3. Pollsters trying to predict the outcomes of national elections always use stratified random samples. Why do you suppose that they do so? If you were conducting such a poll, how would you choose your strata?

Suggested readings

Morris J. Slonim, *Sampling in a Nutshell* (New York: Simon and Schuster, 1960), pp. 1–53.

The following is recommended for the advanced student:

Herman Chernoff and Lincoln E. Moses, *Elementary Decision Theory* (New York: John Wiley & Sons, Inc., 1959).

VI

Coherence

VI.1. INTRODUCTION. The concepts of epistemic and inductive probability were introduced in Chapter I as numerical measures grading degree of rational belief in a statement and degree of support the premises of an argument give its conclusion. In Chapter V we encountered a mathematical characterization of probabilities and conditional probabilities. Why should epistemic and inductive probabilities obey the mathematical rules laid down for probabilities and conditional probabilities? One reason that can be given is that these mathematical rules are *required* by the role that epistemic probability plays in rational decision.

This sort of argument can be made at various levels, depending on what simplifying assumptions are made. The discussion can be elementary and transparent, in the sort of classical gambling situation discussed in V.6. Here we assume that all that is at issue in the decision problem is money. We also assume that money has constant marginal utility across the range of stakes at issue; that is, an extra dollar counts for as much whether it is added to big winnings or big losses. Finally, we assume that if the bettor takes several individual bets as fair (favorable, unfavorable), he takes the result of making them all together as fair (favorable, unfavorable). Under these assumptions we can show that if a bettor violates the rules of the probability calculus, he can have a *Dutch Book* made against him; that is, a clever bookie can make a series of bets with him, all of which he considers fair or favorable, such that he suffers a net loss *no matter what happens regarding the propositions he is betting on.*

The assumptions for a Dutch Book argument can plausibly be held to be true (or approximately true) for typical monetary gambles with small stakes, but each of them breaks down when we consider the problem of rational decision more globally. These considerations lead to a deeper level, which involves the theory of utility and culminates in an analysis that shows how coherent systems of preference can always be represented as having come from probability and utility, with preference going by expected utility. In this chapter we will start with the simpler situation and end with a sketch of the leading ideas of the deeper results.

VI.2. THE PROBABILITY CALCULUS IN A NUTSHELL. In discussing these questions it will be useful to have as concise a characterization of the mathematical conception of a probability as is possible. Here is the classic one, due to Kolmogorov:

Definition 14: A probability (on statements) is a rule assigning each statement, S, a unique probability, Pr(S), such that:

a. No probability is less than zero.
b. If T is a tautology, $\Pr(T) = 1$
c. If P; Q are mutually exclusive, then $\Pr(P \vee Q) = \Pr(P) + \Pr(Q)$.

Let us see how this brief characterization yields the longer one of Chapter V. Rules 1 and 4 of Chapter V are explicitly contained in Definition 14 as 14b and 14c. We saw in Chapter V that these two rules yield the negation rule. Since $P \vee {\sim}P$ is a tautology and $P; {\sim}P$ are mutually exclusive,

$$\Pr(P) + \Pr({\sim}P) = 1$$

Rule 5: $\Pr({\sim}P) = 1 - \Pr(P)$

Since the denial of a contradiction is a tautology, 14b and Rule 5 show that if C is a contradiction:

$$\Pr(C) = 1 - 1$$

Rule 2: $\Pr(C) = 0$

It is not so obvious that Rule 3 ("if two statements are logically equivalent, they have the same probability") is a consequence of Definition 14, but it is. Suppose P is logically equivalent to Q. Then $P; {\sim}Q$ are mutually exclusive for ${\sim}Q$ is true when P is false and false when P is true. (If the foregoing statement is not obvious, review what logical equivalence means and prove it.) By the same token, $P \vee {\sim}Q$ is a tautology. So by 14b and 14c: $\Pr(P) + \Pr(-Q) = 1$. Using the negation rule:

$$\Pr(P) + 1 - \Pr(Q) = 1$$
$$\Pr(P) = \Pr(Q)$$

Rule 3: If P and Q are logically equivalent, then $\Pr(P) = \Pr(Q)$

We have now only to show that Definition 14 restricts the possible probability values to the range from 0 to 1, for everything else in Chapter V was shown to follow from what we have here developed. We will demonstrate this by showing something much more general and interesting. First, a few preliminaries.

Definition 15: Q is a logical consequence of P just if Q is true in every case in which P is true.

So, for example, R is a logical consequence of R & S, as is S; and $R \vee S$ is a logical consequence of R and of S.

	R	S	R & S	$R \vee S$
Case 1:	T	T	T	T
Case 2:	T	F	F	T
Case 3:	F	T	F	T
Case 4:	F	F	F	F

Notice that a tautology is a logical consequence of everything, since a tautology is true in all cases. And everything is a logical consequence of a contradiction, since a contradiction is never true. We will now show the *Logical Consequence Principle*:

If Q is a logical consequence of P, then $\Pr(Q)$ must be at least as great as $\Pr(P)$.

Then, of course, every probability must fall in the interval from 0 to 1.

If Q is a logical consequence of P then either P and Q are true in exactly the same cases or Q is true in the cases where P is, plus some extra cases. In the former instance P and Q are logically equivalent and thus have the same probability. Let us, then, look at the latter. For example:

	P	Q	$(Q \,\&\, -P)$
Case 1:	T	T	F
Case 2:	F	F	F
Case 3:	T	T	F
Case 4:	T	T	F
Case 5:	F	T	T
Case 6:	F	F	F
Case 7:	T	T	F
Case 8:	F	T	T

Here Q is a logical consequence of P. It is true in the cases where P is [1, 3, 4, 7] plus some extra ones [5, 8]. At the right is the statement,

$Q\&{\sim}P$, which is true in *just these extra cases* and false otherwise. Notice that (Q&~P) *and* P *are mutually exclusive.* Next notice that Pv(Q&~P) *is logically equivalent to* Q since $(Q\&{\sim}P)$ adds just the required extra cases to P. Then:

$$\Pr(Q) = \Pr[Pv(Q\&{\sim}P)] = \Pr(P) + \Pr(Q\,\&\,{\sim}P)$$

Since $\Pr(Q\&{\sim}P)$ is at worst zero, $\Pr(Q)$ must be at least as great as $\Pr(P)$.

This result completes the argument that Definition 14 captures the probability calculus as developed in Chapter V. It is, however, of more than incidental importance. It assures the *probabilistic validity of deductive argument.* If

$$\begin{array}{c} P_1 \\ P_2 \\ \cdot \\ \cdot \\ \cdot \\ P_n \\ \hline C \end{array}$$

is a (truth-functionally) deductively valid argument, then the conclusion, C, is a logical consequence of the conjunction of the premises P_1 & P_2 . . . & P_n. Our result shows that our conclusion *must be at least as probable as the conjunction of the premises.*

Think of the disastrous consequences if this were *not* true! We could have good reason for believing that all the premises of an argument are true, deduce the conclusion from them, and *not* have equally good reason for thereby believing the conclusion. Under such circumstances it would be hard to imagine a rationale for applying deductive logic.

The probabilistic validity of deductive argument provides a justification for applying deductive logic to situations where we are entitled to assign high probability to the conjunction of the premises but are not entitled to be certain of their truth. There are strong arguments to the effect that this covers almost all applications of deductive logic to empirical knowledge.

Exercises:

1. Show:
   a. If P is a logical consequence of Q and Q is a logical consequence of P, then P is logically equivalent to Q.

   b. If P is a logical consequence of Q and Q is a logical consequence of R, then P is a logical consequence of R.
   c. If P is logically equivalent to Q, then P and Q are logical consequences of the same statements and have as logical consequences the same statements.

2. Show that if R is a logical consequence of P and of Q and R has as a logical consequence PvQ, then R is logically equivalent to either P or to Q or to PvQ.

3. Study the following truth table:

	P	Q	R	${\sim}P\&Q\&R$	$P\&Q\&{\sim}R$	$({\sim}P\&Q\&R)v(P\&Q\&{\sim}R)$
Case 1:	T	T	T	F	F	F
Case 2:	T	T	F	F	T	T
Case 3:	T	F	T	F	F	F
Case 4:	T	F	F	F	F	F
Case 5:	F	T	T	T	F	T
Case 6:	F	T	F	F	F	F
Case 7:	F	F	T	F	F	F
Case 8:	F	F	F	F	F	F

Then:
   a. Show that for *any* case in *any* truth table you can construct a sentence which is true in that case and false in all other cases.
   b. Show that for any *set* of cases in *any* truth table you can construct a sentence true in those cases and false in all other cases.

VI. 3. *THE LOGICAL CONSEQUENCE PRINCIPLE ALONE IS NOT ENOUGH. The Logical Consequence Principle is both plausible and powerful. It is hardly open to dispute that if Q is true in every case in which P is, Q must have at least as good a chance of being true as P. We should be pleased that Definition 14 of the previous section leads to this result; we should have been dismayed if it had not.

It is interesting to ask here how far the Logical Consequence Principle will take us toward mathematical probability as characterized by Defi-

*This section may be omitted without loss of continuity.

nition 14. It leads immediately to the principle that logically equivalent statements have the same probability since P and Q are logically equivalent just in case each is a logical consequence of the other. It leads to the fact that there must be a maximum probability value, shared by all tautologies, and a minimum probability value, shared by all contradictions. But even if we arbitrarily choose 1 as the maximum and 0 as the minimum probability, the logical consequence principle does not thereby lead to the additivity of probabilities of mutually exclusive sentences, that is:

14 c. If P; Q are mutually exclusive, then $\Pr(P \vee Q) = \Pr(P) + \Pr(Q)$

This can be most easily seen by considering a new quantity, *plausibility*, which is defined in terms of probability as on the graph in Figure 4. (To find a statement's plausibility from its probability, first find its probability on the horizontal axis, go straight up to the curve and straight over to its plausibility on the vertical scale.) Notice that the curve defining plausibility is so drawn that greater probabilities lead to greater plausibilities and greater plausibilities arise only from greater probabilities. That is:

Probability $(A) >$ Probability (B)
if and only if
Plausibility $(A) >$ Plausibility (B)

The upshot of this is that if we were to arrange some statements in order of increasing plausibility, we would place them *in the same order* as we would if we were arranging them in order of increasing probability. A short way of saying this is to say that probability and plausibility are *ordinally similar.*

It is now easy to see that plausibility must satisfy the Logical Consequence Principle. If Q is a logical consequence of P then the probability of Q must be at least as great as the probability of P (since we showed, in the last section, that probability satisfies the logical consequence principle). Since plausibility is ordinally similar to probability, plausibility of Q must be at least as great as plausibility of P.

However, plausibility need not add for mutually exclusive statements. Assume that some statement, P, is as likely as not; so:

$$\Pr(P) = \Pr(\sim P) = \tfrac{1}{2}$$

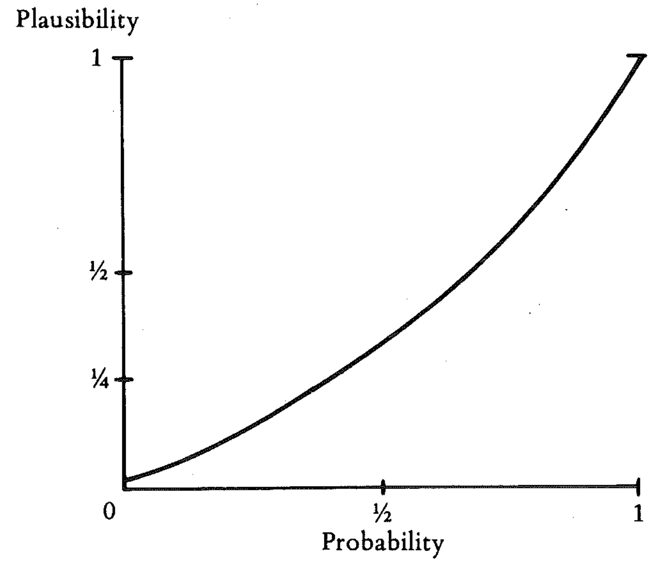

Figure 4

Referring to Figure 4, we see that the *plausibility* values of P and of $\sim P$ are $\frac{1}{4}$. We also see that since $\Pr(P \vee \sim P) = 1$, the plausibility of $(P \vee \sim P) = 1$. So plausibility of $(P \vee \sim P)$ is definitely *not* equal to plausibility of (P) plus plausibility of $(\sim P)$.

Although plausibility satisfies the logical consequence principle, and has a maximum of 1 and a minimum of zero, it is not a probability.

To understand the need for the extra property of additivity which distinguishes a probability we must look to quantitative applications of probability. In the next few sections, we shall consider the oldest and most general quantitative application of probabilities—that of a guide for the intelligent gambler.

Exercises for the advanced student:

1. The following principle has been proposed for any grading of rational degree of belief.

I. For any statements, R, S, T, if R and T are mutually exclusive and S and T are mutually exclusive, then the degree of belief in R is greater than the degree of belief in S, if and only if the degree of belief in $R \vee T$ should be greater than the degree of belief in $S \vee T$.

a. Show that probability satisfies principle I and that any quantity ordinally similar to probability does also.

b. Show that any quantity representing degree of belief which:

i. Satisfies principle I,
ii. Gives logically equivalent sentences the same probability, and
iii. Requires that tautologies have the maximum probability and contradictions the minimum

must satisfy the logical consequence principle.

2. Consider a language containing only two simple statements, P, Q, together with every complex statement which can be built out of them using &, v, ~.

a. Show that every statement in the language is logically equivalent to one of the following sixteen statements:

$P\&{\sim}P$; ${\sim}P\&{\sim}Q$; ${\sim}P\&Q$; ${\sim}P$; $P\&{\sim}Q$; ${\sim}Q$; $(P\&{\sim}Q)\text{v}({\sim}P\&Q)$; ${\sim}P\text{v}{\sim}Q$; $P\&Q$; $(P\&Q)\text{v}({\sim}P\&{\sim}Q)$; Q; ${\sim}P\text{v}Q$; P; $P\text{v}{\sim}Q$; $P\text{v}Q$; $P\text{v}{\sim}P$. (Hint: Look at the truth table for all these statements.)

b. Assuming $\Pr(P\&Q) = .41$; $\Pr(P\&{\sim}Q) = .29$; $\Pr({\sim}P\&Q) = .2$; $\Pr({\sim}P\&{\sim}Q) = .1$; calculate the probability values for each of the foregoing statements.

c. Consider the quantity representing a degree of belief which takes as values just the foregoing probabilities except that $({\sim}P\text{v}{\sim}Q)$ and $({\sim}P\text{v}Q)$ switch values. That is, for any statement R, $D(R) = \Pr(R)$ except that $D({\sim}P\text{v}{\sim}Q) = .71$ and $D({\sim}P\text{v}Q) = .69$

i. Show that D satisfies the Logical Consequence Principle. (Hint: to verify that the Logical Consequence Principle is not violated you need only verify that no sentence of which ${\sim}P\text{v}Q$ is a consequence has a probability greater than .69 and no statement which is a consequence of ${\sim}P\&{\sim}Q$ has a probability less than .71.)

ii. Show that D *violates* principle I. (Hint: ${\sim}P\text{v}Q$ is logically equivalent to $(P\&Q)\text{v}{\sim}P$ and ${\sim}P\text{v}{\sim}Q$ is logically equivalent to $(P\&{\sim}Q)\text{v}{\sim}P$.)

3. Exercises 1 and 2 show that satisfying the Logical Consequence Principle is not a sufficient condition for being ordinally similar to probability. Show that satisfying principle I also fails to be sufficient for ordinal similarity to probability. (This is a *difficult* problem. The required proof can be found in the first of the suggested readings. Conditions which are sufficient to guarantee ordinal similarity to probability may be found in both of the suggested readings.)

Suggested readings

Krantz, Luce, Suppes, and Tversky, *Foundations of Measurement* (New York: Academic Press, 1971).

Richard Cox, *The Algebra of Probable Inference* (Baltimore, Md.: The Johns Hopkins Press, 1961).

VI.4. BETS. A bet on a statement, P, is an arrangement by which the bettor collects a certain sum, a, if P is true and forfeits a certain sum, b, if P is false. The situation can be characterized in a payoff table:

P	Net Gain
T	$+a$
F	$-b$

The total amount involved, $a + b$, is the *stake* and the ratio $\frac{b}{a}$ is the *odds*.

A bet on one statement may also constitute a bet on another statement. In the most trivial case, if two statements are logically equivalent then a bet on one is equally a bet on the other. In the next most trivial case a negative bet on ${\sim}P$ is identical to a positive bet on P.

${\sim}P$	Net Gain	P
T	$-b$	F
F	$+a$	T

If B_1 and B_2 are betting arrangements, their *sum* is an arrangement by which the bettor fulfills his obligations under both B_1 and B_2. For example:

B_1

P	Net Gain
T	a
F	$-b$

B_2

Q	Net Gain
T	c
F	$-d$

Sum of B_1 and B_2

P	Q	Net Gain
T	T	$a + c$
T	F	$a - d$
F	T	$c - b$
F	F	$-(b + d)$

It should be clear that the sum of two bets *need not* be a bet on any given statement. A bet on a statement, S, pays a certain value if S is true and costs a certain value if S is false. There is no differentiation about *ways* that S can be true or *ways* that S can be false. So if in the payoff table for the sum bet, there are at least three different figures under net gain, the sum bet cannot be interpreted as a bet on any statement.

Let us move to a more interesting case of a sum bet being a bet on a statement. Suppose that P and Q are mutually exclusive, that B_1 is a bet on P for stakes $(a + b)$ at odds $\frac{b}{a}$ and B_2 is a bet on Q, for stakes $(c + d)$ at odds $\frac{c}{d}$:

P	Q	B_1	B_2	Sum of B_1 and B_2
T	F	a	$-d$	$a - d$
F	T	$-b$	c	$c - b$
F	F	$-b$	$-d$	$-(b + d)$

Since P and Q are mutually exclusive, there are only three possible combinations of truth values. If the sum bet has a different payoff value in each of the three cases, we know that it is not a bet on any statement. But what if the payoff values in the first two cases are the same (that is $a - d = c - b$)? Then the bettor wins this value in either of these cases, *that is whenever* PvQ *is true* and loses $b + d$ in the third case, *that is when* PvQ *is false.* So if $a - d = c - b$, the sum bet is a bet on PvQ with stakes $(a - d) + (b + d) = (a + b)$ and odds $\frac{b + d}{a - d}$. Under what conditions does this interesting phenomenon occur? It doesn't take much algebra to show that $a - d = c - b$ just in case $a + b = c + d$, that is just in case the stakes of our bets on P and Q are equal. In summary,

> *If* P *and* Q *are mutually exclusive, the sum of bets on* P *and on* Q *at equal stakes is a bet on* PvQ *at the same stakes.*

There is another kind of betting arrangement which is of general interest and which is not a bet on any statement. This is the sort of bet that is called off if certain conditions are not met; call it a *conditional bet.* If the bet is on Q and the conditions to be met are specified by P, then it is called, not surprisingly, a bet on Q conditional on P and gives rise to the following sort of payoff table:

P	Q	Payoff
T	T	a
T	F	$-b$
F	T	0
F	F	0

A little reflection should convince you that many of the betting situations that we get ourselves into are conditional bets. Sometimes we may wish that even more of them were. If so, it should come as good news that we can always construct a betting arrangement conditional on P by a simple hedging strategy.

Consider the sum of two bets, the first being a bet on P&Q and the second being a bet on $\sim P$.

P	Q	P&Q	Bet 1	$\sim P$	Bet 2	Sum of Bets 1 and 2
T	T	T	c	F	$-f$	$c - f$
T	F	F	$-d$	F	$-f$	$-(d + f)$
F	T	F	$-d$	T	e	$e - d$
F	F	F	$-d$	T	e	$e - d$

If we arrange Bet 2 so that our winnings on $\sim P$, e, equal our losses from Bet 1, d, the sum of Bets 1 and 2 will be a bet on Q conditional on P, as follows:

P	Q	Sum of Bets 1 and 2
T	T	$c - f$
T	F	$-(d + f)$
F	T	0
F	F	0

In summary: The sum of two bets, the first on P & Q and the second on $\sim P$ with the winnings on the second being equal to the losses on the first, is a bet on Q conditional on P.

Exercise:

If you bet someone a dollar at 2 to 1 odds that P:

a. What is *your* payoff table for P?

b. What is *your* payoff table for $\sim P$?

c. What is *his* payoff table for P?

d. What is *his* payoff table for $\sim P$?

VI.5. FAIR BETS. Remember from Chapter V that the *expected value* of a betting arrangement is the sum of the quantities obtained by multiplying the payoff in a given case by the probability of that case. For example, the bet:

Bet 1

P	Payoff
T	a
F	$-b$

has an expected value of $a\Pr(P) - b\Pr(\sim P)$ and the betting arrangement:

Bet 2

P	Q	Payoff
T	T	a
T	F	b
F	T	c
F	F	$-d$

has an expected value of $a\Pr(P\&Q) + b\Pr(P\&\sim Q) + c\Pr(\sim P\&Q) - d\Pr(\sim P\&\sim Q)$.

If the expected value of a bet is positive, it is called a *favorable* bet; if negative, it is an *unfavorable* bet; if zero, it is a *fair* bet. Whether a bet is fair, favorable, or unfavorable depends on how the probabilities balance out the odds. Consider Bet 1 on P. It is fair just in case:

$$a\Pr(P) - b\Pr(\sim P) = 0$$
$$a\Pr(P) - b[1 - \Pr(P)] = 0$$
$$a\Pr(P) - b + b\Pr(P) = 0$$
$$a\Pr(P) + b\Pr(P) = b$$
$$\Pr(P)[a + b] = b$$
$$\Pr(P) = \left(\frac{b}{a+b}\right)$$

The quantity $\frac{b}{a+b}$ is called the *betting quotient* for P. So we can say that a bet on P is fair just in case the probability of P equals the betting quotient for P.

Suppose we have fair bets on P and Q:

P	Payoff
T	a
F	$-b$

Q	Payoff
T	c
F	$-d$

Must the sum of these two bets be fair? The sum bet:

P	Q	Payoff
T	T	$a + c$
T	F	$a - d$
F	T	$c - b$
F	F	$-b - d$

is fair if and only if:

$$\Pr(P\&Q)(a + c) + \Pr(P\&\sim Q)(a - d) + \Pr(\sim P\&Q)(c - b) + \Pr(\sim P\&\sim Q)(-b - d) = 0$$

or:

$$a\Pr(P\&Q) + c\Pr(P\&Q) + a\Pr(P\&\sim Q) - d\Pr(P\&\sim Q) + c\Pr(\sim P\&Q) - b\Pr(\sim P\&Q) - b\Pr(\sim P\&\sim Q) - d\Pr(\sim P\&\sim Q) = 0$$

or:

$$a[\Pr(P\&Q) + \Pr(P\&\sim Q)] - b[\Pr(\sim P\&Q) + \Pr(\sim P\&\sim Q)] + c[\Pr(\sim P\&Q) + \Pr(P\&Q)] - d[\Pr(P\&\sim Q) + \Pr(\sim P\&\sim Q)] = 0$$

or:

$$a\Pr(P) - b\Pr(\sim P) + c\Pr(Q) - d\Pr(\sim Q) = 0$$

But since our bet on P is fair, $a\Pr(P) - b\Pr(\sim P) = 0$; and since our bet on Q is fair, $c\Pr(Q) - d\Pr(\sim Q) = 0$. So if bets on two statements are fair their sum is fair.[1]

[1]Note that this is not the only way that the sum bet can be fair. If the expected value of one bet is $+e$ and that of the other is $-e$, then the sum bet will be fair.

The argument is summarized in Figure 5. Each square contains the payoff for one case of the sum bet multiplied by the probability of that case (e.g., in the upper left-hand square $\Pr(P\&Q)(a + c) = a\Pr(P\&Q) + c\Pr(P\&Q)$. Thus, the expected value of the sum bet is just the sum of everything in all the squares. The squares are divided into triangles to suggest a way of adding. The quantities in the lower left triangles are added downward and summed at the bottom of the columns. The quantities in the upper right triangles are added to the right. The sum of the quantities at the bottom of the columns is the expected value of the

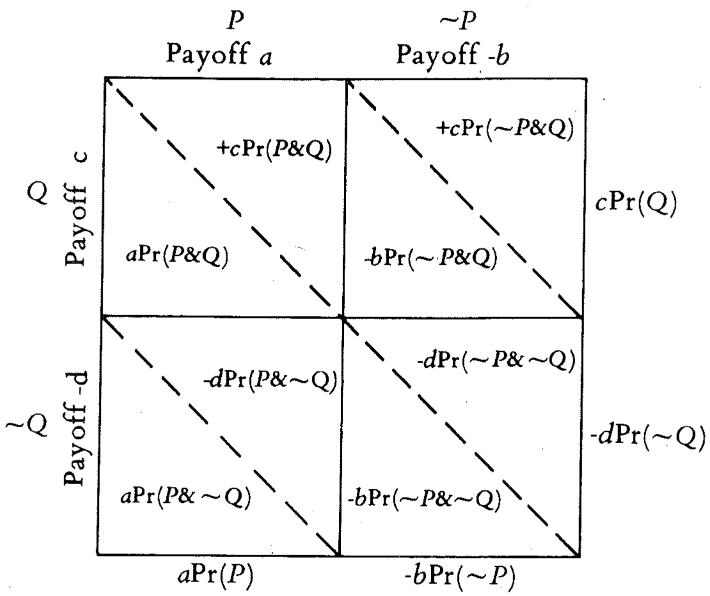

Figure 5

bet on P; the sum of the quantities a the right of the rows is the expected value of the bet on Q. *So we have shown that the expected value of the sum of bets on* P *and on* Q *is the sum of the expected values of those bets.* Of course, then, if two bets are fair, their sum bet is fair.

So far, we have only talked about the sum of two fair *bets*, rather than two fair *betting arrangements*. Remember, a bet on a statement, P, admits of only two possibilities—P is true or P is false—and specifies a unique payoff in each case. As we saw in the last section, the sum of the two bets (on statements) is a betting arrangement, which typically is not a bet on any statement. Thus, we have still to ask whether the sum of any two fair *betting arrangements* is a fair betting arrangement. The

answer, happily, is yes—by the same sort of argument we used in the simpler case. The argument is indicated in Figure 6. Again the lower left triangles are summed downwards and when they are added at the bottom we find that their sum equals the expected value of betting arrangement 1. Likewise the sum of all the contents of all the upper right triangles equals the expected value of betting arrangement 2. But the sum of all the triangles is just the expected value of the sum betting arrangement.

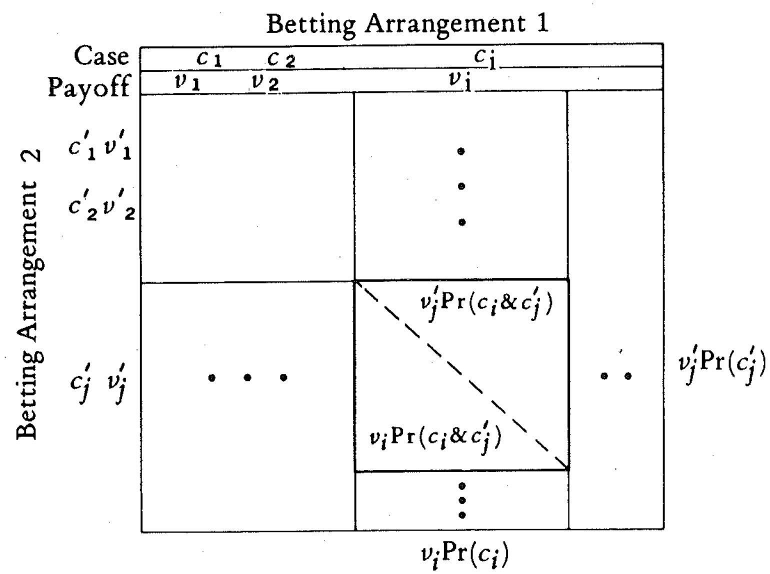

Figure 6

So we know in general that:

> The expected value of a sum of betting arrangements is the sum of the expected values of the individual betting arrangements.

and in particular that:

> A sum of fair betting arrangements is fair.

Perhaps you may think that this is not very surprising and that we have been, perhaps, belaboring the obvious. Well and good! Think then how surprising it would be if this were not true—in fact how *disastrous* it would be for making decisions under uncertainty. We could undertake a series of fair risks and yet have no assurance that the total arrangement was not unfair.

What is surprising is not that probabilities, as measures of belief, lead to such well-behaved concepts of expected value and fair bets but that probabilities are the *only* kinds of measures of belief which will do so.

Imagine us using some *new* measure of belief, call it plausibility (Pl), to take the place of probability. Again, we will call a bet on a statement S

S	Payoff
T	a
F	$-b$

fair just in case the betting quotient for S, $\frac{b}{a+b}$, equals the plausibility of S.

Now remember that the foregoing bet of S *is* also a bet on $\sim$S with negative payoff $(-b)$ if $\sim$S is true and a negative loss $(-a)$ if $\sim$S is false.[2]

$\sim$S	Payoff
T	$-b$
F	$-(-a)$

The betting quotient for $\sim$S is

$$\frac{-a}{(-a)+(-b)} = \frac{a}{a+b}$$

Now suppose S has a certain plausibility, Pl(S). If a bet on S is fair, the betting quotient for S must equal that plausibility:

$$\text{Pl}(S) = \frac{b}{a+b}$$

Since this bet is also a bet on $\sim$S and since it is fair, the betting quotient for $\sim$S must equal its plausibility:

[2] Remember that a bet on a statement, *P*, is an arrangement by which the bettor collects a certain sum, *a*, if *P* is true and forfeits a certain sum, *b*, if *P* is false. These quantities may be negative. Thus a bet of *P* is literally also a bet on $\sim P$. This argument thus depends on *fairness* being a property of the arrangement rather than on the way it is described (as bet on *P* or bet on $\sim P$, etc.).

$$\text{Pl}(\sim S) = \frac{a}{a+b}$$

Notice that plausibility is beginning to resemble probability since:

$$\text{Pl}(\sim S) = 1 - \text{Pl}(S)$$

Now let's look at the case where we have two mutually exclusive statements, *P*; *Q*. Suppose we find the proper betting quotients $\frac{b}{a+b}$ and $\frac{d}{c+d}$ to assure fair bets. Keeping to these quotients, we can choose the stakes so that they are equal on the bets $(a + b = c + d)$. In section VI.4 we saw that under such circumstances the sum of these bets is a bet on *PvQ* with betting quotient $\frac{b+d}{(a-d)+(b+d)}$. (If you don't understand where this came from, go back to section VI.4 and work it out.) This is just $\frac{b+d}{a+b} = \frac{b}{a+b} + \frac{d}{a+b}$. Since we assumed the stakes were equal, $a + b = c + d$. So the betting quotient for *PvQ* is equal to $\frac{b}{a+b} + \frac{d}{c+d}$; that is, to the sum of the betting quotients for *P* and for *Q*.

At this point we need to add one modest assumption about fairness to make any headway:

If a bet is a sum of fair bets it is fair.

Given this assumption it follows that:

If P and Q are mutually exclusive then Pl(PvQ) = *Pl*(P) + *Pl*(Q).

Taken together with the foregoing fact about negation [i.e., that Pl $(\sim S) = 1 - \text{Pl}(S)$] this shows us the Pl($Pv\sim P$) = 1 and Pl($P$&$\sim P$) = 0. Remember now from VI.4 that if two statements are logically equivalent, a bet on one is equally a bet on the other. Thus every tautology must have plausibility 1 and every contradiction plausibility 0. It would be hard to imagine anything less worthy of belief than P&$\sim P$ or more worthy than Pv$\sim P$. If we make this final assumption,

A tautology has the maximum plausibility and a contradiction the minimum

we have insured that all the clauses of Definition 14, section VI.2, have been met and thus that plausibility must, in fact, be probability.

To sum up: *If plausibility meets the following conditions, it is probability:*

(i) A bet on a statement, S, is fair if and only if the betting quotient for S equals the plausibility of S.
(ii) If two bets are fair, their sum bet is fair.
(iii) A tautology has the maximum plausibility value and a contradiction the minimum.

Exercises:

1. I bet on P with you, with the betting quotient for P being $\frac{a}{a+b}$ and the stakes being $a + b$ with a and b both positive quantities. Describe your bet on $\neg P$ with me.

2. More precisely, when we bet *with* someone he enters into an arrangement where our winnings are his losses, and *vice versa.* That is, the entries on his payoff table are the negative of the corresponding payoffs in our table. Call his bet the *complement* of our bet.

Consider the principle that *a bet is fair if and only if its complement is fair.* Using the results of Exercise 1, show that this principle can replace the use of negative winnings and losses in the argument for:

$$\text{Pl}(S) + \text{Pl}(\sim S) = 1$$

3. What kinds of bets give rise to betting quotients greater than 1 or less than 0? (Hint: if a and b are both positive, $\frac{a}{a+b}$ must be between 0 and 1. If a is positive and b is zero, $\frac{a}{a+b} = 1$. If b is positive and a is zero, $\frac{a}{a+b} = 0$. What if a and b are both negative?)

4. Give an intuitive argument as to why the kinds of bets that give rise to betting quotients less than 0 or greater than 1 could not reasonably be described as fair no matter how plausible or implausible the statement in question. Remember that a bet that is favorable or unfavorable is not fair.

5. Suppose we allowed some statements to have greater plausibility than a tautology, and some less than a contradiction, but kept to the other restrictions. So, for example: $\text{Pr}(P) = 2$ and $\text{Pr}(\sim P) = 1 - \text{Pr}(P) = -1$. Suppose also that we calculate the expected value of a bet on P in the normal way as $a\,\text{Pr}(P) - b\,\text{Pr}(\neg P)$.

a. Show that these plausibility values would give an expected value of 0 to some of the types of bets discussed in Exercises 3 and 4.
b. Show that these plausibility values would violate the following principle:

If Bet 1 and Bet 2 differ only in that Bet 2 has a greater payoff in one case than Bet 1, then Bet 2 has at least as great an expected value as Bet 1.

6. When we showed that if fair bets and expected values are to work reasonably, plausibility must be probability, our demonstration was based on minimal assumptions about fair bets. If we assume more about fair bets, the argument becomes very short. Assume that the expected value of a *betting arrangement* is the sum of the products of the plausibility of a case and the payoff in that case. We will consider only unfair bets.

a. Consider:

Bet 1		Bet 2	
P	Payoff	$P\vee\sim P$	Payoff
T	a	T	a
F	a	F	$-b$

Show that they are the same bet.

b. Assume the following *sure-thing* principle: *If a bet pays off a fixed amount,* a, *in every possible case, then the expected value of that bet must be* a. Under this assumption, show that the plausibility of tautology must be 1 and the plausibility of $\sim S$ must be $1 - \text{Pl}(S)$.

c. Suppose that P and Q are mutually exclusive. Consider the betting arrangement:

P	Q	Payoff
T	F	a
F	T	a
F	F	a

Using the rule for calculating expected value and the sure-thing principle, relate $\text{Pl}(P\&\neg Q)$; $\text{Pl}(\neg P\&Q)$; $\text{Pl}(\neg P\&\neg Q)$. Assuming that logically equivalent statements have the same plausibility, show that $\text{Pl}(P) + \text{Pl}(Q) = \text{Pl}(P\vee Q)$

d. Show that logically equivalent statements must have the same plausibility.

VI.6. THE DUTCH BOOK. If you are so foolish and your bookie is so clever that you conclude a series of bets with him such that he wins the sum bet *no matter what happens* he is said to have made a *Dutch Book* against you.

The following striking fact is often cited as a justification for the assumption that epistemic probabilities should obey the rules of the probability calculus:

> If you count as fair any bet on S if the betting quotient for S equals the plausibility of S, and if you are willing to make any series of bets each of which you regard as fair, then *if your plausibility values do not obey the rules of the probability calculus a Dutch Book can be made against you.*

Against the background of the previous two sections, the reasons for this theorem should be fairly transparent. Let us take the conditions for being a probability in order.

14a: No probability is less than zero.

If you have done the exercises you have already discovered the unpleasant results of taking plausibilities less than 0. Such plausibilities would lead me to regard a bet on P as fair, which would result in a loss whether P is true or false. For example, a plausibility of $-.10$ would lead me to regard the following bet as fair

P	Payoff
T	-110
F	-10

since the betting quotient is $\frac{b}{a+b} = \frac{10}{(-110)+10} \doteq -.10$. In general, any plausibility value, ϵ, for S will justify as fair a bet on S with winnings a if S is true and losses b if S is false just in case $\frac{a}{b} = \frac{1-\epsilon}{\epsilon}$. (*Exercise:* Show that this is true.) Thus, if ϵ is negative, it will justify a bet with negative winnings (a) and positive losses (b).

14b. If T is a tautology, $\Pr(T) = 1$.

Any plausibility greater than 1 will get us into what we just discussed because if ϵ is greater than 1, $\frac{1-\epsilon}{\epsilon}$ is negative. Suppose, on the other hand, we underestimate a tautology and give it plausibility less than 1. Then there will be some odds at which we consider it fair to bet against T (i.e., bet on T with negative winnings and negative losses). This is a bet we are sure to regret. For example, suppose we assign a plausibility of .75 to a tautology. This justifies a bet where $a = \$-25$ and $b = \$-75$. My bookie need only do a truth table to collect my $25.

14c: If P; Q are mutually exclusive, then $\Pr(P\text{v}Q) = \Pr(P) + \Pr(Q)$.

We have already learned that if P and Q are mutually exclusive, the sum of bets on P and on Q of equal stakes is a bet on $P\text{v}Q$ such that the betting quotient on $P\text{v}Q$ is the betting quotient on P plus the betting quotient on Q. Since I am committed to accepting any series of bets, each member of which I consider as fair, my bookie can always compel me to act *as if* $\text{Pl}(P\text{v}Q) = \text{Pl}(P) + \text{Pl}(Q)$ by placing separate bets at equal stakes on P and Q. He makes a bet on P which I consider fair. This means that the betting quotient for P equals what I take to be $\text{Pl}(P)$. Likewise with Q at equal stakes. The sum of these bets is a bet on $(P\text{v}Q)$ which I would consider fair if and only if I took $\text{Pl}(P\text{v}Q)$ to equal $\text{Pl}(P) + \text{Pl}(Q)$. But we are assuming that I take $\text{Pl}(P\text{v}Q)$ to have *another* value which establishes a different betting quotient for what I take to be a fair bet. I am offering my bookie two separate sets of odds on $P\text{v}Q$! Obviously, the thing for him to do is bet on $P\text{v}Q$ at one set of odds and against $P\text{v}Q$ at the other, choosing the most lucrative odds and catching me in the middle. For example, suppose the effective fair betting ratio on $P\text{v}Q$ resulting from separate bets on P and Q is .6 while the betting ratio I judge directly to be fair is .5. Then I will judge to be fair separate bets on P and Q whose sum will pay me $4 if $P\text{v}Q$ is true and cost me $6 if $P\text{v}Q$ is false. I will also judge a bet to be fair which costs me $5 if $P\text{v}Q$ is true and pays me $5 if $P\text{v}Q$ is false. If my bookie makes all these bets he will win $1 from me no matter what happens. The whole story is in the following table:

P	Q	Bet on P	Bet on Q	Sum of P; Q	Bet on $P\text{v}Q$	Sum of Bets on P; Q; $P\text{v}Q$
T	F	7	-3	4	-5	-1
F	T	-3	7	4	-5	-1
F	F	-3	-3	-6	$+5$	-1

If you understand the principles at work, you should be able to now show for yourself how this can be done in general.

The moral of the story is important. The Dutch Book being made against me results from my having two different effective betting quotients for $P\text{v}Q$. If we regard the odds that a person is willing to give on P a measure of his degree of belief on P, my problems stem from my having two incompatible degrees of belief. The most extreme case of this disease would occur if I gave a proposition degrees of belief 0 and 1, thus in effect believing with certainty both P and $\sim P$. Thus, if degrees of belief are held to be tied to betting quotients and betting behavior[3] in the manner indicated, the additivity requirement for probability is a kind of *consistency requirement* for degrees of belief. The fact that we do use our epistemic probabilities as weights for determining what risks to take in uncertain situations makes this the strongest argument to the effect that epistemic probabilities *are* probabilities.

The argument lacks one step of being complete. We have shown that if you violate the rules of the probability calculus you lay yourself open to a Dutch Book. But we have not shown that compliance with those rules protects you against a Dutch Book. Does it? Stop now, if you do not know, and think about the answer.

The answer is, of course, implicit in section VI.5. When someone makes a Dutch Book against you he entices you into a sum bet such that it is a sum of individual bets which you consider to be fair, but which itself guarantees you a loss in every case. Now if you are dealing in genuine probabilities (rather than some wacky plausibilities with values less than 0 or greater than 1) you will consider the sum bet to be unfair. The expected value of the sum bet is the sum of the products of the payoffs and the corresponding probabilities. Some of the probabilities will be positive; none will be negative. All the payoffs will be negative; so will the expected value.

Now we proved in section VI.5 that if we are using genuine probabilities to define fairness, *if two bets are fair, their sum bet is fair, and if two betting arrangements are fair their sum is fair.* It follows that if you are using genuine probabilities, no sequence of bets you consider fair can constitute a Dutch Book against you.

Exercises:

1. Suppose someone assigns a plausibility to $P \vee Q$ different from the sum of the plausibilities he assigns to P and to Q. Give explicit instructions for making a Dutch Book against him.

2. Show that if you are using genuine probabilities, no sequence of bets you consider fair or favorable can constitute a Dutch Book against you.

3. If you conclude a series of bets such that there is no possible circumstance under which you can win on the sum bet and there is some possible circumstance under which you can lose on the sum bet, we will say that you are the victim of a semi-Dutch Book. (This concept is due to Shimony, as are the following theorems.)

a. Prove that if you assign probability 1 to any statement other than a tautology, you lay yourself open to a semi-Dutch Book.

b. Show that if you adhere to the rules of the probability calculus and assign probability 1 only to tautologies, you are not open to a semi-Dutch Book.

[3] That a series of bets may be made if the bets individually are fair.

Suggested readings

Abner Shimony, "Scientific Inference," in *The Nature and Function of Scientific Theories*, ed. Robert Colodny (Pittsburgh: University of Pittsburgh Press, 1970), pp. 79–172.

Richard Jeffrey, *The Logic of Decision* (New York: McGraw-Hill Book Company, 1965).

VI.7. CONDITIONALIZATION. In the preceding sections we have discussed why *epistemic* probabilities should, in fact, be probabilities. The question of *inductive* probabilities has been left open. In this section, let us approach the question of inductive probabilities from rather a different angle than that of Chapter 1; that is, via the rule of conditionalization.

Let us assume, for the moment, that we are operating within a certainty model. We get to know more and more things with certainty and these items pile up, so to speak, in an ever-growing stock of knowledge. Suppose we have a certain "initial" set of epistemic probabilities; $\Pr_i$; and our senses toss a new item of knowledge, P, into our stock of knowledge. How are we to move to a "final" set of epistemic probabilities, $\Pr_f$, which accommodate our new item of knowledge in a rational fashion? The rule of conditionalization gives this answer:

> **Rule C:** For any statement Q, take its new probability to be its old probability conditional on the new item of knowledge, i.e., $\Pr_f(Q) = \Pr_i(Q \text{ given } P)$.

Note that Rule C gives P the new value of 1, a status commensurate with its new-found certainty. Rule C, however, can also effect a change in the epistemic probability of nearly every other statement. What justification is there for making these changes according to this rule? For the answer we must go back to bets once more.

Remember that the sum of bets on $P\&Q$ and on $\sim P$ is, if the stakes are right, a bet on Q conditional on P with a payoff table as represented below:

P	Q	Payoff
T	T	a
T	F	$-b$
F	T	0
F	F	0

This betting arrangement is fair just in case its expected value is zero, that is:

$$a\Pr(P\&Q) - b\Pr(P\&{\sim}Q) = 0$$

or:

$$\frac{\Pr(P\&Q)}{\Pr(P\&{\sim}Q)} = \frac{b}{a}$$

Note that several sets of probability values will render such a bet fair. For example, if $a = \$1$ and $b = \$2$, then $\Pr(P\&Q) = \frac{2}{3}$ and $\Pr(P\&{\sim}Q) = \frac{1}{3}$ renders the bet fair, as does $\Pr(P\&Q) = \frac{2}{6}$ and $\Pr(P\&{\sim}Q) = \frac{1}{6}$. The first set of values makes $\Pr(P) = 1$ and the second set makes $\Pr(P) = \frac{1}{2}$. In fact, *any* value of P is compatible with our conditional bet so long as it is[4] divided up into $\Pr(P\&Q)$ and $\Pr(P\&{\sim}Q)$ in the same ratio as b to a.

In other words, the bet is fair just in case:[5]

$$\Pr(P\&Q) = \frac{b}{a+b}\Pr(P)$$

$$\Pr(P\&{\sim}Q) = \frac{a}{a+b}\Pr(P)$$

or

$$\frac{\Pr(P\&Q)}{\Pr(P)} = \Pr(Q \text{ given } P) = \frac{b}{a+b}$$

$$\frac{\Pr(P\&{\sim}Q)}{\Pr(P)} = \Pr({\sim}Q \text{ given } P) = \frac{a}{a+b}$$

[4] If not zero.

[5] Or $\Pr(P) = 0$ or $a = b = 0$.

If we call $\frac{b}{a+b}$ the *betting quotient on* Q *conditional on* P, which seems reasonable, we can now say that a conditional bet is fair when the conditional betting quotients equal the corresponding conditional probabilities.

Now the interesting thing to notice about all this is the connection between a conditional bet's *remaining* fair under a belief change and that change taking place by conditionalization. Suppose a bet on Q conditional on P with conditional betting quotient $\frac{b}{a+b}$ is fair on a set of initial probabilities $\Pr_i$. Then $\frac{\Pr_i(P\&Q)}{\Pr_i(P)} = \Pr_i(Q \text{ given } P) = \frac{b}{a+b}$. Suppose now that a change to a new set of probabilities is made by conditionalizing on P. Then $\Pr_f(P\&Q) = \Pr_i(P\&Q \text{ given } P) = \Pr_i(Q \text{ given } P)$ and $\Pr_f(P) = \Pr_i(P \text{ given } P) = 1$. Thus, $\frac{\Pr_f(P\&Q)}{\Pr_f(P)} = \frac{\Pr_i(Q \text{ given } P)}{1} = \frac{b}{a+b}$. So, *if beliefs are changed by conditionalization on* P, *fair bets conditional on* P *remain fair.*

Very nice. But what's nicer is that conditionization is the *only* method for changing beliefs under these circumstances[6] which has this property. Suppose a bet on Q conditional on P is fair before and after a belief change from $\Pr_i$ to $\Pr_f$. Then $\frac{\Pr_i(P\&Q)}{\Pr_i(P)} = \frac{b}{a+b} = \frac{\Pr_f(P\&Q)}{\Pr_f(P)}$. If this belief is a result of P becoming *certain*, then $\Pr(P) = 1$. Furthermore, $\Pr_f(P\&Q)$ must equal $\Pr_f(Q)$ for $\Pr_f(Q) = \Pr_f(P\&Q) = \Pr_f({\sim}P\&Q)$ and $\Pr_f({\sim}P\&Q)$ must be 0, since $\Pr_f(P) = 1$. So $\Pr_f(Q) = \frac{\Pr_i(P\&Q)}{\Pr_i(P)}$ and the change has taken place by conditionalization. *The only method of changing beliefs such that* P *becomes certain and bets conditional on* P *which are fair remain fair is conditionalization on* P[7]. If this is true, a Dutch Book argument cannot be far away. We have shown that if someone does *not* change his beliefs according to Rule C, the conditional betting quotient which he regards as assuring a fair bet on Q conditional on P will *change* upon the acquisition of P as an

[6] Certainty model.

[7] Excepting cases where $\Pr(P) = 0$ or $a = b = 0$.

item of knowledge. If the bookie knows *how* the bettor will change his betting quotients he is clearly in a position to guarantee a profit if P occurs. By making conditional bets before and after that occurrence, he is essentially betting on Q at two different sets of odds. We have already seen how a bookie can assure himself a profit in such a situation. If $Pr_i(Q \text{ given } P)$ is less than $Pr_f(Q \text{ given } P)$, he will bet initially *on* Q conditional on P and finally *against* Q conditional on P. If $Pr_i(Q \text{ given } P)$ is greater than $Pr_f(Q \text{ given } P)$, he will bet initially *against* Q conditional on P and finally *on* Q conditional on P. Choosing the stakes correctly, he guarantees himself a profit if P occurs. Furthermore, he breaks even if P does not occur, since all bets are conditional on P. Only one more small step is required to achieve a proper Dutch Book. The bookie considers the amount, a, that he has guaranteed he will win if P occurs, and makes a side bet of $\frac{1}{2}a$ on $\sim P$, guaranteeing himself net winnings, come what may.

The virtues of conditionalization having been firmly established, now let us look a little more closely at the workings of the certainty model with the rule of conditionalization. As we travel through life, with our eyes open, we come to know more and more things. This growth of knowledge is represented by the adding of statements $(O_1, O_2, \ldots O_n)$ to our stock of knowledge. Upon the addition of a new item of knowledge, O_n, to our stock of knowledge, we revise our belief structure by passing from old epistemic probabilities, Pr_n to new epistemic probabilities $Pr_{(n+1)}$ by conditionalization *on* O_n.[8] So for any statement Q, $Pr_{(n+1)}(Q) = Pr_n(Q \text{ given } O_n)$, and

$$Pr_{(n+2)}(Q) = Pr_{n+1}(Q \text{ given } O_{n+1})$$

but

$$Pr_{n+1}(Q \text{ given } O_{n+1}) =$$

$$\frac{Pr_{n+1}(Q\&O_{n+1})}{Pr_{n+1}(O_{n+1})} =$$

$$\frac{Pr_n(Q\&O_{n+1} \text{ given } O_n)}{Pr_n(O_{n+1} \text{ given } O_n)} =$$

$$\frac{Pr_n(Q\&O_{n+1}\&O_n)/Pr_n(O_n)}{Pr_n(O_{n+1}\&O_n)/Pr_n(O_n)} =$$

[8] We are assuming that $Pr_n(O_n) > 0$, so that the conditional probabilities are well defined.

$$\frac{Pr_n(Q\&O_{n+1}\&O_n)}{Pr_n(O_{n+1}\&O_n)} =$$

$$Pr_n(Q \text{ given } O_n\&O_{n+1}).$$

So two steps can be compressed into one. First conditionalizing on O_n and moving from the resulting distribution by conditionalizing on O_{n+1} is equivalent to moving from the original distribution by conditionalizing on the conjunction $O_n\&O_{n+1}$. It follows that we can compress any finite number of steps into one.

The set of epistemic probabilities, $Pr_{(n+1)}(Q)$ *arrived at by successive conditionalizations on items in a stock of knowledge* $(O_1, O_2 \ldots O_n)$ *is identical to the set of probabilities which would be arrived at by conditionalization on the conjunction of all those items of knowledge,* $Pr_1(Q \text{ given } O_1\&O_2\& \ldots \&O_n)$.

$Pr_1(Q \text{ given } O_1\&O_2\& \ldots \&O_n)$ is a measure of the firmness with which $O_1\&O_2\&\ldots\&O_n$ supports Q. Since Pr_1 is not the result of a conditionalization, it does not depend upon the contents of our stock of knowledge. This suggests that we might identify it as the *inductive* probability of the argument:

$$\begin{array}{c} O_1 \\ O_2 \\ \cdot \\ \cdot \\ \cdot \\ O_n \\ \hline Q \end{array}$$

This identification is vouchsafed by Definition 5 of Chapter I:

> In the certainty model the epistemic probability of a statement is the inductive probability of that argument which has that statement as its conclusion and whose premises consist of all the observation reports which comprise our stock of knowledge.

and it answers the question with which we began this section. *"Inductive probabilities" must, in fact, be conditional probabilities.*

In Chapter I, we started with inductive probabilities and, in the certainty model, defined epistemic probabilities in terms of them. In this section we started with epistemic probabilities and, within the assumptions of the certainty model, recovered inductive probabilities.

The approach of Chapter I is that of Carnap; that of this section is associated with the Bayesian school. That they coincide to such an extent is a pleasant and informative fact.

Exercise:

Show that if I move from an initial set of probabilities, first by conditionalizing on P, then on Q, then on R, to a final set of probabilities, then for any statements S and T:

$$\frac{\Pr_f(P\&Q\&R\&S)}{\Pr_f(P\&Q\&R\&T)} = \frac{\Pr_i(P\&Q\&R\&S)}{\Pr_i(R\&Q\&R\&T)}$$

(provided the initial probabilities are positive).

VI.8. FALLIBILITY. *A man would be rash indeed if his acts of observation all resulted in *certainty* in an associated observation statement. In fact, there are reasons to believe that it is never rational to be certain (in the sense of assigning epistemic probability of 1) of any observational statement. The first reason is Shimony's argument that assigning probability 1 or 0 to any statement not a logical truth or a contradiction respectively lays us open to a quasi–Dutch Book. The other reasons have emerged from much threshing about by epistemologists in this century. The threshing is perhaps not yet over, and no brief summary of its results is likely to be regarded as fair by all sides. Nevertheless, what I take to be the heart of the matter is this: no matter what language we use to describe our observations, the act of observation and the act of believing a sentence attributing a certain character to that observation are distinct.[9] Doing one does not entail doing the other. The link between them is causal, not logical. If I am of sound mind and body, adopt a modest observation language, and am proficient in its use, this causal process may be highly reliable as a means for generating true beliefs. But there is no reason whatsoever to believe that it is infallible.

In such circumstances it is hard to see how it would be reasonable to be certain. Remember that certainty for us means an epistemic probability equal to 1. And if $\Pr(P) = 1$, the bet

*This section deals with an advanced topic and may be omitted without loss of continuity.

[9] N.B. "Distinct" means "not identical." It does not mean "disjoint"; it does not mean "unrelated."

P	Payoff
T	0
F	$-b$

is fair no matter how great b is. It is common folk knowledge that someone who says he is certain and who even *feels* certain, may shrink from putting his money where his mouth is. Certainty of the sort in which we are interested involves the willingness to risk *everything* if you are wrong over against no gain if you are right.

If all this has not convinced you that certainty is never warranted for contingent statements, I hope it has at least convinced you that there are *some* times when we wish to change our beliefs under the pressure of new evidence where the certainty model is inappropriate. We need, then, a way of changing our epistemic probabilities when an observation raises our degree of belief in a statement, without raising it all the way to 1.

Suppose that an observation causes us to change our degree of belief in P from $\Pr_i(P)$ to $\Pr_f(P)$. We might hope that our rule for changing beliefs in such a situation would be such that bets conditional on P and bets conditional on $\sim P$ which are fair before the change remain fair. We saw in section VI.7 that bets conditional on P remain fair just in case the ratio of $\Pr(P\&Q)$ to $\Pr(P\&\sim Q)$ remains constant. And this ratio remains constant just in case the conditional probabilities $\Pr(Q \text{ given } P)$ and $\Pr(\sim Q \text{ given } P)$ remain constant. By the same token, fair bets conditional on $\sim P$ remain fair just in case the conditional probabilities $\Pr(Q \text{ given } \sim P)$ and $\Pr(\sim Q \text{ given } \sim P)$ remain constant. Thus if fair bets conditional on P and on $\sim P$ are to remain fair:

$$\Pr_f(P\&Q) = \Pr_f(P)\Pr_i(Q \text{ given } P)$$

$$\Pr_f(P\&\sim Q) = \Pr_f(P)\Pr_i(\sim Q \text{ given } P)$$

$$\Pr_f(\sim P\&Q) = \Pr_f(\sim P)\Pr_i(Q \text{ given } \sim P)$$

$$\Pr_f(\sim P\&\sim Q) = \Pr_f(\sim P)\Pr_i(\sim Q \textit{ given } \sim P)$$

and

$$\Pr_f(Q) = \Pr_f(P\&Q) + \Pr_f(\sim P\&Q)$$

Putting these together we have:

Jeffrey's Rule: If our new information is represented as a change in degree of belief in P from $\Pr_i(P)$ to $\Pr_f(P)$, then for any statement Q, take:

$$\Pr_f(Q) = \Pr_f(P)\Pr_i(Q \text{ given } P) + \Pr_f({\sim}P)\Pr_i(Q \text{ given } {\sim}P)$$

Notice that Jeffrey's rule is a generalization of Rule C. In the special case where $\Pr_f(P) = 1$, Jeffrey's rule reduces to Rule C. Notice also that Jeffrey's rule can be viewed as a weighted average of Rule C to both P and to ${\sim}P$. Conditionalizing on P, $\Pr_f(Q)$ would be $\Pr_i(Q$ given $P)$. Conditionalizing on ${\sim}P$, $\Pr_f(Q)$ would be $\Pr_i(Q$ given ${\sim}P)$. Averaging these results, weighting the average by $\Pr_f(P)$ and $\Pr_f({\sim}P)$, gives us Jeffrey's rule.

We have, then, a viable fallibility model for changing from one set of epistemic probabilities to another. But now it is not so easy as it was in the certainty model to represent an epistemic probability as the result of an inductive probability operating on a stock of knowledge. What observation gives us now is not a set of certain sentences $0_1, 0_2, \ldots$, but rather a set of observational probabilities, $\Pr_o(0)_1)$; $\Pr_o(0_2)$; .... The observational probabilities are to be the outcome solely of the observation, not of inductive reasoning, for the point is to separate out the factors of observation and induction.

In the certainty model we showed that conditionalizing first on 0_1, then on 0_2, etc., gave the same result as conditionalizing on their conjunction 0_1 & 0_2. Hence the possibility of "factoring" our epistemic probability into a stock of knowledge and a set of inductive probabilities. In general there is no long conjunction and associated probability to which we can apply Jeffrey's rule and get the same set of epistemic probabilities as we would have gotten from successive applications of that rule.

Suppose we attempt to define our epistemic probability as the result of applying Jeffrey's rule successively to each item in our stock of knowledge, taking inductive probabilities as the conditional probabilities in the first step, the resulting epistemic conditional probabilities as conditional probabilities for the next step, and so on. This will not do, for several reasons. The first is that the final result differs depending on the *order* in which the items in our stock of knowledge are taken in this process. This will not work, since the *same data* coupled with the same inductive probabilities should generate the same epistemic probability. The second reason is that at each stage in this process the

observational probability is taken as the final probability. In the certainty model if P is observed, it becomes certain. Well and good. Its final probability becomes 1. But if we are sophisticated enough to realize that observations may fall short of certainty, we should be sophisticated enough to realize that observational probability need not be the only factor influencing final probability. Final probability is rather the result of the interaction of observational probability with theories which we may hold on the basis of previous observations.

Let me illustrate. Suppose I see a bird at twilight which I clearly identify as a raven. Because the light is not so good, the probability I can assign to him being black on the basis of that observation is only .8. Suppose further that I hold the theory that all ravens are black and that this theory is buttressed by massive numbers of previous observations. In such a situation the final probability I assign to the statement that the raven is black will be higher than the observational probability, and quite properly so. Otherwise I could disconfirm lots of theories just by running around at night.

All right, my *theory* (which is really the conduit of the force of previous observations) *pulls up* the observational probability in this case. It is just as easy to think of cases where *it pulls it down*, say where I think I see a water buffalo on the San Bernardino Freeway at 3 A.M.

Can we have an analysis of the interaction of theory and observation along these lines? Is there a valid Dutch Book argument for Jeffrey's rule? These are controversial questions under current investigation. If you find them of compelling interest, you may want to follow up the suggested readings.

Exercises:

1. Start with initial probabilities $\Pr(P\&Q) = 1/3$; $\Pr(P\&{\sim}Q) = .001$; $\Pr({\sim}P\&Q) = 1/3$; $\Pr({\sim}P\&{\sim}Q) = 1/3$. Apply Jeffrey's rule taking $\Pr_f(P) = .99$. Calculate $\Pr_f(P\&Q)$; $\Pr_f(P\&{+}Q)$; $\Pr_f({\sim}P\&Q)$; $\Pr_f({\sim}P\&Q)$. Now taking this set of probabilities as initial probabilities, apply Jeffrey's rule taking $\Pr_f(Q) = .99$. Calculate the final probabilities of all the same statements.

Now repeat the process in opposite order; that is, first apply Jeffrey's rule on Q at $\Pr_f(Q) = .99$, then on P at $\Pr(P) = .99$. Compare this set of final probabilities with the previous one.

2. Suppose we move from $\Pr_i$ to $\Pr_f$ by applying Jeffrey's rule to P, taking $\Pr_f(P)$ to have some value between 0 and 1. Suppose also that $\Pr_i(S) = 1$ only if S is a tautology and $\Pr_i(S) = 0$ only if S is a contradiction.

   a. Show that $\Pr_f(S) = 1$ if S is a tautology and $\Pr_f(S) = 0$ only if S is a contradiction.

b. Show that for any contingent statements S and T,

$$\frac{\Pr_i(P\&S)}{\Pr_i(P\&T)} = \frac{\Pr_f(P\&S)}{\Pr_f(P\&T)}$$

and

$$\frac{\Pr_i(\sim P\&S)}{\Pr_i(\sim P\&T)} = \frac{\Pr_f(\sim P\&S)}{\Pr_f(\sim P\&T)}$$

c. If $\Pr_i(P) = a$ and $\Pr_f(P) = b$, show that first applying Jeffrey's rule to $\Pr_i$ with $\Pr_f(P) = b$ and then applying it to that set of probabilities with $\Pr_f(P) = a$ gets you back to the initial set of probabilities.

Suggested readings

Brad Armendt, "Is There a Dutch Book Argument for Probability Kinematics?" *Philosophy of Science* 47 (1980): 583–588.

Rudolf Carnap, "Inductive Logic and Rational Decisions," in *Studies in Inductive Logic and Probability* I, ed. Carnap and Jeffrey (Los Angeles: University of California Press, 1971), pp. 5–31.

Hartry Field, "A Note on Jeffrey Conditionalization," *Philosophy of Science* 45 (1978): pp. 361–367.

Richard Jeffrey, *The Logic of Decision* (2nd ed.) (Chicago and London: University of Chicago Press, 1983), chap. 11.

Brian Skyrms, "Higher Order Degrees of Belief," in *Prospects for Pragmatism: Essays in Honor of F. P. Ramsey*, ed. D. H. Mellor (Cambridge: Cambridge University Press, 1980).

Paul Teller, "Conditionalization and Observation," *Synthese* 26 (1973): 218–258.

VI.9. UTILITY. We have been operating so far within a set of assumptions that often approximate the truth for monetary gambles at small stakes. It is time to take a more global viewpoint and question these assumptions.

An extra hundred dollars means less to a millionaire than to an ordinary person. But if I win a million, I'm a millionaire. So the difference between winning a million + 100 dollars and winning a million means less to me than the difference between winning 100 dollars and winning nothing. In the terminology of economics, money has decreasing rather than constant marginal utility for me.

The idea of utility was introduced into the literature on gambling in this connection by Daniel Bernoulli in 1738. Bernoulli was concerned with the St. Petersburg game. In this game, you flip a fair coin until it comes up heads. If it comes up heads on the first toss, you get \$2; if on the second toss, \$4; if on the

third toss, \$8; if on the nth toss, $\$2^n$. The expected dollar value of this game is infinite. (*Exercise:* check this!) How much would you pay to get into this game? Bernoulli's idea was that if the marginal utility of money decreased in the right way,[10] the St. Petersburg game could have a reasonable finite expected *utility* even though the monetary expectation is infinite.

When we consider decisions whose payoffs are in real goods rather than money, there is another complication we must take into account. That is, the value of having two goods together may not be equal to the sum of their individual values because of interactions between the goods. If a man wants to start a pig farm, and getting a sow has value b for him and getting a boar has value c, then getting both a sow and a boar may have value greater than $b + c$. The sow and the boar are, for him, *complementary* goods. Interaction between goods can also be negative, as in the case of the prospective chicken farmer who wins two roosters in two lotteries. The presence of an active market reduces, but does not eliminate the effect of complementarities. The second rooster is still of more value to a prospective chicken farmer in Kansas than to Robinson Crusoe. The farmer can, for example, swap it for a hen; or at least sell it and put the money toward a hen. Because of complementarities, we cannot in general assume that if a bettor makes a series of bets each of which he considers to be fair, he will judge the result of making them all together as fair. Where payoffs interact, the right hand may need to know what the left is doing.

The preceding points about how utility works are intuitively easy to grasp. But it is harder to say just what utility is. We know how to count money, pigs, and chickens; but how do we measure utility? Von Neumann and Morgenstern showed how to use the expected utility principle to measure utility if we have some chance device (such as a wheel of fortune, a fair coin, a lottery) for which we know the chances. We pick the best payoff in our decision problem and give it (by convention) utility 1; likewise, we give the worst payoff utility 0. Then we measure the utility of a payoff, P, in between by judging what sort of a gamble with the worst and the best payoffs as possible outcomes has value equal to P. For instance, farmer Jones wants a horse, a pig, a chicken, and a husband. Her current decision situation is structured so that she will get exactly one of these. She ranks the payoffs:

1. Horse
2. Husband
3. Pig
4. Chicken

---

[10] Utility = log Money.

She is indifferent between (1) a lottery that gives $\frac{4}{5}$ chance of a horse and $\frac{1}{5}$ chance of a chicken, and one that gives a husband for sure; and (2) a lottery that gives $\frac{1}{2}$ chance of a horse and $\frac{1}{2}$ chance of a chicken, and one that gives a pig for sure. Thus, her utility scale looks like this:

	Utility
Horse	1
Husband	.8
Pig	.5
Chicken	0

If her decision situation were structured so that she might end up with all these goods, and if they didn't interfere with one another, then her utility scale might have a different top: Horse and Husband and Pig and Chicken. If it were structured so that she might end up getting none of these goods, after going to some expense, there might be a different bottom, which would have utility 0.

Utility, as measured by the von Neumann–Morgenstern method, is *subjective* utility, determined by the decision maker's own preferences. There are, no doubt, various senses in which a decision maker can be wrong about what is good for him. However, such questions are not addressed by this theory.

From a decision maker's utilities we can infer his degrees of belief. Farmer Smith has bought two tickets to win for a race at the county fair, one on Stewball and one on Molly. If he holds a ticket on a winning horse, he wins a pig; otherwise he gets nothing. We assume that he does not care about the outcome of the horse race per se; it is important to him only insofar as it does or does not win him a pig. He is indifferent to keeping his ticket on Stewball or exchanging it for an objective lottery ticket with a known 10 per cent chance of winning; likewise for Molly or an objective lottery ticket with a 15 per cent chance of winning.

Farmer Smith's utility scale looks like this:

	Utility
Pig	1
Ticket on Molly	.15
Ticket on Stewball	.10
Nothing	0

If he maximizes expected utility, his expected utility for his bet (ticket) on Molly is:

Degree of Belief (Molly Wins) Utility (Pig) +
Degree of Belief (Molly Loses) Utility (Nothing)

This is just equal to his degree of belief that Molly wins. Then his *degree of belief* that Molly wins is .15; in the same way, his degree of belief that Stewball wins is .10. Subjective degrees of belief are here recovered from subjective utilities in an obvious and simple way. (Things would be more complicated if farmer Smith cared about the outcome of the race over and above the question of the pig, but as we shall see in the next section, his subjective degrees of belief could still be found.)

Exercises:

1. A decision maker with declining marginal utility of money is *risk averse* in monetary terms. He will prefer $50 for sure to a wager that gives a chance of $\frac{1}{2}$ of winning $100 and a chance of $\frac{1}{2}$ of winning nothing, because the initial $50 has more utility for him than the second $50. Suppose that winning $100 is the best thing that can happen to him and winning nothing is the worst.
   a. What is his utility for winning $100?
   b. What is his utility for winning nothing?
   c. What is his utility for a wager that gives a known objective chance of $\frac{1}{2}$ of winning $100 and $\frac{1}{2}$ of winning nothing?
   d. What can we say about his utility for getting $50?
   e. Draw a graph of utility as against money for a decision maker who is generally risk averse.

2. Suppose farmer Smith has one ticket on each horse running at the county fair, and thus will win a pig no matter which horse wins. Let $U(\text{pig}) = 1$ and $U(\text{nothing}) = 0$. Suppose farmer Smith's preferences go by expected utility.
   a. Farmer Smith believes that all his tickets taken together are worth one pig for sure. What does this tell you about his degrees of belief about the race?
   b. Suppose that farmer Smith also believes that each of his tickets has equal utility. What does this tell you about his degrees of belief about the race?

3. (Advanced) Let us say that the *physical sum* of two bets, B_1; B_2, pays off, at each case, both the physical goods that B_1 pays off and the physical goods that B_2 pays off. However, let us say that a bet B_3 is the *mathematical sum* of B_1 and B_2 if it pays off, at each case, a physical good whose utility is equal to the *sum of the utilities* of the physical payoffs of B_1 and B_2 in that case. Show the following:
   a. For someone who cares only about gold, and whose marginal utility for gold is constant, the physical sum of two bets (with payoffs in gold) is the mathematical sum.
   b. For someone who cares only about gold but whose marginal utility for gold is declining, the physical sum of two bets need not equal their mathematical sum. (Hint: see Exercise 1.)
   c. For payoffs in arbitrary physical goods, the physical sum of two bets may fail to equal their mathematical sum as a result of complementarities.

d. Suppose that propositions p and q are incompatible and that one has a betting arrangement that has a payoff of utility x if p is true (and q false), and again a payoff of utility x if q is true (and p false), and a utility 0 if p and q are both false ($x > 0$). This betting arrangement can be correctly described in two ways: (i) as a bet with payoff of utility x if p *or* q is true, utility 0 otherwise; (ii) as the *mathematical* sum of two bets, one of which yields a payoff of utility x if p is true, utility 0 otherwise; the other of which yields a payoff of utility x if q is true, utility 0 otherwise. Reconsider the Dutch Book arguments in light of the foregoing.

Suggested readings

R. D. Luce and H. Raiffa, *Games and Decisions* (New York: Wiley, 1957), chap. 2.

J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior* (2nd ed.) (Princeton: Princeton University Press, 1947).

VI.10. RAMSEY. The von Neumann–Morgenstern theory of utility is really a rediscovery of ideas contained in a remarkable essay, "Truth and Probability," written by F. P. Ramsey in 1926. In the essay, Ramsey goes even deeper into the foundations of utility and probability. The von Neumann–Morgenstern method requires that the decision maker know some objective chances, which are then used to scale his subjective utilities. From his subjective utilities and preferences, information about his subjective probabilities can be recovered. Ramsey starts without the assumption of knowledge of some chances, and with only the decision maker's preferences.

Ramsey starts by identifying propositions that, like the coin flips, lotteries, and horse races of the previous section, have no value to the decision maker in and of themselves, but only insofar as certain payoffs hang on them. He calls such propositions "ethically neutral." A proposition, p, is ethically neutral for a collection of payoffs B if it makes no difference to the agent's preferences, that is, if he is indifferent between B with p true and B with p false. A proposition, p, is ethically neutral if p is ethically neutral for maximal collections of payoffs relevant to the decision problem. The nice thing about ethically neutral propositions is that the expected utility of gambles on them depends only on their probability and the utility of their outcomes. Their own utility is not a complicating factor.

Now we can identify an ethically neutral proposition, H, as having probability $\frac{1}{2}$ for the decision maker if there are two payoffs, $A;B$, such that he prefers A to B but is indifferent between the two gambles: (1) Get A if H is true, B if H is false; (2) get B if H is true, A if H is false. (If he thought H was more likely than $\sim H$, he would prefer gamble 1; if he thought $\sim H$ was more likely than H, he would prefer gamble 2. For the purpose of scaling the decision maker's utilities, such a proposition is just as good as the proposition that a fair coin comes up heads.

The same procedure works in general to identify surrogates for fair lotteries. Suppose there are 100 ethically neutral propositions, $H_1;H_2;\ldots;H_{100}$, which are pairwise incompatible and jointly exhaustive. Suppose there are 100 payoffs, $G_1;G_2;\ldots;G_{100}$, such that G_1 is preferred to G_2, G_2 is preferred to G_3, and so forth up to G_{100}. Suppose the decision maker is indifferent between the complex gamble:

If H_1 get G_1 &
If H_2 get G_2 &
.
If H_i get G_i &
.
If H_{100} get G_{100}

and every other complex gamble you can get from it by moving the Gs around. Then each of the H_is gets probability .001, and together they are just as good as a fair lottery with 100 tickets for scaling the decision maker's utilities.

A rich enough preference ordering has enough ethically neutral propositions forming equiprobable partitions of the kind just discussed to carry out the von Neumann–Morgenstern type of scaling of utilities described in the last section to any desired degree of precision. Once the utilities have been determined, the degree of belief probabilities of the remaining ethically neutral propositions can be determined in the simple way we have seen before. The decision maker's degree of belief in the ethically neutral proposition, p, is just the utility he attaches to the gamble: *Get* G *if* p, B *otherwise,* where G has utility 1 and B has utility 0.

With utilities in hand, we can also solve for the decision maker's degrees of belief in non-ethically neutral propositions, although things are not quite so simple here. Suppose that farmer Smith owned Stewball and wanted his horse to win, as well as wanting to win a pig. Then "Stewball wins" and "Molly wins" are not ethically neutral for him. Now suppose we want to determine his degree of belief in the proposition that Molly wins. Given our conventions, we can't set up a gamble that gives utility 1 if Molly wins because what farmer Smith desires most and gives utility 1 is: "Get a pig and Stewball wins." But we know that the expected utility of the wager "Pig if Molly wins, no prize if she loses" is equal to:

Pr(Molly wins) U(get pig and Molly wins) +
1 − Pr(Molly wins) U(no prize and Molly loses)

If we know the utility of the wager, of "Get pig and Molly wins," and of "No prize and Molly loses," we can solve for Pr(Molly wins).

For a rich and coherent preference ordering over gambles, Ramsey has conjured up both a subjective utility assignment and a degree of belief probability assignment such that preference goes by expected utility. This sort of representation theorem shows how deeply the probability concept is rooted in practical reasoning.

Exercises:

1. Suppose that the four propositions, *HH;HT;TH;TT*, are pairwise incompatible (at most one of them can be true) and jointly exhaustive (at least one must be true). Describe the preferences you would need to find to conclude that they are ethically neutral and equiprobable.

2. Suppose that farmer Smith owns Stewball and that "Molly wins" is not ethically neutral. His most preferred outcome is "Get pig and Stewball wins"; his least preferred is "No pig and Molly loses"; therefore, these get utility 1 and 0, respectively. The propositions, *HH;HT;TH;TT* are as in Exercise 1. Farmer Smith is indifferent between "Get pig and Molly wins" and a hypothetical gamble that would ensure that he would get the pig and Stewball would win if *HH* or *HT* or *TH* and that he would get no pig and Stewball would lose if *TT*. (What does this tell you about his utility for "Get pig and Molly wins"?) He is indifferent between "No pig and Molly loses" and the hypothetical gamble that would ensure that he would get the pig and Stewball would win if *HH* and that he would get no pig and Stewball would lose if *HT* or *TH* or *HH*. He is indifferent between the gamble "Pig if Molly wins; no pig if she loses" and the gamble "Get pig and Stewball wins if *HH* or *HT*, but no pig and Stewball loses if *TH* or *TT*."
   a. What are his utilities for "Get pig and Molly wins"; "No pig and Molly loses"; the gamble "Pig if Molly wins; no pig if Molly loses"?
   b. What is his degree of belief probability that Molly will win?

Suggested readings

F. P. Ramsey, "Truth and Probability," in *The Foundations of Mathematics and Other Logical Essays*, ed. R. B. Braithwaite (London: Routledge and Kegan Paul, 1931) and in *Studies in Subjective Probability*, ed. H. Kyburg and H. Smokler (Huntington, N.Y.: Krieger, 1980).

For the advanced student:

Peter Fishburn, "Subjective Expected Utility, a Review of Normative Theories" *Theory and Decision* 13 (1981), pp. 139–199.

Terrence Fine, *Theories of Probability* (New York and London: Academic Press, 1973), chap. VIII.

VII

Kinds of Probability

VII.1. INTRODUCTION. Historically, a number of distinct but related concepts have been associated with the word *probability*. These fall into three families: rational degree of belief, relative frequency, and chance. Each of the probability concepts can be thought of as conforming to the mathematical rules of probability calculus, but each carries a different meaning. We have, in one way or another, met each of these probability concepts already in this book. A biased coin has a certain objective *chance* of coming up heads. If we are uncertain as to how the coin is biased and what the objective chance really is, we may have a rational *degree of belief* that the coin will come up heads that is unequal to the true chance. If we flip the coin a number of times, a certain percentage of the tosses will come up heads; that is, the *relative frequency* of heads in the class of tosses will be a number in the interval from 0 to 1. The relative frequency of heads may well differ from both our degree of belief that the coin will come up heads and the objective chance that the coin comes up heads. The concepts are distinct, but they are closely related. Observed *relative frequencies* are important evidence that influences our *rational degrees of belief* about *objective chances*. If, initially, we are unsure whether the coin is biased 2 to 1 in favor of heads or 2 to 1 in favor of tails (degree of belief $\frac{1}{2}$), and then we flip the coin 1000 times and get 670 heads, we will have gotten strong evidence indeed that the coin is biased toward heads. Along just these lines, Cicero evaluated divination as a statistical theory and found it unworthy of a high degree of belief. In our own time, microphysics consists of theories that postulate chances, and that are largely tested against frequentist evidence. This final chapter is devoted to a review of these conceptions of probability and a sketch of their interrelation.

VII.2. RATIONAL DEGREE OF BELIEF. Belief is not really an all or nothing affair; it admits of degrees. You might be reasonably sure that the president was guilty without being absolutely certain. You might be extremely dubious about the plaintiff's supposed whiplash injury without being certain that he is malingering. You might think of it as only slightly more likely than not that the cause of a sore throat is a virus. Degrees of belief can be represented numerically, with larger numbers corresponding to stronger beliefs. What should the mathematics of these numbers be for a rational agent?

Chapter VI introduced some of the reasons why it has been held that *rational* degrees of belief should admit a numerical representation that obeys the