. Imagine a man deliberating about whether or not to smoke. There are two, equally likely hypotheses (according to his beliefs) for explaining the statistical correlation between smoking and cancer: (1) a genetic disposition to cancer is correlated with a genetic tendency to the sort of nervous disposition which often inclines one to smoke. (2) Smoking, more or less, causes cancer in some cases. If hypothesis (1) is true, he has no independent way to find out whether or not he has the right sort of nervous disposition. In such a case, it seems clear that the probability of the conditional (if I were to smoke, I would get cancer), and not the conditional probability is what is relevant . . . .