REFERENCES

* This essay was written with partial support from the National Science Foundation, Grant No. GS-2050.
1 R. Carnap [1].
2 *Ibid.*, p. 309.
3 P. Suppes [7].
4 In addition to the paper mentioned see [6] as well as [8], which is Suppes' comment on R. C. Jeffrey's paper [3]. Jeffrey's reply is instructive.
5 See P. A. Samuelson [4], p. 257, and footnote 8, p. 263.
6 Whether strict coherence is plausible or whether more coherence should be required is debatable. However, adoption of the weaker assumption would complicate the exposition by requiring qualification of the claim that the set of sentences accepted as evidence is identical with the set of sentences accorded a credence value of 1. No point essential to this discussion would be lost, however, if the qualification were to be made.
7 See, for example, L. S. Savage [5], pp. 46–55.
8 Patrick Suppes has emphasized this point. See [6], p. 35.
9 Of course, X might be mistaken in reporting that $c(H, X\text{'s total evidence}) = r$. Indeed, the metalinguistic predicate $c(H, \_\_) = r$ may be true of no non empty set of sentences in X's language. Observe, however, that X can be mistaken in what he takes to be his explicitly given evidence. The shift from explicit to implicit characterizations of evidence involves no change in this respect.
10 See, for example, R. C. Jeffrey [2], pp. 156–7.
11 *Op. cit.*, Ch. 11. Also [3].
12 Jeffrey [3].
13 *Op. cit.*, p. 179.
14 *Ibid.*, p. 176.

RICHARD C. JEFFREY

DRACULA MEETS WOLFMAN: ACCEPTANCE VS. PARTIAL BELIEF

One of the things I'd like to see come out of this conference is a clarification of the issues between Isaac Levi and me.[1] I take it that Levi's scruples about partial belief and probability kinematics are not idiosyncratic, nor are my scruples about his work on acceptance, so that the matter may be of general interest. When Dracula meets Wolfman in the movies it is not simply I-and-thou: They gibber and slaver for all vampires and all werewolves everywhere. So let it be with us.

I. MY KIND OF PRAGMATISM

Levi and I both seem to be pragmatists of sorts – of different sorts. Thus, he quotes with approval from Peirce ([11], p. 208), "The settlement of opinion is the sole end of inquiry", but my kind of pragmatism has deliberation, not opinion, as its focus: It tries to see how much mileage it can get out of the fact that man is an agent. Opinion and valuation come together in the matrix of deliberation, from which action emerges.

I do not mean to deny that men value understanding, and seek it for its own sake as well as for its practical value; but I do not see understanding as a matter of relief from agnosticism, nor do I find it illuminating to characterize scientists as seeking to replace doubt by true belief. (Still less do I find it illuminating to consider such questions as, 'Am I a *true* pragmatist?' and 'Have I betrayed fallibilism?" – see [11], p. 209.)

Deliberation can take a variety of forms, and has done, in the history of mankind. Our notion of what counts as rational decision-making keeps evolving – anyway, changing. The view that I find most satisfactory is very dimly adumbrated in the *Protagoras* (ca. 356) and Aristotle's *De Anima* (434^a 7–10); more sharply in *The Port-Royal Logic* (last Chapter), and more sharply still in Daniel Bernoulli's *Exposition of a New Theory on the Measurement of Risk* (1738); and has finally come into its own in the past 30 or 40 years through the work of Ramsey, De Finetti, Savage, and others. As an important social force, it has come into its own only in the last decade. L. J. Savage's *The Foundations of Statistics* was a radical book, when it appeared in 1954 – it took some years for his 'Personalistic' standpoint to be recognized as a serious and important one – but now the Ramsey-Savage view of decision-making is much in vogue, in the corridors of governmental and economic power. Students at the Harvard Business School learn decision theory from a book [12] by Howard Riaffa, who is Frank P. Ramsey Professor of Managerial Economics and a member of the Graduate School of Public Administration. According to *The New York Times Magazine* ([16], p. 39)

> Decision theory, as it is taught – and, to some extent, created – at the business school, rests on the concept of "subjective probabilities", which allows for the decision-maker's attitude toward risk. Through a device called "the preference curve" (which assesses, say, how much it stings a man to lose $50000 on a chance of gaining $275000) this attitude may be translated into numbers and incorporated into decision-tree calculations.

I have no proof that my own account of rational deliberation is right, or uniquely right.[2] Rather, that account aims at being a fair representation of what currently passes as sensible deliberation. Mind you, part of the reason why what passes, passes is that it fits into such coherent frameworks as Ramsey and his successors have given, for deliberation – we like to feel we know what we are doing, or why we are doing it, in that sense. I am pretty happy with (say) my account in the sense that I think that something like it will continue for a while to serve as an adequate-seeming framework. On the other hand, I am pretty sure that it is not quite satisfactory as it stands; I expect a more satisfactory account to appear soon. But you see what I mean: A currently satisfactory account of practical reason is satisfactory in being a coherent description, not of all practice, but of the instances of practice now thought patently sound – and an account succeeds as a description partly because of its persuasive force.

Perhaps Savage would put the matter more strongly. In his book he presents seven postulates which he hopes everyone will agree are true of all sensible preference rankings of possible acts – everyone: Personalists and those who have never heard of Personalism and those who have heard of it but reject it. He then gives a mathematical proof that anyone whose preference ranking satisfies the postulates is already, in effect, a Personalist, i.e. for any such person there is a probability measure p and a utility function u – where p is determined uniquely by the preference ranking, and u is determined uniquely except for choice of a zero and a unit –, which characterize the ranking. In particular, one act will be ranked with, above, or below another accordingly as its expected utility as computed via p and u is equal to, greater than, or less than the expected utility of the other.

One might find Savage's argument interesting and valuable, but less than compelling. Someone who sees himself as believing in various special decision methods that are incompatible with Personalism and who also sees himself as accepting Savage's postulates as norms will see that he has an inconsistent view of himself, after he follows Savage's proof. Likely as not, he will then have a hard look at the individually plausible-seeming postulates, in which so much proves to be buried when they are viewed collectively; and he may find that certain of them are unacceptable after all. This is a situation of the well-known sort in which you take yourself to believe p, q, and $-r$, and then discover that p and q together imply r. The logical alarm has rung: you will want to give up one (or more) of $p, q, -r$, but logic does not tell you which. By the way, the account of decision-making that I espouse is like deductive logic in that it does not tell you what to do, or what to prefer to what; but it *does* ring a quasi-logical alarm e.g. when you find that you prefer p to q, and q to $p \vee q$, and then discover that p and q are logically incompatible; for in such a case, $p \vee q$ must be somewhere between p and q in your preference ranking, if the 'logic of decision' is not to be violated.

Ethan Bolker[3] has given a set of postulates that bear the same relation to my account of probabilities and utilities that Savage's postulates bear to his. Of course, I find Bolker's postulates acceptable, individually and *in toto*; but I, like Bolker, find that in the case of one of the postulates, I must translate it into terms of probabilities and utilities in order to see that I accept it. But *that* does not discourage me. The conclusion of a valid inference is always lurking somewhere in the premises, and to the extent that the conclusion is a surprise, one has not fully understood the (implications of the) premises! Furthermore, the conclusion is generally weaker than the conjunction of the premises, and therefore more likely to be true; then *prima facie* it seems a bit unpromising to try to convince someone of the conclusion indirectly, by convincing him that the premises are all true, and the inference valid. In the inference under discussion, from Bolker's postulates, I find the conclusion eminently reasonable isimply because the probability/utility way of looking at decision problems has succeeded so well in making sense of a wide diversity of kinds of decisions, some of which were initially rather puzzling. As far as I know, this Personalistic or 'Bayesian' account is unique in its ability to accommodate the whole range of decision problems, while clearly marking out the domains of applicability of various other, more special accounts, e.g. that in which one seeks to minimize maximum disutility. I am a Bayesian because I see the matter in this way – not conversely.[4]

But what has any of this to do with belief? Well, the Bayesian thesis might be summarized by the statement that when A and B are propositions you are sure do not both hold, i.e., when your belief function p is such that $p(A\&B)=0$, and when $p(A \vee B)$ is not 0, then the utility $u(A \vee B)$ which you ascribe to one or the other of A, B's holding ought to be a weighted average of the utilities you ascribe to them separately:

$$u(A \vee B) = w_1\, u(A) + w_2\, u(B),$$

where the weights w_1, w_2 add up to 1 and are proportional to your degrees of belief in A and in B. Then we must have $w_1 = p(A)/p(A \vee B)$ and $w_2 = p(B)/p(A \vee B)$. Here, the promise of $A \vee B$ is viewed as a gamble with two possible outcomes: A, B. Your valuation of $A \vee B$ is then determined by your valuations of A and of B, and by your beliefs about A and about B, where 'belief' here is understood as an attribute of your attitude toward various risks. In particular, in terms of your attitudes toward various gambles in which only small gains and losses are possible (so that your utility curve is fairly linear) your degree of belief in a proposition A is p if and only if you think it fair to bet on A at odds of $1-p:p$ (e.g., to pay $\$p$ to get $1 if A is true, and lose your $\$p$ if A is false).

It is one of the charms of the Bayesian position that it uses a notion of partial belief which is clearly and simply related to the agent's attitudes toward risk. The charm is twofold. For one thing, it gives one fairly clear, quasi-operational criteria for determining whether someone has a definite degree of belief in a certain proposition at a certain time and, if he has, for discovering what that degree of belief is; and it gives one a rather beautiful way of showing that degrees of belief ought to satisfy the axioms of the elementary probability calculus, viz., that degree of belief ought to be a normalized, non-negative, additive function of propositions.[5] This notion of degree of belief accords well enough with one strand in

our talk about belief, but it is not intended as a bit of ordinary language analysis or as a bit of phenomenology. This notion of belief is the theorist's, and need not be the agent's. Thus, suppose the agent insists that he believes that Black Beauty will win the first race at Suffolk Downs – insists that he fully believes this, is certain of it, etc.; and suppose that he will not give any longer odds than 9:1 on Black Beauty, but *will give* odds of 9:1 or less. I would conclude that his degree of belief in Black Beauty's winning is .9; I would also conclude that he is using the terms 'full belief' and 'certainty' in ways that elude me. But no matter – I shall give him the word 'belief', if he insists, and use 'degree-of-belief' in the clear sense described above, as a technical term.

It need not be easy to tell whether someone *has* a degree of belief in a certain proposition at a certain time, e.g. his attitude toward risks in which that proposition figures might be a rather complex one. One mode of complexity that is compatible with the Bayesian attitude was studied by Friedman and Savage [5]: The bettor's utility function for money may be such that the magnitudes of the possible monetary gains and losses, and not just the odds, *gain*:*loss*, figure in his decision to accept or reject bets. If this is the case, we shall be able to find it out by examining his attitude toward various gambles and certainties, and rationalize his behavior by observing that although the monetary odds may be the same in two bets, the utility odds are different. Thus, an even-money bet for $100 on the toss of a fair coin might be unattractive because the utility gain

(a) $u(\text{being } \$100 \text{ richer}) - u(\text{status quo})$

is less than the utility loss

(b) $u(\text{status quo}) - u(\text{being } \$100 \text{ poorer})$

so that the utility odds, $a{:}b$, are poor although the money odds 100:100 are even. Concretely, plotting the utilities of various incomes against the dollar sizes of those incomes as in Figure 1, the expected utility of a fair, even-money bet for $100 is represented by the height of the midpoint M_1 of the dotted line segment connecting the points L and G_1 which represent the two possible outcomes: lose $100 and gain $100. But the sure income of which the utility is equal to the expected utility of the bet is $-\$50$: Having made the bet, but not yet knowing the outcome, the bettor whose utility function is as shown in Figure 1 would be willing to pay up to $50 to be free of the situation – so much (according to Figure 1) does his dread of losing outweigh his hope of winning, even though he thinks himself as likely to win as lose. In general, where the utility function is concave downward, the agent is adverse to risk. But where the stakes are different, the same bettor may find risk attractive, as in a situation where he will be $100 poorer if a fair coin falls tail up, and $300 richer if it falls head up. Again, the expected utility of the gamble is represented by the height of the midpoint M_2 of the dotted line segment connecting the points L (lose $100) and G_2 (gain $300). The sure income of which the

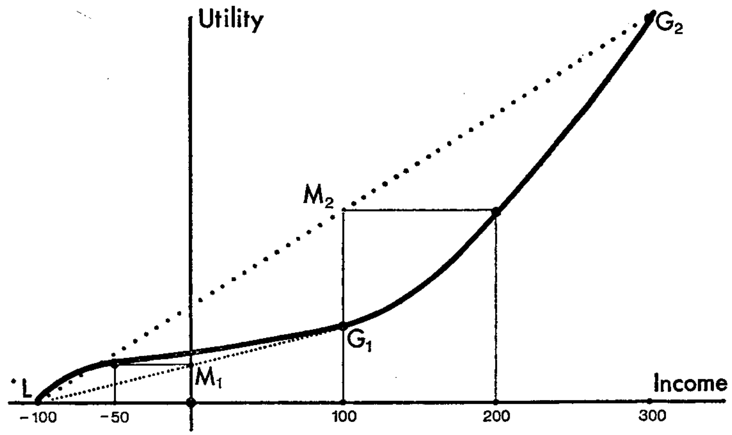

Fig. 1. Rationalizing a complex attitude toward risk.

utility is equal to the expected utility of this bet is $200: Having made the bet, and not knowing the outcome, the bettor would not sell his rights in the bet for less than $200, although the actuarial value of the gamble is only $100. Then apparently clashing attitudes toward risk can make perfect sense, if the agent's utility function has the right sorts of nonlinearity. The particular sort of nonlinearity shown in Figure 1 would be appropriate to someone who could be $50 richer or poorer without radical change in his situation, but for whom greater gains or losses would appear disproportionately attractive or repulsive – greatly expanding or contracting his horizons. (Perhaps with another $200 he could

invest in a very profitable enterprise, while with $200 less he would be bankrupt.)

In other cases, especially where the payoffs are not monetary, the truth or falsehood of the proposition gambled upon may influence the utility of winning or losing in a confusing way. Thus, I may think the weather tomorrow as likely to be foul as not and, quite apart from that, I might rate a ticket to a football game tomorrow as high above the *status quo* as I rate a dull day around the house below it. Abstractly, then, a bet in which I get the ticket if I win and the boredom if I lose ought to look fair to me if I have a degree of belief 1/2 in the proposition bet upon; yet, the bet would be unattractive to me if I were betting on foul weather tomorrow, since that proposition would detract from the utility of winning, while its denial would leave the utility of losing as it was. (See Figure 2.)

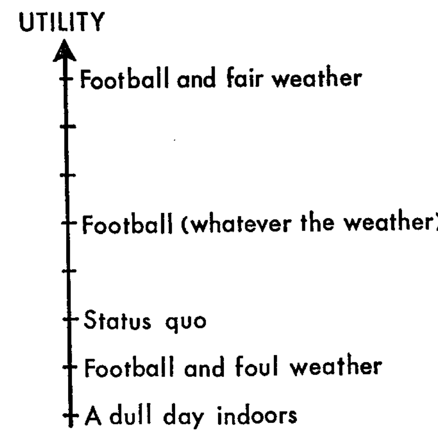

Fig. 2. Enhancement.

To determine whether $p(A)$ exists, for a certain agent, we may have to go far afield, and examine his attitude toward risks in which propositions other than A figure. Imputations of beliefs and values to agents face the tribunal of the agent's attitudes toward risk as a body, although some imputations are especially easy to isolate and test.

The intimacy of the connection between valuation and belief is very evident when we let $B = -A$ in the relationship

$$u(A \vee B) = \frac{p(A)}{p(A \vee B)} u(A) + \frac{p(B)}{p(A \vee B)} u(B),$$

which holds when A is logically incompatible with B and $p(A \vee B)$ is not zero. Since $p(A \vee -A) = 1$ and $p(-A) = 1 - p(A)$, we then have an equation which can be solved for $p(A)$ to get

$$(*) \qquad p(A) = \frac{u(T) - u(-A)}{u(A) - u(-A)}$$

as long as $u(A) \neq u(-A)$. Here, T is the necessary proposition: $T = A \vee -A = B \vee -B = C \vee -C$, etc. Intuitively, T or $A \vee -A$ is a gamble on the proposition A in which the agent gets A if he wins and $-A$ if he loses. In this as in other cases, the gamble is fair (by the agent's lights) if and only if he takes his probability of winning to be

$$\frac{u(\textit{status quo}) - u(\text{losing})}{u(\text{winning}) - u(\text{losing})}.$$

A final illustration: Suppose you like A and B equally well; does it follow that you dislike $-A$ and $-B$ equally? No, you dislike $-A$ more than $-B$ if $p(A)$ is greater than $p(B)$. If you had to choose between $-A$ for sure and $-B$ for sure you should choose $-B$, for then you have a better chance of getting one-or-the-other of the equal goods A, B, than if you had chosen the denial of A, the more probable good. If you work it out algebraically via (*) you will get the same result.

II. RATIONALIZED PARTIAL BELIEF

It seems pretty clear that conformity with the laws of the elementary probability calculus is a necessary condition for reasonableness of a set of partial beliefs, and it should be equally clear that the condition is not sufficient. But let us have a closer look at both claims, especially the second.

Doubts about the first claim have been expressed by Suppes [15] and Hacking [6]. According to Suppes,

> A theory of rationality that does not take account of the specific human powers and limitations of attention, memory and conceptualization may have interesting things to say, but not about human rationality.

Thus, according to the elementary probability calculus, we must have $p(A) \leqslant p(B)$ whenever we have $A \vdash B$ – where $\vdash$ is the deducibility relation in some complete formalization of first-order logic, say. But since there is no general decision procedure for the relation $\vdash$, it seems inappropriate to count someone as irrational whenever $p(A)$ and $p(B)$ both exist and $p(A) > p(B)$ but $A \vdash B$. Better, indeed, to count his degrees of belief as incoherent, and add that with the best will and the best brain in the world, the poor fellow cannot be expected to be aware of all entailments. Of course that is true, and of course we were not thinking of locking him up as a mental incompetent. The point is merely that whether he knows it or not, his beliefs suffer from a logical failing. We would count *him* irrational only if it had been demonstrated to him that $A \vdash B$ and $p(A) > p(B)$, but he regarded the situation with equanimity. And presumably the fellow, being rational, will not want decision theory to be so permissive as to neglect to classify the situation we have been envisaging as a fault, on the ground that it was not *his* fault.

As to the second claim, that conformity with the laws of the elementary probability calculus is not sufficient for rationality of a set of beliefs – I take this to be clearly correct. I mention this fact because Isaac Levi seems to base some of his arguments against my position on the premise that I think conformity with the laws of the probability calculus is sufficient as well as necessary for reasonableness. This passage is a case in point:

> Thus, if coherence is the only obligation imposed on a rational agent in assigning probabilities to propositions, acquisition of new evidence dictates very little to him regarding how he is to revise his probabilities. New evidence is virtually irrelevant to the revision of probabilities ([11], p. 198).

I conjecture that Levi is reasoning as follows: "The only conditions Jeffrey states, as necessary for rationality, are those of coherence, i.e., conformity with the laws of the elementary probability calculus. If he thought there were any further necessary conditions, he would surely have stated them. Since he has not, it must be that the conditions he has stated as necessary are collectively sufficient, in his view."

What I really think is this: Coherence is far from sufficient for rationality, but I do not know how to state any further conditions, except in fuzzy or circular or inaccurate ways; so I have stated no further conditions. Yet, I am sure that further conditions exist. It is a bit as if I were writing an instruction manual for bicycle riding, and gave a number of suggestions, e.g. 'If you feel yourself falling to the left (right) while moving forward, turn the front wheel a bit to the left (right)', and other such tips. I am quite aware that conformity to all suggestions will not make someone a competent bicycle rider, and I would make further suggestions which would complete the set of tips if I knew how; but I do not.

And the situation is a bit like cycling in another respect, I think. Bicycles are made for man. People generally *are* able to learn to ride them, on the basis of such tips as I have mentioned, after some practice. Our nervous and skeletal systems and musculature and sensoria have various special features, common to men and different from the analogous features in dogs, mackerel, etc., and in virtue of some of these features we are in fact able to learn to ride bicycles even though we are unable to give complete explanations of how we do it. Similarly, I think, for the business of forming our systems of partial beliefs. Being the sorts of creatures we are, we are sharply limited (although in ways I cannot clearly describe) in our capabilities for partial belief. And when we add the requirement that our partial beliefs form a coherent probabilistic structure, the limitations become even sharper, although I cannot give a tight general account of how, or why. But I can give some examples:

The Chevalier de Méré. In a letter to Fermat dated July 29, 1654 (see [17], p. 11), Pascal mentions

> a difficult point which astonished M. so greatly, ...
> The odds on throwing a six with a die in 4 trials are 671 to 625.
> But the odds are *against* throwing a double six with a pair of dice in 24 trials. Nevertheless, 24 is to 36 (which is the number of faces of two dice) as 4 is to 6 (which is the number of faces of one die).

We do not quite know the background, but it may have been this: The Chevalier had been betting on double six in sequences of twenty-four tosses of a pair of dice, and losing, on the whole, even after many changes of dice. Pascal pointed out to him that if (as he said he did) Méré regarded all 36 possible outcomes of a single throw of a pair of dice as equally probable, and if he regarded successive throws as independent (as he said he did) then he must regard two sixes in twenty-four throws of a pair of dice as rather less probable than one six on four throws of a single die, and even as a bit less probable than having an even number of points turn up on one throw of one die. In fact, on the stated assumptions, we must have *p*(*one six on four tosses of a die*) $=671/1296>1/2$, and

p(*two sixes on 24 tosses of a pair of dice*) $=[36^{24}-35^{24}]/36^{24}<1/2$, despite the misleading symmetry, viz., 24:36::4:6. In the face of all this, the Chevalier *might* drop the assumption that all 36^{24} possible outcomes of 24 tosses of a pair of dice are equally likely; but being human, he was influenced by the observed frequencies, which conform better with the attitude that *p*(*two sixes on 24 tosses*) $<1/2$ than with his original attitude. Still, according to Pascal ([17], p. 11) the discrepancy

> was his great scandal, which made him say haughtily that the theorems were not consistent and that arithmetic was demented. But you will easily see the reason by the principles you have.

These principles were rather arcane, in 1654, but today they are the common property of millions of high school seniors.

The Birthday Problem [1]. Would you accept an even-money bet for a dollar on at least two people in this room[6] having the same birthday? Perhaps some of you would, and some would not, and some of the discrepancy might be accounted for by different nonlinearities in your utility functions for income. But there might remain a basic disagreement, evinced perhaps in the advice you would give to someone whose utility curve for money was linear in the interval $\pm\$1$ about his *status quo*: some of you would call it a good bet, others a bad one, perhaps. The fact is that it is a good bet if there are at least 23 people in the room. When I make this plonking statement, I am making use of facts we all have, about human births, to say that if any one of you will think the thing out in the light of those facts, you will find that your degree of belief in there being two or more people with the same birthday in this room is greater than 1/2 if you count and find that there are 23 or more people here. The facts are, that there is no reason to suppose that the manner of our selection and self-selection for presence here was biased against sameness of birthdays; that the period about 9 months after the Christmas season is one in which more people in this country have their birthdays than other periods of the same length in the year; and the like. Then if we suppose that each of us is as likely to have his birthday on one day of the year as on any other, and that the probability of one of us having a certain birthday, given that some other of us has some (same or other) definite birthday, is the same as the absolute probability of the first person's having the birthday in question, we shall be making a set of assumptions on which, if anything, it is less likely that two of us have the same birthday than is really the case, according to our more accurately articulated belief functions. And on those assumptions, the probability is a bit over .5 that with 23 people here, two or more of us have the same birthday, while the probability is a bit under .5 if there are 22 people here.

Here, then, are two examples which illustrate how a determination to make one's beliefs conform to the calculus of probabilities can result in what will generally be regarded as a correction of one's original belief system. First, a factual or logical alarm rings: One's beliefs are at odds with observed frequencies (factual alarm), or with each other, if one accepts the laws of probability (logical alarm). In neither case can one say quite generally how we think we should respond after due consideration. One may finally regard a discrepancy between observed frequencies and the frequencies one takes to be most likely as a matter of chance, and for the time, anyway, keep to one's belief function, as when a coin is tossed ten times and there are seven heads; this is rather unlikely, but not astonishing, on the usual views about coin-tossing; and it may be that the coin appears quite normal, so that one has no inclination to adopt an extraordinary view of what is going on. Similarly when the logical alarm rings, it is not generally clear how one will think one should revise one's beliefs; but in practice there is a striking uniformity of response to similar situations. Here's a final example.

The Maturity of Chances. How do you go about convincing someone that he is wrong in thinking that, e.g., with a normal coin, the sequence *hhh* is more likely to continue with a *t* (tail) than with an *h* (head)? Sometimes, you cannot. I mean, sometimes you cannot produce reasons or data that will make him adopt the belief function which assigns probability 1/16 to each possible string of four *h*'s and *t*'s. But usually, such people do not have an internally coherent view of what goes on in coin-tossing according to which $p(hhht/hhh)$ is, say, greater than 1/2. If such a person will agree that any set of degrees of belief in such propositions ought to be extendible to a set which is coherent and assigns degrees of belief to each of *hhhh*, *hhht*, *hhth*, ..., *tttt*, you can generally rely on any particular such extension's having features which disturb him in a way that overbalances his satisfaction at having $p(hhht/hhh)$ be as he originally had it. You might try tossing coins with him, and showing him that among the runs of three heads, about the same number are followed by tails as by heads, but that need not work, partly because you have to toss a large number of times to accumulate much data of that sort, and partly because he might have $p(hhht/hhh)$ only very slightly above 1/2, so that it might take months to collect data he regards as significant. In that case, of course, the two of you might agree to disagree, reasoning that there is little or no practical difference between your views.

I have spoken at length about coins and dice because these are cases where there is broad intersubjective agreement – and even there, deductively reasonable men may differ, although they seldom do for long. In this domain it is entirely possible to go strikingly far beyond mere coherence as a statement of what constitutes a sensible belief function, e.g. one might state that concerning an ordinary penny, tossed in an ordinary manner, no belief function is rational unless it attributes nearly identical probabilities to any two equally long strings of *h*'s and *t*'s. The belief functions that have this feature are rational in the sense that there are commonly available empirical data and arguments about coherence which fairly uniformly have been found to convince people who care to think about it that, after all, their beliefs are as described by one of the 'rational' belief functions; and in the sense that such people regard themselves as having good reasons for abandoning their old belief functions and adopting one of the 'rational' ones.

Then as a practical matter, I think one *can* give necessary conditions for reasonableness of a set of partial beliefs that go beyond mere coherence – in special cases. The result is a patch-work quilt, where the patches have frayed edges, and there are large gaps where we lack patches altogether. It is not the sort of seamless garment philosophers like to wear; but (we ragged pragmatists say), the philosophers are naked! Indeed we have no proof that no more elegant garb than our rags is available, or ever will be, but we have not seen any, yet, as far as we know. We will be the first to snatch it off the racks, when the shipments come in. But perhaps they never will. Anyway, for the time being, we are dressed in rags, tied neatly at the waist with a beautiful cord – probabilistic coherence. (It is only the cord that visibly distinguishes us from the benighted masses.)

What might the seamless garment look like? Well, Carnap has an idea: It will be a conditional probability measure $c(\ /\ )$ defined on a unified language of science which we shall also use in daily life. In that language we shall be able to report the weather, meter readings, and upset stomachs as well as the laws of physics-chemistry-biology as we shall then take them to be. For any sentences p and q in that language, $c(p/q)$ will have a definite value (unless $c(q, p \vee -p)$ is zero) which will presumably be computable as accurately as we please, with the aid of superfast, complex machinery then available. 'In principle', we shall then be able to separate the inductive logical component (represented by the function c) from the experiential component (represented by an enormously long conjunction e of observation reports) in anyone's beliefs, as they ought to be. Of course, there are difficulties, e.g. to accumulate the enormous conjunction e such that for any hypothesis h, $c(h, e)$ is the degree of belief in h that is justified by one's experiences, one would have to spend all one's time writing in one's diary, and one would consequently have little of interest to write; and there are other difficulties, connected with experiences one cannot clearly formulate in the language. But, Carnap would say, the philosophical point is that in principle we then have an empiricist account of justified partial belief, whatever the difficulties of carrying it out in practice.

Indeed, at such a stage in the development of mankind, practice might well be much different from what it is now, e.g. the Academy of Unified Science might be constantly feeding certified evidence-sentences into the machine, to go into a common pool of data. Individuals might also transmit other data for storage in their personal portions of the great Memory Banks via their pocket transceivers, on which they could also request and receive up-to-date rational (for them) degrees of belief in hypotheses that interest them at the moment. It might even be possible to wire transceiving equipment into people's heads, to do much of this automatically. Perhaps in the end we should then have evolved into a colonial form of life, where to a great extent we share a supersensorium and a superconsciousness – one for all.

That is as may be – it is not my object to predict how the race will develop in the long run, if there is one. But as far as I can see now, the c-function built into that supercomputer will be the (or *a*) rationally correct one in at best the evolutionary sense that, when the computer is built, it is the survivor (or one of the survivors) in the process of proposing new candidates for the position and eliminating some of the contestants as having features that almost nobody finds acceptable. In fact, *this* process is a far more tenuous one than the process illustrated above in the case of the Chevalier de Méré etc., because in this process we have no factual alarms – only logical ones. The value of $c(s_1/s_2)$ is what it is independently of what we have observed; that is what makes c a purely 'logical' function, so that the entire experiential component in our rational belief is represented by the particular choice of $s_2 = e =$ one's diary to date.

Then here is where I stand, at the moment. The human race is evolving culturally as well as physically, and part of what is evolving culturally is our notion of what constitutes sensible decision-making and what constitutes rational belief, where by 'belief' I mean the thing that goes along with valuation in decision-making: degree-of-belief, or subjective probability, or personal probability, or grade of credence. I do not care what you call it because I can tell you what it is, and how to measure it, within limits, no matter what you want to call it. (On the other hand I think it an instructive comment on the way we normally talk about belief, to call it 'degree of belief' – but that is secondary. The analysis of ordinary language is not my object either.)

God knows, we may be wrong in having the view we have, of rationality in these matters. To demonstrate this would be to give a persuasive analysis of what constitutes right and wrong in this context, and then show that our current view (to the extent that there is a single current view) is wrong. Better yet, one would provide a new view, that is demonstrably right! But if and when this happens, it will (as far as I can see) be just another instance of cultural evolution: The persuasive analysis will have persuaded the best minds (let us call them) to change their minds, and the Harvard Business School will follow, and government and industry will follow them. And then we shall be in a position rather like the one we are in now: Relative to what will then be called 'now', *as far as we can see now, such-and such is the best account of deliberation and rational belief.* There will, of course, be this difference: at that time we shall presumably have a more copious account, within which we can see the limits and the limited virtues and defects of the present account. That will no doubt be a large ingredient in the confidence men will then feel, in keeping to their account instead of regressing to ours.

Nor would I bet very heavily on our notions of utility and degree of belief being usefully discernible in men's view of rational deliberation, a century from now, except for certain vestiges; nor am I disturbed by the fact that our ordinary notion of *belief* is only vestigially present in the notion of degree of belief. I am inclined to think Ramsey sucked the marrow out of the ordinary notion, and used it to nourish a more adequate view. But maybe there is more there, of value. I hope so. Show me; I have not seen it at all clearly, but it may be there for all that.

III. ON THE DYNAMICS OF PARTIAL BELIEF

In Physics, Dynamics is a contrary of Kinematics as well as of Statics: it is the first contrariety that I had in mind when I called Chapter 11 of *The Logic of Decision*, 'Probability Kinematics'. Take a see-saw, with fulcrum 2/3 of the way toward your end. If you push your end down two feet, the other end will go up three. That is kinematics: You talk about the propagation of motions throughout a system in terms of such constraints as rigidity and manner of linkage. It is the physics of position and time, in terms of which you can talk about velocity and acceleration, but not about force and mass. When you talk about forces – *causes* of accelerations – you are in the realm of dynamics.

So the kinematics of partial belief is concerned with the question, 'Suppose you have a definite coherent set of degrees of belief in various propositions, and then your degree of belief in one or more of them changes; how must that change be propagated throughout the entire system of your beliefs, if coherence is to be preserved?' There is no general answer, any more than there is a general answer to the question, 'Suppose point A rises, on the rigid bar shown in Figure 3. How will point C move?'

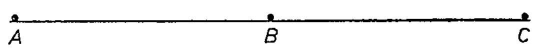

Fig. 3.

You need to know a bit more than that, in order to get an answer, e.g. it would be enough to know that point B remains stationary.

The cases of kinematics that I treated in my book were of this sort: We are given that degrees of belief in propositions

$$(\text{'Basis'}) \quad A_1, A_2, \ldots, A_n$$

change from their present values, $p_0(A_1)$, $p_0(A_2), \ldots, p_0(A_n)$, to new values, $p_1(A_1)$, $p_1(A_2), \ldots, p_1(A_n)$; we are also given that the A's form

a partitioning, i.e., for all $i, j = 1, 2, \ldots, n$,

$$(\text{Partitioning}) \quad \begin{cases} p_1(A_i) \neq 0, \\ p_1(A_i \,\&\, A_j) = 0 \text{ if } i \neq j, \\ p_1(A_1 \vee A_2 \vee \cdots \vee A_n) = 1. \end{cases}$$

Finally, we are given that all conditional probabilities, relative to the individual A's, are the same after the change as they were before: For all $i = 1, 2, \ldots, n$, and for all B for which $p_0(B)$ exists, we have

$$(\text{Rigidity}) \quad p_1(B/A_i) = p_0(B/A_i).$$

If all these conditions hold, and if both p_0 and p_1 satisfy the laws of the elementary probability calculus, then the function p_1 is determined: For any proposition B for which $p_0(B)$ exists, we have ('K' for 'Kinematics')

$$(\text{K}) \qquad p_1(B) = p_1(A_1)\, p_0(B/A_1) + \cdots + p_1(A_n)\, p_0(B/A_n).$$

The proof of all this is as trivial as the proof of Bayes' theorem. It is enough to verify that by the elementary probability calculus, (Partitioning) implies

$$p_1(B) = p_1(A_1)\, p_1(B/A_1) + \cdots + p_1(A_n)\, p_1(B/A_n),$$

and then substitute via the rigidity conditions, to get (K).

An obvious sort of application is that in which some passage of experience leads the agent to change his degrees of belief in the A's, but does not lead him to change his conditional degree of belief in any proposition B relative to any of the A's. Then the rigidity conditions hold, and if the A's form a partitioning and his belief function is coherent both before and after the change, his new belief function is determined: It is as in (K). On this, Levi [11] comments that coherence is a static condition, which may be met by the agent's beliefs both before and after the passage of experience even though (K) fails: Both belief functions may be coherent, and we may indeed have $p_1(A_1), \ldots, p_1(A_n)$ as his new degrees of belief in the A's which form a partitioning, but nothing in the coherence conditions forces the function p_1 to be as in (K). (I *think* that is what he is saying in the first part of part II of [11].) Now that is true enough, *but if the rigidity conditions also hold* (as they must, if K is to be applicable), *then it is a matter of logic and high-school algebra that K holds too.*

Starting on p. 199, Levi [11] considers a situation in which the agent is sure that a certain urn, from which he draws a disc, is composed of blue discs together, perhaps, with some green ones. Let A_1 be the proposition that the disc drawn is blue, and A_2 that it is green. These two A's form a partitioning. Now Levi considers three cases:

(1) The agent fails to observe the disc, but somehow or other (perhaps because of a change in his blood chemistry) comes to have new degrees of belief $p_1(A_1)=.9$, $p_1(A_2)=.1$ in the A's anyway.

(2) The agent observes the disc 'and in response admits... $[A_1]$... into his evidence'. I take it that this implies a change in the agent's degree of belief in A_1 to a new value of $p_2(A_1)=1$, although I am not sure. 'Acceptance as evidence' is a technical term in Levi's account ([10], pp. 28–9) of rational belief.

(3) The agent observes the disc (presumably, under less than ideal lighting conditions, or through sunglasses) and in response his degrees of belief in the A's change to new values, $p_3(A_1)=.9$, $p_3(A_2)=.1$.

Now let's consider these cases in turn and, while we are at it, let us consider what Levi says about them.

Case 1: The agent does not observe the disc, but for some reason comes to have new degrees of belief in the A's anyway. His new belief function is p_1, which reflects these changes and perhaps others as well. (Some other changes will be required if p_1 is to be coherent, unless the set of propositions about which the agent has beliefs is very simple indeed.) *Questions:* (a) Are the agent's new judgements, represented by p_1, rational? (b) Was the shift from p_0 to p_1 rational? (c) Was it irrational *of the agent* to change from p_0 to p_1? Levi distinguishes (b) from (c) in a way that I do not: As far as I can see, rationality of a set of beliefs or of a change in beliefs is always relative to an agent at a time – unless you can manage to give a thorough description of the mental state and the experiences which, in that or any other agent, support a verdict of 'rational' or 'irrational' on the beliefs or on the change. Certainly, I see no hope of judging from an examination of p_0 by itself whether it is a rational set of partial beliefs, or of judging from an examination of p_0 and p_1 whether a shift from p_0 to p_1 is a rational one! I still do not want to call the agent himself irrational when, e.g., his belief function p is such that $p(A)>p(B)$ even though one can show by a complicated proof that $A \vdash B$. In that case I would call p irrational even though the agent may be highly rational, as human beings go. Similarly when the agent's belief function p is coherent, but reflects his missing of some abstruse sensory cue that is available to him, I would not call the agent irrational, but would want to make some deprecatory remark about p as a belief function for him at that time, or perhaps as a belief function for anybody who is in the relevant respects as the agent is at that time. The deprecatory remark need not contain the word 'irrational'. It would be more helpful to be more precise, e.g. to remark that the agent seems to have missed a certain sensory clue; and similarly in the first case, it would be more helpful to point out that, unbeknownst to the agent, his belief function p is incoherent, than to give the less informative description, 'It is irrational.' And in the case of the missed sensory clue, it would probably be misleading to characterize the situation even vaguely as one in which p is *irrational* for the agent at that time; 'insensitive' might be a better term. My general evolutionary view of our changing conceptions of what constitutes sensible decision-making and sensible ways of adjusting one's beliefs to one's experiences makes me chary of talk of 'rationality', as does the weight I assign to tacit elements in the basis for our judgements of such matters.

To get back to case 1: I simply do not know what to say about the agent's shift from p_0 to p_1 in the absence of any information about what caused it. If, as Levi suggests, the cause was a change in his body chemistry quite unrelated to any ratiocination or perception on the agent's part, then I would call the shift mysterious and a-rational. Furthermore, if I were informed that p_0 was quite a sensible belief function for the agent at the time, I would deplore the shift to p_1, and view p_1 as unsatisfactory; I would not even balk, here, at calling p_1 irrational for that agent at that time, and saying that the agent himself acted irrationally.

Case 2: The agent observes that A_1 is true. Then his degrees of belief in the A's change to $p_2(A_1)=1$, $p_2(A_2)=0$. (Let me discuss this particular reading of Levi; if it is a misleading reading, Levi will correct me.) If this is indeed the whole of what he takes himself to have learned from the observation, then (as far as we can tell from *this* story), nothing has changed his degrees of belief in any propositions conditionally on A_1. We then have $p_2(B/A_1)=p_0(B/A_1)$ and since $p_2(A_1)=1$, the elementary probability calculus gives

$$p_2(B)=p_0(B/A_1)$$

for all B's for which p_1 is defined. (He no longer has any conditonal degrees of belief relative to A_2, i.e., $p_2(B/A_2)=0/0$ – undefined. Alternatively and, by and large, harmlessly and uselessly, we might stipulate that $p_2(B/A_2)=p_0(B/A_2)$.) But maybe there is more to the story. By augmenting the story in various ways, we might make it plausible that p_2 is not as described above. Thus, we might imagine that seeing the disc to be blue sets off or is anyway accompanied by some train of ratiocination which makes him doubt the soundness of his original judgement, that $p_0(B/A_1)=.7$, say, for some particular B. Perhaps he takes $p(B/A_1)=.4$ to be a sounder judgement, and in consequence forms p_2 by a process more messy than conditionalization of p_0 relative to A_1. Then I would say that conditionalization would have been inappropriate after all – just as I would call (K) inappropriate if one of the rigidity conditions failed. This inappropriateness is of the clear-cut kind, where the new belief function would be incoherent; if the chain of ratiocination leads the fellow to the attitude that $p_2(B/A_1)=.4\neq p_0(B/A_1)$ and he forms p_2 by conditionalizing p_0 relative to A_1, then p_2 will not be single-valued! (Clearer, perhaps, to put it this way: the change in conditional probability of B relative to A_1 simply *prevents* him from forming a new, coherent belief function by conditionalizing p_0 relative to A_1.)

This is the sort of thing I find unsatisfactory about Levi's examples: He tells a little story about what the agent has or has not observed, and tells you that he has changed his degrees of belief in the A's in a certain way. Now it is natural to suppose that unless the story calls out for completion (e.g. by 'the change was caused by an attack of indigestion' or 'the change was caused by a brilliant new idea that just happened to come to him then' in case 1) then we have been told everything relevant; for we are asked, on the basis of the story, to say whether the change was a reasonable one – small wonder that we suppose we have been given all the relevant facts. Especially in case 2, when we are told that the agent observes the disc and consequently accepts A_1 as evidence, it is natural to suppose that nothing else, e.g. ratiocination, is going on that might bear on the appropriateness of the simplest change, from p_0 to $p_2=p_0(\ /A_1)$. But then Levi faults that simple change by pointing out that it is not the only one that will eventuate in a coherent function p_2; and indeed, there are plausible stories one can tell when conditionalization relative to A_1 would be inappropriate, e.g. one could continue the story of case 1 by saying, 'and the agent also observed that B is true', whereupon (if *that is* the whole story) we should have $p_2(C)=p_0(C/A_1\,\&B)$ for all C for which p_0 is defined.

The real use Levi makes of this ploy – that the truth need not be the whole truth – is to set a scene in which I seem committed to the truth of the antecedent of this conditional:

> if coherence is the only obligation imposed on a rational agent in assigning probabilities to propositions, acquisition of new evidence dictates very little to him regarding how he is to revise his probabilities. New evidence is virtually irrelevant to the revision of probabilities.

But as I have said repeatedly above, I do not take coherence to be sufficient for reasonableness of someone's beliefs. The situation is rather that coherence is the only condition I can think of which is *necessary* for reasonableness of sets of beliefs in all cases – regardless of what the agent's mental and experiential condition may be. In case 1, where the story is clearly incomplete, the only thing I can tell you about the reasonableness of the agent's beliefs without hearing the rest of the story is that if they are incoherent, they are unreasonable. Tell *me* more and maybe I shall be able to tell *you* more; but please do not imagine that because I cite no conditions beyond coherence as necessary in all cases, I think that coherence is sufficient in any case! Now, finally, let us look at

Case 3: The agent observes the disc. He admits no new proposition into his evidence but changes to new degrees of belief $p_3(A_1)=.9$, $p_3(A_2)=.1$ in the A's. Levi says ([11], p. 203)

> One might be tempted to suggest that perhaps case (3) is not completely described.

Indeed I am sorely tempted, and herewith yield to the temptation, although not (as Levi notes) by continuing the story, 'These degrees of belief in the A's are simply $p_3(A_1)=p_0(A_1/\text{the chip appears to be blue})$ and $p_3(A_2)=p_0(A_2/\text{the chip appears to be blue})$, where "The chip appears to be blue" is a phenomenological report which describes the relevant aspects of the agent's observation.' My objection to *this* move is that I see nothing in the statement, 'The chip appears to be blue', to warrant $p_0(A_1/\text{the chip appears to be blue})=.9$ instead of .7, say – the statement is too vague for that.

Rather, I would complete the story by saying, 'And nothing in what the agent saw or thought on that occasion moved him to revise his degrees of belief in any proposition relative to either of the A's.' Indeed, if you had to guess the rest of the story, you would guess some such thing simply because you expect the narrator to have said something about it, if in fact something did move the agent to change some of those conditional probabilities. Anyway, with the story completed in that way, p_3 *must* come from p_0 via (K) if both belief functions are to be coherent.

My general point is this: To judge the soundness of a shift from p to p' we must look at more than the kinematics; we must not only look at the two belief functions, and their differences; we must also inquire into the forces which prompted the change – the dynamics of the situation. This is another way of saying that although coherence is the only condition I can formulate that seems necessary for reasonableness of belief functions in all cases, each particular case must be examined with an eye to the agent's particular situation. In particular, the question, whether (K) is an appropriate kinematical relation in a particular case must be answered by finding out what is moving the agent in that case. If he has just examined a chip in sunlight while wearing his sunglasses and he offers 9:1 but no longer odds on the chip's being blue, and wants to *be given* those same odds on its being green, then it is *plausible* to suppose that the rigidity conditions are met, and that (K) is appropriate with .9 and .1 substituted in the right places. Of course, one may be wrong, and one can describe various observations one might make that would convince one of that.

Of course one can, in a rather question-begging way, give necessary and sufficient conditions for appropriateness of (K): it is necessary and sufficient that it be sensible for the partitioning conditions to hold of the agent's original belief function, and that the rigidity conditions hold between his old and new belief functions – in view of whatever the agent's situation happens to be at the time of the change. Similar question-begging conditions can be given for conditionalization relative to E in response to an observation: Everything relevant to the agent's beliefs that can be said about the observation is expressed by the statement, 'The agent has, sensibly, come to have degree of belief 1 in E.' But none of this is much real help. What is of help is our practice, of using conditionalization – or, perhaps, (K) – in various cases more or less tacitly seen to be the right ones. Any *one* of these cases can be discussed non-circularly with some profit – in terms of facts about *that* case which perhaps are generally seen as supporting or undermining the thesis that conditionalization is indeed sensible, there. But the genus of cases in which conditionalization is

appropriate is one I do not know how to characterize clearly and non-circularly. Ditto for cases in which (K) is appropriate.

IV. THE AUTHORITY OF REASON

According to Aristotle[7] humans of both sexes are rational animals, but in women, reason is without authority. ('I contradict myself? Very well, I contradict myself.')

In what, if anything, does the authority of reason consist? In particular, does not my evolutionary account of our canons for reasonableness of deliberation and belief tend to undermine the authority of reason? What reason is there to think the current vogue any better than the one before that? None – if you require a reason to be certifiable as such by certifiably eternal canons of rationality that transcend all vogues. But if you are less exacting, the answer is not far to seek: The reasons in favor of the current vogue are the considerations which prompted us to see it as satisfactory, or as more satisfactory than any other canons that were in the running at the time the current vogue swept the field. To inquire into the source of the authority of reason is to ask the question wrong way round: What we currently take to be reason has the status of reason in our eyes because of the authority it has come to have, for us. Granted, we can rebel against the authority of what currently passes for reason; but remember that man is on the whole a reason-seeking animal, shunning frank irrationalism; and remember that not just anything can pass as a serious candidate for reason, given the stage of cultural evolution we happen to be at. The evidence for the thesis that man is persistently a reason-seeker is a bit sobering; part of the evidence is that he has resorted to astrology and the like, so strong is his desire for rational-seeming authority over deliberation. Then my kind of pragmatism sees man not simply as an agent, but as a would-be rational agent: A deliberator. I am sure that what we now take for rationality in deliberation will appear rather seriously flawed, in the not-too-distant future. We are in something like the situation of the gambler who, when it was pointed out to him that the roulette game he was playing was biased against him said, 'I know, but it is the only wheel in town.'

We cannot help but act, we cannot help but deliberate; when our deliberations are coherent, we act as if we had probabilistic degrees of belief in various propositions. Robots might be programmed to have all sorts of bizarre probabilistic belief functions, and to change them in all manner of bizarre ways in response to their sensory inputs. We are not robots; but if you wish to think of us as robots or as like robots in important respects, note that we are robots of a particular narrow range of different designs. For us, some belief functions are so far out, and some modes of change are so far out as to be inaccessible. When I walk through a soaking rain I can no more believe that it is a fine, sunny day than I can fly; and in general, in any evidential situation, I have little latitude in what I am free to believe. Indeed, I seldom *choose* my beliefs – I generally simply find I *have* them, willy nilly, and generally, that is no defect. Nor is ratiocination an exception to this rule. Perhaps I am free to deliberate or not, but when I elect to deliberate I engage in an activity which, to the extent that it is successful, will pretty much force certain partial beliefs upon me, even though I may not be able to quote explicit rules that I am following.

Similarly for observation and experiment, which are ingredients in (anyway, raw materials for) deliberation. Part of the business of learning English is the matter of learning that for certain questions there are certain fairly definite ways in which one can become an authority. Take the statement, 'The sun is shining'. Part of knowing one's way about in English is knowing that by going outdoors and opening one's eyes one will have a sensibly arrived-at degree of belief in that statement which will in all probability be close to 0 or to 1. Sense-perception has its authority, too.

Man is an explaining animal as well as a deliberating animal, and the two activities are intertwined. Part of reasonableness is a readiness to survey one's own beliefs with an explanatory eye. To some extent, this is a matter of keeping track of what led us to believe what. Example: we do well to keep some track of the kinematical history of our beliefs. Perhaps I have a succession p_0, p_1, p_2, p_3 of belief functions, where each arose from its predecessor via (K). Suppose the bases of the successive changes were A, $-A$ (I got from p_0 to p_1 by changing my degree of belief in A), B, $-B$ (to get from p_1 to p_2) and finally C, $-C$ – where in each case the rigidity conditions were met. So far, so good. But I might well wonder whether I would have got p_3, starting with p_0, by a single application of (K) if I had made all three changes at once. If not, I might do well to think the whole thing over, and make some adjustments. (What adjustments?

That depends on what the propositions A, B, C are, and on what my situation is.) There is some interesting lore, here. Suppose the changes were, $p_0(A)\rightarrow p_1(A)=a$; $p_1(B)\rightarrow p_2(B)=b$; $p_2(C)\rightarrow p_3(C)=c$ (with $1-a$, $1-b$, $1-c$ as the corresponding values for the denials, of course.) If the rigidity conditions were met and A, B, C were independent relative to p_0 in the sense that $p_0(A\&B)=p_0(A)\,p_0(B)$, $p_0(A\&B\&C)=p_0(A)\,p_0(B)\,p_0(C)$, etc., then those three must be independent relative to the other three belief functions as well, and we can get directly from p_0 to p_3 via (K), taking $n=8$, $A_1=A\&B\&C$, $A_2=A\&B\&-C$, ..., $A_8=-A\&-B\&-C$, and $p_3(A\&B\&C)=abc$, $p_3(A\&B\&-C)=ab(1-c),\ldots$, $p_3(-A\&-B\&-C)=(1-a)\ (1-b)\ (1-c)$. (Of course, to apply formula (K) above, the subscript '1' should be changed to '3' throughout.)

But if A, B, C are not independent in this way, this cannot be expected to work. In that case, if the new values are a, b, c again, there is considerable latitude within which the new values of $A\&B\&-C$, etc., might lie. And by *first* changing degree of belief in A to a, *then* changing that in B to b, and only then changing that in C to c, we may well be in a position where, although we do have c as the value of $p_3(C)$, we do not have $p_3(B)=b$ (rather, $p_2(B)=b$, and the change $p_2(C)\rightarrow p_3(C)=c$ may have induced a change from $p_2(B)=b$ to $p_3(B)$=something else) and we do not have $p_3(A)=a$. This is a reflection of the fact that since no special logical or probabilistic relationship is postulated, among A, B, and C, they do not themselves form a basis for a single jump from p_0 to p_3. The relevant basis would consist of such of the eight conjunctions of those letters and denials of those letters as have positive probability, relative to p_0. It would be idle to seek a rule telling us how to assign new probabilities to those conjuncts, given only the new probabilities of A, B, and C: Any such rule would work only in special cases, e.g. the rule that we ought to assign abc, $ab(1-c),\ldots,(1-a)\,(1-b)\,(1-c)$ respectively to $A\&B\&C$, $A\&B\&-C,\ldots,-A\&-B\&-C$ would be appropriate if and only if A, B, and C are fully independent, relative to p_0.[8]

One final comment. Part of the business of keeping an explanatory eye on our beliefs (our own and those of other agents) is, as I say, the matter of keeping track of the forces that initiate kinematical changes – changes via K. In one's own case, one may have a clear notion of what the basis was, in a certain case, i.e., of what the propositions $A_1,\ldots,A_n$ were, in which the change initiated. Now Levi falsely says that (aside from such private insights, I suppose),

> No shift to some nonextreme probability value can be marked off as initial by appeal to the value reached. Jeffrey lacks a way of identifying initial shifts to be used to justify other shifts ([11], p. 204).

It is the conjunction of these two statements I am calling false: the first conjunct is true enough, but Levi seems to suggest that its truth explains the truth of the second conjunct, which it cannot, since the second conjunct is false. Suppose p_0 and p_1 are successive belief functions and suppose that, in the course of examining them with an explanatory eye, I notice the existence of a partitioning $A_1, \ldots, A_n$ which meets the rigidity conditions. Then p_1 is related to p_0 as in (K) – assuming that both functions are coherent. I will then do well to explore the hypothesis that the changes $p_0(A_1) \to p_1(A_1), \ldots, p_0(A_n) \to p_1(A_n)$ were the initiating changes. The hypothesis might be false, of course; in testing it, I would want to think about the propositions $A_1, \ldots, A_n$ to see on what sensible basis the agent in question might have changed his degrees of belief in them. But perhaps I shall notice that they are all reports on the color of some disc which, at the time of the change, the agent was regarding through sunglasses. I would then be pretty confident of the explanatory hypothesis, that the A's were the basis for the change. On the other hand, I might notice that by choosing $A'_1 = A_1 \vee A_2$, $A'_2 = A_3 \vee A_4 \vee A_5$ (suppose $n=5$), I again get a basis, consisting of the simpler set, A'_1, A'_2. In such a case I would become interested in the hypothesis that the fellow was only concerned with a rather coarse color observation; that he has two categories, not five. Of course it *might* be that he had five categories in mind, and as luck would have it, the change worked out precisely as it would have done if he had had only the two in mind; but in this as in any case of explanation we are attracted by simplicity, and in the absence of evidence to the contrary might opt for the simpler explanation as likelier to be true and, at the same time, making a tighter explanatory package.

V. CONCLUSION

Wolfman and Dracula live in disjoint, unreal worlds; when Dracula meets Wolfman in the movies, the two fantasies generally destroy each other – Wolfman cannot breathe the musty air of Castle Dracula, nor can Dracula survive the winds of Wolfman's moors. Now that Levi and I have confronted each other, I see that my title may have been inappropriate: perhaps our worlds can mix, without reducing each other to absurdity. Perhaps, indeed, the disparity between them is the disparity of viable accounts of different parts of the real world, not that of incompatible fantasies.

Before the confrontation my dissatisfaction with Levi's positive proposals – his Bayesian account [10] of the business of accepting and rejecting hypotheses – was centered on the very notions of acceptance and rejection. I observed [9] that while he gave methods for deciding which (if either) of the two acts *accept H* and *reject H* one ought to perform, he provided no account of how one is to go about performing those acts. The notions of belief and disbelief are familiar enough but, I find, unclear. In contrast, I find the notion of subjective probability, for all its (decreasing) unfamiliarity, to be a model of clarity – a clarity that it derives from its association with the concepts of utility and preference within the framework of Bayesian decision theory. But as I understand him, Levi now takes the notions of acceptance and rejection to be in need of clarification, which he seeks to provide by giving them a place in an account of inquiry of the sort discussed in his contribution to this conference. I take it that this account is still in the programmatic stage, and I await its elaboration with interest. It would be good to make systematic sense of our talk of acceptance and rejection of hypotheses, and the *prima facie* difficulty of giving a coherent account is no reason for thinking that no such account will be forthcoming. Meanwhile, in trying to square theory with practice, *I* continue to avoid talk about knowledge and acceptance of hypotheses, trying to make do with graded belief, as in [8]. That, too, is programmatic, but it is the program that strikes me as most promising. For both programs I think the crucial task is that of squaring a Bayesian account of practical deliberation, which both Levi and I seem to accept, with the facts about how theory is generally thought to legitimately impinge upon practice. 'My kind of pragmatism' would be shattered by a wedge driven between theory and practice.[9]

University of Pennsylvania

BIBLIOGRAPHY

[1] G. E. Bates, *Probability*, Reading, Mass., 1965.
[2] Ethan Bolker, *Functions Resembling Quotients of Measures*, Dissertation, Harvard University, 1965.
[3] Ethan Bolker, 'Functions Resembling Quotients of Measures', *Trans. American Mathematical Society* **124** (1966) 292–312.
[4] Ethan Bolker, 'A Simultaneous Axiomatization of Utility and Subjective Probability', *Philosophy of Science* **34** (1967) 333–40.
[5] M. Friedman and L. J. Savage, 'The Utility Analysis of Choices Involving Risk', *Journal of Political Economy* **56** (1948) 279–304.
[6] Ian Hacking, 'Slightly More Realistic Personal Probability', *Philosophy of Science* **34** (1967) 311–25.
[7] Richard C. Jeffrey, *The Logic of Decision*, New York 1965.
[8] Richard C. Jeffrey, 'Probable Knowledge', in *The Problem of Inductive Logic* (ed. by I. Lakatos), Amsterdam 1968, pp. 166–80.
[9] Richard C. Jeffrey, 'Review of I. Levi's *Gambling With Truth*', *Journal of Philosophy* **65** (1968) 313–22.
[10] Isaac Levi, *Gambling With Truth*, New York and London 1967.
[11] Isaac Levi, 'Probability Kinematics', *British Journal for the Philosophy of Science* **18** (1967) 197–209.
[12] Howard Riaffa, *Decision Analysis: Introductory Lectures on Choices under Uncertainty*, Reading, Mass., 1968.
[13] F. P. Ramsey, *The Foundations of Mathematics and other Logical Essays*, London and New York 1931.
[14] Leonard J. Savage, *The Foundations of Statistics*, New York 1954.
[15] Patrick Suppes, 'Rational Changes of Belief', in *The Problem of Inductive Logic* (ed. by I. Lakatos), Amsterdam 1968, pp. 186–8.
[16] Richard Todd, 'Where Money Grows on Decision Trees', *New York Times Magazine*, November 10, 1968.
[17] I. Todhunter, *A History of the Mathematical Theory of Probability*, New York 1949.

REFERENCES

1 Those issues were indeed greatly clarified during the conference; but it seems best to publish the paper just as it was presented at the conference even unto the title – except for minor corrections throughout, and for some brief concluding remarks in Section V.

2 'My' account is in [7]; it is a modification and, as I see it, an improvement of the accounts of Ramsey [13] and Savage [14].

3 See [2], [3], and [4]. Bolker's work was prior to mine: A case in which pure mathematics had an unexpected application.

4 For some striking illustrations of the flexibility of the Bayesian framework, see [5], where some puzzling and rather complex behavior is rationalized via a simple hypothesis about the shape of the utility curve for income.

5 This is the celebrated argument from coherence: If his degrees of belief do not satisfy the probability axioms there will be a set of bets, each of which looks fair to the agent, but on which he will with logical necessity suffer an overall loss. See [7], p. 49 and references given there.

6 The room in which this was read had some 40 people in it. Imagine that the date was September 25.

7 *Politics* 1260^a 13.

8 One may take this as a reply to the objection, 'One difficulty immediately leaps to the eye…', just below the formula in [11], p. 205.

9 In [10], p. 13, Levi observes that his "critical cognitivism renders asunder, at least partially, what many philosophers have endeavored to join together – theoretical and practical wisdom".