ads for Tylenol. How could a healthy man have arthritis in his thirties? Dr. Clarfield, a hockey physician, said, "He played enough hockey for three careers, so *the trauma in his joints was probably more than it was for most players."*

6 *Landing on Mars.* A NASA spokeswoman said that *the first person to set foot on Mars would in all probability be an engineer.*

7 *Lunatic coin tossers.* A fair coin is being tossed. We say that after N tosses, heads is in the lead if heads has come up more times than tails. Conversely, tails is in the lead if tails has come up more often. (With an even number N, they could be tied.) We say that the lead changes at toss N if heads was in the lead at N−1, but tails is in the lead at N+1. *The probability that in 10,000 tossings of a fair coin, the lead never changes, is about .0085.*

8 *An amazing answer.* Here you are asked about the meaning of probability in a question. The answer to this question is given in the answer section, but you are not expected to know the answer here. But make a guess. *What is the probability that in 10,000 tosses, one side is in the lead for more than 9,930 tosses, while the other side is in the lead for fewer than 70 tosses?*

KEY WORDS FOR REVIEW

Frequency theory	Personal theory
Propensity theory	Logical theory

13 Personal Probabilities

How personal degrees of belief can be represented numerically by using imaginary gambles.

Chapters 1–10 were often *deliberately ambiguous* about different kinds of probability. That was because the basic ideas usually applied, across the board, to most kinds of probability.

Now we develop ideas that matter a lot for belief-type probabilities. They do not matter so much from the frequency point of view.

THE PROGRAM

There are three distinct steps in the argument, and each deserves a separate chapter.

- *This chapter* shows how you might use numbers to represent your degrees of belief.
- *Chapter 14* shows why these numbers should satisfy the basic rules of probability. (And hence they should obey Bayes' Rule.)
- *Chapter 15* shows how to use Bayes' Rule to *revise* or *update* personal probabilities in the light of new evidence. This is the fundamental motivation for the group of chapters, 13–15.

In these chapters we are concerned with a person's degrees of belief. We are talking about personal probabilities. But this approach can be used for other versions of belief-type probability, such as the logical perspective of Keynes and Carnap.

Because Bayes' Rule is so fundamental, this approach is often called *Bayesian.* "Belief dogmatists" are often simply called Bayesians because the use of Bayes' rule as a model of learning from experience plays such a large part in their philosophy. But notice that there are many varieties of Bayesian thinking. This perspective ranges from the personal to the logical.

THOUGHT EXPERIMENTS

We are concerned with degrees of belief. There is a close relation between belief and action. For example:

> You are walking in unfamiliar countryside. On the way home you come to an unmarked fork in a path. You do not know whether to go left or right. You favor going left but are uncertain. You decide, you take a calculated risk. You translate your belief into action. You vote with your feet. You go left rather than right.

How confident are you that this is a good decision? To what degree do you believe that the left fork goes home? A thought experiment provides a clue on how to represent degrees of belief using numbers.

For example:

> If you toss a fair coin to decide which way to go, you must think that it is as probable that home is to the left, as that it is to the right. Your personal probability, for each fork, is ½.
>
> But if you would roll a die, and go right if and only if you rolled a 6, then your personal probability, that home is to the right, must be about 1/6.

This chapter is about thought experiments. Not experiments for your instructor or your neighbor, but experiments for *you*.

FIRST THOUGHT EXPERIMENT: GIFTS

The first thought experiments that you will conduct involve choosing between "gambles." The gambles don't cost you anything. You take no risks.

> Think of being offered a small but desirable prize, a free gift.
>
> It does not matter whether the gift is $10, or a good grade in the course, or being sent roses next Valentine's Day by someone you care about.
>
> Pick your own gift.
>
> Be real. Write down here a small gift that you would like:

____

RISK-FREE GAMBLES

Now think of two possible events. Here we use a trite example—the weather on the first day of spring.

> Event 1: Heavy snow falls in Toronto next March 21.
>
> Event 2: Next March 21 in Toronto is a balmy and springlike day.

A gamble is a choice to be made between two possibilities. Say you have chosen $10 as your gift. Take your choice:

- $10 if event 1 occurs, and nothing otherwise, or
- $10 if event 2 occurs, and nothing otherwise.

Would you rather "gamble" on event 1 occurring? Or on event 2 occurring?

Be real. Write down your answer:

____

If you are indifferent between these two options, then you must think that events 1 and 2 are equally likely.

But if you would rather get your gift if event 1 occurs, you seem to think it is more likely that event 1 will occur, than event 2. You attach a higher personal probability to event 1, than to event 2.

Conversely, if you would rather get your gift if event 2 occurs, you seem to think it is more likely that event 2 will occur, than event 1. You attach a higher personal probability to event 2, than to event 1.

EVENTS AND PROPOSITIONS

We have just been using the event language. As usual, we can also use the proposition language. Consider:

> Proposition 3: The right claw of a healthy lobster is almost always larger than the left claw.
>
> Proposition 4: The left claw of a healthy lobster is almost always larger than the right claw.

Would you rather have your $10 if proposition 3 is true, or if proposition 4 is true? If you prefer (3) to (4), then you have a higher personal probability for (3) than for (4).

REALITY CHECK: SETTLE SOON

Gambles are worthwhile only if they are settled soon. I have no interest in a gamble on the weather, with the prospect of a gift of $10, a million years from now. Think of this:

5: Heavy snow falls in Toronto on January 1, 2075.
6: Next January 1 is balmy and springlike in Toronto.

I myself think that (5) is much more likely than (6), greenhouse effect or no.

All the same, I would rather have my gift if (6) occurs, than if (5) occurs. I will be dead by the time (5) is settled, so whatever happens, I myself win nothing, and by 2075 ten dollars may be worth nothing anyway. But if I choose (6), there is some small chance that I will get my $10 gift.

From now on, imagine that all gambles and bets are to be settled quite soon.

That may require *a lot* of imagination. Think of the lobsters. Propositions (3) and (4) say "almost always." How do we settle this bet? Or recall Pascal's wager. Think of this gamble:

7: God created exactly this Universe.
8: This universe evolved by chance from an uncaused Big Bang.

Some people—"creationists"—might prefer to gamble on getting their gift if (7) is true. Other people—atheists—might gamble on getting their gift if (8) is true. But few of us believe that the question will be settled in our lifetimes.

REALITY CHECK: UNRELATED PRIZES

The value of the prize used for comparing two events should not be affected by the occurrence of events.

For example, if your chosen gift were a new parka, you would value your prize more if it snowed on March 21 than if spring arrived. So that is not a suitable gift for thinking about events 1 and 2.

Likewise, "eternal bliss" is not a suitable gift for comparing the imaginary gamble involving (7) and (8). For if you think that (7) is very improbable, eternal life will not, for you, be a real possibility at all.

SECOND THOUGHT EXPERIMENT: USE FAIR COINS

So far we have been qualitative. We can use the same procedure to be quantitative.

Choose what you believe to be a fair coin. Think of the following gamble, assuming that the coin is tossed next March 21:

\$10 if (1): Heavy snow falls in Toronto next March 21, or
\$10 if (9): This fair coin falls heads when first tossed next March 21.

If you prefer option (1), then you must think it more likely that it will snow, than that the coin will fall heads. You assess the probability of snow on March 21 as greater than ½.

Suppose you prefer (1). Then try the following imaginary gamble:

\$10 if (1): Heavy snow falls in Toronto next March 21, or
\$10 if (10): This fair coin falls heads at least once in the first two tosses made next March 21.

The probability of (10) is ¾. If you prefer (10) over (1), then your personal probability for (1) is between ½ and ¾.

If you prefer (1) to (10), your personal probability for snow on March 21 is more than ¾.

If you are indifferent between (1) and (10), then your personal probability for snow in Toronto next March 21 is about ¾.

You can narrow down your range by choosing to compare to three coins thrown in a row. Or you can think of an urn with k green balls and $n - k$ red balls, to *calibrate* your beliefs as finely as you please.

Thus you can think of these experiments as using an artificial randomizer to calibrate your personal degrees of belief.

You can even explain when, in your personal opinion, a coin is fair. You are indifferent between betting on heads or tails, so long as the stakes and payoffs are the same.

REPRESENTING, NOT MEASURING

After the second thought experiment, we said that when people are indifferent between (1) and (10), then their personal probability is "*about* 3/4"? Can't we be exact?

No. Belief is not the sort of thing that can be measured exactly. This method can never determine your personal degrees of belief to very many places of decimals, but it can be as fine as makes any sense, for you. For example, suppose some people prefer

\$10 if (11): This fair coin falls seven heads in a row on the first seven tosses made next March 21,

to

\$10 if (12): Heavy snow falls in Columbia, South Carolina, next March 21.

But they prefer \$10 if (12) to

\$10 if (13): This fair coin falls eight heads in a row during the first eight tosses made next March 21.

Then we could calibrate their personal probability, p, of snow in Columbia next March 21:

$1/256 \leq p \leq 1/128$. Or $0.0039 \leq p \leq 0.0078$.

By this time we have lost any feel for the numbers. If we wanted to *represent* the personal probability by a number, we could say, perhaps, that the personal probability "is," say, 0.006. This is not because we have measured the person's degree of belief. It is because we might find it convenient to represent the person's degree of belief by a number.

BETTING

Our thought experiments never involve taking a risk. You will never lose, in any of these experiments. But now think of more risky situations. Imagine making bets, which you can win or lose.

You can imagine betting on almost anything, so long as you also imagine that the bet will be settled reasonably soon.

Imagine that you and I bet on an event or proposition A—any of the possibilities in (1) to (10) above will do.

Suppose we are betting for money (but any other utility will do).

Say you bet $X *on* A, and I bet $Y *against* A.

This means that if A occurs, or turns out to be true, you collect my $Y.

But if A does not occur, or turns out to be false, I collect your $X.

In such a bet, the sum of the two bets, $(X+Y), is called the stake.

BETTING RATES

Suppose that you and I are betting. You bet that it is going to snow on March 21. I bet with you that it is not going to snow. You bet $1 *on* snow. I bet $3 *against* snow.

If it snows, you win the whole $4, for a profit of $3. I lose $3.

If it does not snow, I get the $4, a gain of $1 for me, and a loss of $1 for you.

Under these conditions, your betting rate on snow is 1/4. My betting rate against snow is 3/4.

> Your betting rate = (Your bet) ÷ (Stake)

PAYOFF

The easiest way to display a bet is in a table called a payoff matrix. This table shows payoffs when the stake is $S, and your betting rate is p.

You bet $$p$S on A.

A person who bets against you bets $S(1−$p$) = $(S−$p$S).

The left column lists the possible outcomes of the bet: A occurs, ~A occurs.

	Payoff for bet on A, given B	Payoff for bet against A, given B
A	$(1−$p$)S	$−(1−$p$)S
~A	$−$p$S	$$p$S

In the future, we will usually omit the dollar signs.

THIRD THOUGHT EXPERIMENT

Suppose that neither you nor your friend knows whether event E will occur. I offer the following deal, for nothing. You can have either.

Option (E): A chance of winning (1−p)($10) if E occurs,

or

Option (~E): A chance of winning p($10) if E does not occur.

Now you have to divide the cake—choose p. Your friend then chooses either option (E) or option (~E). If in your opinion p is a fair betting rate, then neither option has an advantage.

If, after choosing p, you think that option (E) is preferable to option (~E), then you should increase p.

This argument is just a generalization of our technique for calibrating personal probabilities.

FAIR BETTING RATES

In choosing between those two options, you did not take a risk. But the same argument will apply to anyone who bets. Suppose that the stake is $10. If p is, in your opinion, the fair rate for betting on E, then you should be indifferent between:

- A bet *on* E at rate p: you win (1−p)($10) if E occurs, just as in option (E) above.
- A bet *against* E at rate (1−p): you win p($10) if E does not occur, just as in option (~E) above.

If you think that a betting rate of p is fair, you see no advantage in betting one way—on E at rate p—rather than the other—against E at rate (1−p).

A quick way to say this:

> The fair rate for betting against E is the *reverse* of the fair rate for betting on E.

NOT REAL-LIFE BETTING

Bettors want to profit from their bets. They don't offer $1 to win $3 if E occurs, and also offer $3 to win $1 if E does not occur. They want an *edge*.

So if they think that ~E is three times as likely as E, they offer $1 to win *more* than $3 if E occurs. They want to win $(3+$x$), so their rate for betting *on* E is

$1/(1+3+x)$. This is because they think this gives them an advantage. They want the bet to be "unfair"—in their favor.

On the other hand, if they are betting against E, and they put up $3, they want to win *more* than $1. So they offer $3 to win $(1+x) if E does not occur. So their rate for betting *against* E is $3/(1+3+x)$.

If you go to a bookmaker, and ask to bet on the next America's Cup challenge race for large yachts, you can get rates on:

The Australian yacht will win the next America's Cup.
The Australian yacht will not win the next America's Cup.

You can be sure that the two rates you are offered will not add up to 1!

FAIR MEANS NO EDGE

We are not going to the races. We are doing thought experiments, to find out how to represent our *own* degrees of belief. There is no "edge." A fair personal betting rate p is one where the person thinks there is no advantage in betting either way, at rate p on E, or rate $(1-p)$ against E.

ODDS

Gamblers don't talk about betting rates. They talk about the odds (against something happening).

You stake $1, hoping to make a profit of $3. Gamblers say the odds are 3 to 1 against your winning.

A betting rate is the ratio of your stake to the total amount staked, namely $1/4 = 1/(3+1)$. Thus if the odds against E are $y{:}x$, the betting rate on E is $x/(x+y)$.

Betting rates and odds are two ways to express the same basic idea. We use betting rates here because they look more like probabilities.

Moreover, gamblers just love to create a fantasy world of fancy talk. For example, betting on (American) football can use either point spreads (which use even odds) or money-line bets. In money-line bets you either "give" the odds or "take" the odds. If you are familiar with this talk, it is all very simple, but the rest of us just get confused. So we use exactly one idea, betting rates.

CONDITIONAL BETS

Betting rates correspond to personal probabilities. What about conditional probability? We need the idea of a *conditional bet.*

A conditional bet is a bet made on a condition, with all bets cancelled if the condition does not hold.

Examples:

A bet that the Toronto Maple Leafs will win the Stanley Cup, on the condition that they get into the final round.
A bet that it will snow in Toronto next March 21, on the condition that the temperature at the Toronto airport drops below 0°C that night.
A bet that you will go to Heaven, conditional on your following the practices of a certain religion.
A bet that within ten years there will be radical breakthroughs in the treatment of inheritable diseases, thanks to successes in the human genome project.

To stick to the weather, you bet $1, say, that it will snow conditional on a subzero temperature, and I bet $2 that it will not. The stake is $3.

If the temperature stays above zero, *the bet is off.* There is no winning or losing.

If the temperature drops below zero, *the bet is on,* and either you make a net gain of $2 (it snows) or I make a net gain of $1 (it doesn't).

CONDITIONAL BETTING RATES

Conditional betting rates are like ordinary betting rates.

Your betting rate on A, conditional on B, is:

(Your conditional bet) ÷ Stake.

In the above example:

Your conditional bet was $1.
My conditional bet was $2.

The stake was $3.

Your conditional betting rate was $1/3$.
My conditional betting rate was $2/3$.

To generalize:

If the stake is $S, and you bet on A conditional on B, at rate p, then you bet $$p$S, and:

If B does not occur, the bet is off.
If B and A occur, you win $$(1-p)$S.
If B occurs, but A does not occur, you lose $$p$S.

Likewise, if you bet *against* A, conditional on B, when the betting rate on A given B is p, then you bet $$(1-p)$S, and:

If B does not occur, the bet is off.
If B and A occur, you lose $$(1-p)$S.
If B occurs but A does not occur, you gain $$p$S.

THE CONDITIONAL PAYOFF MATRIX

Leaving out the dollar signs, here is a payoff matrix for conditional bets, when the betting rate on A given B is p, and when the total sum staked is S:

	Payoff for bet on A, given B	Payoff for bet against A, given B
A&B	$(1-p)S$	$-(1-p)S$
(~A)&B	$-pS$	pS
~B	0	0

THE ARGUMENT

The Bayesian program has three stages. This chapter has gone through the first stage. Chapters 14 and 15 will explain the second and third stages.

The first stage—this chapter—gets you to represent your personal degrees of belief by betting rates. It seems very simple.

The second stage—Chapter 14—has a rigorous argument.

At the third stage—Chapter 15—we seem to have a simple logical consequence of the first two stages.

In one sense, Chapter 15 is the hardest chapter, because it has a quite precise argument. But in another sense, the present chapter is the hardest one. This is because all the steps are merely plausible steps, quite slippery. You may have the feeling that the whole thing is a confidence trick! So we should say that what really makes the argument interesting is that there are many different versions of it, many more subtle than anything in this book, and all leading to the same conclusions.

EXERCISES

1 *Nuclear power.* You are offered the following gamble that costs you nothing: \$100 if

(a) Energy from commercial nuclear fusion becomes available towards the end of the twenty-first century.

Or

\$100 if a coin of probability p of heads falls heads on the next toss.

Could you realistically use this gamble to calibrate personal probabilities for (a)? Could you use a similar gamble to calibrate for (b)? \$100 if

(b) Nuclear energy is abandoned in North America within two years.

2 *Chocoholic.* Alice loves chocolate, and for her, a box of fine chocolates is a great gift. She is sitting in her doctor's waiting room. Could she realistically use this prize to calibrate her beliefs in these possibilities?

(a) She will be told by her physician to enter a weight-loss program immediately.

(b) She will be told by her physician that she can eat anything she wants, her metabolism can handle it.

3 *Intelligent aliens.* In the summer of 1996, after the announced discovery of evidence for "life" on a Martian meteorite, a large betting company named Ladbroke's "cut the odds on finding intelligent alien life within the next year from 250-to-1 to 50-to-1. For bettors to collect, the United Nations must confirm the existence of alien life forms capable of communication with Earth."

(a) Would a bet with these conditions satisfy the "settle-soon" requirement?

(b) Odds of 50-to-1 on intelligent alien life being confirmed within a year is the same as a betting rate of $1/51$: you give Ladbroke's \$1 and they will pay you back \$51 (net profit of \$50) if you win. Suppose that Ladbroke's allows you to bet the other way at the same odds: you can bet \$50 that intelligent alien life will not be confirmed, for a net profit of \$1 if you win. No one except a true alien-freak believes that it is at all likely that the UN will certify intelligent alien life within twelve months. Many would agree with skeptical Skuli and say that their personal betting rate could be as much as $999/1{,}000$ that the UN will not confirm intelligent alien life. Skuli says he would be willing to bet \$999 to win \$1 right now, if the payoff were immediate. Explain why Skuli, being a rational man, does not bet a mere \$50 to win \$1.

4 *Bets.* I bet \$9 that you will get an A in this course. You bet the opposite way. The stake is \$12. How much did you bet? What is your betting rate *against* getting an A?

5 *Raising the ante.* Now suppose the total stake, with the same betting rates, is \$100. You bet against your getting an A. What is your "payoff," if you do get an A?

6 *Fair bet.* If you consider 1/4 as the fair rate for betting that you get an A, what would you consider the fair rate for betting against your getting an A?

7 *Make-up tests.* In a large class with frequent tests, you are allowed to do a "make-up" if you missed a test because of illness.

You hear through the grapevine that make-up tests are harder than the original tests.

What might be your own personal betting rate that you will get at least a B on the next test?

What might be your betting rate conditional on your being sick and taking a make-up examination?

Draw a payoff matrix for the conditional bet, if the stake is \$10.

8 *"Not" reverses preferences.* We argued that if your fair personal betting rate for betting on E is p, then your rate for betting on ~E should be $1-p$. In general, you "reverse" your bet when you turn from E to ~E. Here are a few exercises to fix this idea. Let (1) and (2) be any events (or propositions). Suppose that you prefer

\$10 if (1) occurs, over
\$10 if (2) occurs.

Would you prefer

\$10 if (1) does not occur, over
\$10 if (2) does not occur?

KEY WORDS FOR REVIEW

Thought experiments	Betting rate
Degrees of belief	Conditional bets
Stake	Conditional betting rate
Representing by numbers	"Fair" betting rate

14 Coherence

Consistent personal betting rates satisfy the basic rules of probability. Consistency in this sense is called coherence.

Personal probabilities and betting rates are all very well. But thus far they have no structure, no rules—in fact, not much meaning. Now we give one argument that betting rates ought to satisfy the basic rules for probability. We have already had three thought experiments in Chapter 13. Here are two more.

FOURTH THOUGHT EXPERIMENT: SETS OF BETTING RATES

A group of beliefs can be represented by a set of betting rates.

> Imagine yourself ***advertising*** a set of betting rates. For each of the propositions A, B, . . . K in the set, you offer betting rates p_a, p_b, p_c, . . . , p_k.

In this imaginary game, you are prepared to bet, say,

> on A at rate p_a, or
> against A at rate $(1-p_a)$.

FIFTH THOUGHT EXPERIMENT: SIMPLE INCONSISTENCY

These are *personal* betting rates. Couldn't you choose any fractions you like?

Of course. But you might be inconsistent.

For example, suppose you are concerned with just two possibilities, B (for "below zero") and ~B:

B: On the night of next March 21, the temperature will fall below 0°C at the Toronto International Airport meteorological station.