2 Belief as Probability

2.1 Subjective and objective probability

Beliefs vary in strength. I believe that the 37 bus goes to Waverley station, and that there are busses from Waverley to the airport, but the second belief is stronger than the first. With some idealization, we can imagine that for any propositions *A* and *B*, a rational agent is either more confident in *A* than in *B*, more confident in *B* than in *A*, or equally confident in both. The agent's belief state then effectively sorts the propositions from 'least confident' to 'most confident', and we can represent a proposition's place in the ordering by a number between 0 ('least confident') and 1 ('most confident'). This number is the agent's degree of belief, or **credence**, in the proposition. My credence that the 37 bus goes to Waverley, for example, might be around 0.8, while my credence that there are busses from Waverley to the airport is around 0.95.

The core assumption that unifies "Bayesian" approaches to epistemology, statistics, decision theory, and other areas is that rational degrees of belief obey the formal rules of the probability calculus. For that reason, degrees of belief are also called **subjective probabilities** or even just **probabilities**. But this terminology can give rise to confusion because the word 'probability' has other, and more prominent, uses.

Textbooks in science and statistics often define probability as relative frequency. On this usage, the probability of an outcome is the proportion of that type of outcome in some base class of events. For example, on the textbook definition, to say that the probability of getting a six when throwing a regular die is $1/6$ is to say that the proportion of sixes in a large class of throws is (or converges to) $1/6$.

Another use of 'probability' is related to determinism. Consider a particular die in mid-roll. Could one, in principle, figure out how the die will land, given full information about its present physical state, the surrounding air, the surface on which it rolls, and so on? If yes, there's a sense in which the outcome is not a matter of probability. Quantum physics seems to suggest that the answer is no: that the laws

of nature together with the present state of the world only fix a certain probability for future events. This kind of probability is sometimes called 'chance'.

Chance and relative frequency are examples of **objective probability**. Unlike degrees of belief, they are not relative to an agent; they don't vary between you and me. You and I may have different opinions about chances or relative frequencies; but that would be an ordinary disagreement. At least one of us would be wrong. By contrast, if you are more confident that the die will land six than me, then your subjective probability for that outcome really is greater than mine.

In this course, when I talk about credence or subjective probability, I do not mean belief about objective probability. I simply mean degree of belief. Our Bayesian model here diverges from *frequentist* or *objectivist* models that define expected utility in terms of objective probability. In these models, the MEU Principle is then restricted to cases in which the agent knows the relevant objective probabilities. (I mentioned this under "Sources and Further Reading" in the previous chapter.) On the Bayesian conception of probability, the MEU Principle does not presuppose knowledge of probabilities; it only presupposes that the agent has a definite degree of belief in the relevant states.

2.2 Probability theory

What all forms of probability, objective and subjective, have in common is a certain abstract structure, a structure that is studied by the mathematical discipline of probability theory.

Mathematically, a **probability measure** is a certain kind of function – in the mathematical sense: a mapping – from some objects to real numbers. The objects that are mapped to numbers are usually called 'events', but in philosophy we call them '**propositions**'.

The main assumption probability theory makes about propositions (the objects that are assigned probabilities) is the following.

> **Booleanism**
>
> Whenever some proposition A has a probability (possibly 0), then so does its negation $\neg A$ ('not A'); whenever two propositions A and B both have a prob-

ability, then so does their conjunction $A \wedge B$ ('*A* and *B*') and their disjunction $A \vee B$ ('*A* or *B*').

(Here and henceforth, I use upper-case letters A, B, C, etc. as schematic variables for arbitrary propositions.)

In our application, Booleanism implies that if an agent has a definite degree of belief in some propositions, then she also has a definite degree of belief in any proposition that can be construed from these in terms of negation, conjunction, and disjunction.

What sorts of things are propositions? Probability theory doesn't say. In line with our discussion in the previous chapter, we will informally understand propositions as possible states of the world. This is not a formal definition, since I haven't defined 'possible state of the world'. But I'll make a few remarks that should help clarify what I have in mind.

Different sentences can represent the very same state of the world. Consider the current temperature in Edinburgh. I don't know what it is. One possibility (one possible state of the world) is that it is 10°C. There is also a possibility that it is 50°F. How are these related? Since 10°C = 50°F, the second possibility is not an *alternative* to the first. It is the very same possibility, expressed with a different unit. The sentences 'It is 10°C in Edinburgh' and 'It is 50°F in Edinburgh' are different ways of picking out the same (possible) state of the world.

Like sentences, possible states of the world can be negated, conjoined, and disjoined. The negation of the possibility that it is 10°C is the possibility that it is *not* 10°C. If we negate that negated state, we get back the original state: the possibility that it is *not not* 10°C is nothing but the possibility that it is 10°C. In general, if we understand propositions as possible states of the world, then logically equivalent propositions are not just equivalent, but identical.

Possible states of the world can be more or less specific. That the temperature is 10°C is more specific than that it is between 7°C and 12°C. It is often useful to think of unspecific states as sets of more specific states. We can think of the possibility that it is between 7°C and 12°C as a collection of several possibilities, perhaps as the set { 7°C, 8°C, 9°C, 10°C, 11°C, 12°C }. The unspecific possibility obtains just in case one of the more specific possibilities obtains. A maximally specific state is called a **possible world** (in philosophy, and an 'outcome' in many other disciplines).

We will sometimes model propositions as sets of possible worlds.

I should warn that the word 'proposition' has many uses in philosophy. In this course, all we mean by 'proposition' is 'object of credence'. And 'credence', recall, is a semi-technical term for a certain quantity in the model we are building. It is pointless to argue over the nature of propositions before we have spelled out the model in more detail. Also, by 'possible world' I just mean 'maximally specific proposition'. The identification of propositions with sets of possible worlds is not supposed to be an informative reduction.

Exercise 2.1 †

First a reminder of some terminology from set theory. The *intersection* of two sets A and B is the set of objects that are in both A and B. The *union* of two sets A and B is the set of objects that are in one or both of A and B. The *complement* of a set A is the set of objects that are not in A. A set A is a *subset* of a set B if all objects in A are also in B.
Now, assume propositions are modelled as sets of possible worlds. Then the negation $\neg A$ of a proposition A is the complement of A.

(a) What is the conjunction $A \wedge B$ of two propositions, in set theory terms?
(b) What is the disjunction $A \vee B$?
(c) Which logical relationship corresponds to a proposition A being a subset of a proposition B?

Exercise 2.2 ††

Not all objects of probability are possible states of the world. Booleanism entails that at least one object of probability is *impossible*. Can you explain why?

Let's continue with the mathematics of probability. A probability measure, I said, is a function from propositions to numbers that satisfies certain conditions. These conditions are called **probability axioms** or **Kolmogorov axioms**, because their canonical statement was given by the Russian mathematician Andrej Kolmogorov in 1933.

The Kolmogorov Axioms

   (i) For any proposition A, $0 \leq \text{Cr}(A) \leq 1$.
   (ii) If A is logically necessary, then $\text{Cr}(A) = 1$.
   (iii) If A and B are logically incompatible, then $\text{Cr}(A \vee B) = \text{Cr}(A) + \text{Cr}(B)$.

I used have 'Cr' here for the probability measure, as we will be mostly interested in subjective probability or credence. '$\text{Cr}(A)$' is read as 'the (subjective) probability of A' or 'the credence in A'. Strictly speaking, we should add subscripts, '$\text{Cr}_{i,t}(A)$', to make clear that subjective probability is relative to an agent i and a time t; but we're mostly dealing with statements that hold for all agents at all times, so we can omit the subscripts.

Understood as a condition on rational credence, axiom (i) says that credences range from 0 to 1: you can't have a degree of belief greater than 1 or less than 0. Axiom (ii) says that if a proposition is logically necessary – like *it is raining or it is not raining* – then it must have subjective probability 1. Axiom (iii) says that the subjective probability of a disjunction equals the sum of the probability of the two disjuncts, provided these are logically incompatible, meaning they can't be true at the same time. For example, since it can't be both 8°C and 12°C, your credence in the disjunctive proposition 8°C $\vee$ 12°C must be $\text{Cr}(8°C) + \text{Cr}(12°C)$.

We'll ask about the justification for these assumptions later. First, let's derive a few consequences.

2.3 Some rules of probability

Suppose your credence in the hypothesis that it is 8°C is 0.3. Then what should be your credence in the hypothesis that it is *not* 8°C? Answer: 0.7. In general, the probability of $\neg A$ is always 1 minus the probability of A:

The Negation Rule

$\text{Cr}(\neg A) = 1 - \text{Cr}(A)$.

This follows from the Kolmogorov axioms. Here is how. Let A be any proposition.

Then $A \lor \neg A$ is logically necessary. By axiom (ii),

$$Cr(A \lor \neg A) = 1.$$

Since A and $\neg A$ are logically incompatible, axiom (iii) tells us that

$$Cr(A \lor \neg A) = Cr(A) + Cr(\neg A).$$

Combining these two equations yields

$$1 = Cr(A) + Cr(\neg A).$$

From that, simple algebraic rearrangement give us the Negation Rule.

Next, we can prove that logically equivalent propositions always have the same probability.

The Equivalence Rule

If A and B are logically equivalent, then $Cr(A) = Cr(B)$.

Proof: Assume A and B are logically equivalent. Then $A \lor \neg B$ is logically necessary; so by axiom (ii),
$$Cr(A \lor \neg B) = 1.$$

Moreover, A and $\neg B$ are logically incompatible, so by axiom (iii),

$$Cr(A \lor \neg B) = Cr(A) + Cr(\neg B).$$

By the Negation Rule,
$$Cr(\neg B) = 1 - Cr(B).$$

Putting all this together, we have

$$1 = Cr(A) + 1 - Cr(B).$$

Subtracting $1 - Cr(B)$ from both sides yields $Cr(A) = Cr(B)$.

Above I mentioned that if we understand propositions as possible states of the world, then logically equivalent propositions are identical: $\neg\neg A$, for example, is

the same proposition as A. The Equivalence Rule shows that even if we had used a different conception of propositions that allows distinguishing between logically equivalent propositions, these differences would never matter to an agent's subjective probabilities. If an agent's credences satisfy the Kolmogorov axioms, then she must give the same credence to logically equivalent propositions.

Exercise 2.3 †††

Prove from Kolmogorov's axioms that $Cr(A) = Cr(A \wedge B) + Cr(A \wedge \neg B)$. (Like the proofs above, each step of your proof should either be an instance of an axiom, or an application of the rules we have already established, or it should follow from earlier steps by simple logic and algebra.)

Next, let's show that axiom (iii) generalizes to three disjuncts:

Additivity for three propositions

If A, B, and C are all incompatible with one another, then $Cr(A \vee B \vee C) = Cr(A) + Cr(B) + Cr(C)$.

Proof sketch: $A \vee B \vee C$ is equivalent (or identical) to $(A \vee B) \vee C$. If A, B, and C are mutually incompatible, then $A \vee B$ is incompatible with C. So by axiom (iii), $Cr((A \vee B) \vee C) = Cr(A \vee B) + Cr(C)$. Again by axiom (iii), $Cr(A \vee B) = Cr(A) + Cr(B)$. Putting these together, we have $Cr((A \vee B) \vee C) = Cr(A) + Cr(B) + Cr(C)$.

The argument generalizes to any finite number of propositions $A, B, C, D, ...$: the probability of a disjunction of n mutually incompatible propositions is the sum of the probability of the n propositions. This has the following consequence, which is worth remembering:

Probabilities from worlds

If the number of possible worlds is finite, then the probability of any proposition is the sum of the probability of the worlds at which the proposition is true.

Suppose two dice are tossed. There are 36 possible outcomes ("possible worlds"),

which we might tabulate as follows.

$$
\begin{array}{cccccc}
(1,1) & (1,2) & (1,3) & (1,4) & (1,5) & (1,6) \\
(2,1) & (2,2) & (2,3) & (2,4) & (2,5) & (2,6) \\
(3,1) & (3,2) & (3,3) & (3,4) & (3,5) & (3,6) \\
(4,1) & (4,2) & (4,3) & (4,4) & (4,5) & (4,6) \\
(5,1) & (5,2) & (5,3) & (5,4) & (5,5) & (5,6) \\
(6,1) & (6,2) & (6,3) & (6,4) & (6,5) & (6,6)
\end{array}
$$

Suppose you give equal credence $1/36$ to each of these outcomes or worlds. What credence should you then give to the hypothesis that both dice land on a number less than 4? Looking at the table, we can see that there are nine possible worlds at which the hypothesis is true: the top left quarter of the table. The hypothesis is equivalent to the *disjunction* of these possible worlds. Both dice land on a number less than 4 iff the outcome is (1,1) *or* (1,2) *or* (1,3) *or* (2,1) *or* (2,2) *or* (2,3) *or* (3,1) *or* (3,2) *or* (3,3). All of these outcomes are incompatible with one another. (The dice can't land (1,1) and (1,2) at the same time.) The rules of probability therefore tell us that the probability of our target hypothesis is the sum of the probability of the individual worlds. Since each world has probability $1/36$, and there are nine relevant worlds, your credence that both dice land on a number less then 4 should be $9 \cdot 1/36 = 1/4$.

> **Exercise 2.4** †
>
> What credence should you give to the following propositions, in the scenario with the two dice?
>
> (a) At least one die lands 6.
> (b) Exactly one die lands 6.
> (c) The sum of the numbers that will come up is equal to 5.

Some thorny technical problems arise if there are infinitely many worlds. It would be nice if we could say that the probability of a proposition is always the sum of the probability of the worlds that make up the proposition. If there are too many worlds, however, this turns out to be incompatible with the mathematical structure of the real numbers. The most one can safely assume is that the principle holds if the number of worlds is *countable*, meaning that there are no more worlds than there are natural numbers 1,2,3,…. To secure this, axiom (iii) – which is known as the axiom of

Finite Additivity – has to be replaced by an axiom of **Countable Additivity**. In this course, we will try to stay away from troubles arising from infinities, so for our purposes the weaker axiom (iii) will be enough.

> **Exercise 2.5 †††**
>
> Prove from Kolmogorov's axioms that if A entails B, then $\mathrm{Cr}(A)$ cannot be greater than $\mathrm{Cr}(B)$. (You may use the rules we have already derived.)

2.4 Conditional probability

To continue, we need two more concepts. The first is the idea of **conditional probability** or, more specifically, **conditional credence**. Intuitively, an agent's conditional credence reflects her degree of belief in a given proposition on the supposition that some other proposition is true. For example, I am fairly confident that it won't snow tomorrow, and that the temperature will be above 4°C. Yet, on the supposition that it will snow, I am not at all confident that the temperature will be above 4°C. My *unconditional credence* in temperatures above 4°C is high, but my *conditional credence* in the same proposition, on the supposition that it will snow, is low.

Conditional credence relates two propositions: the proposition that is supposed, and the proposition that gets evaluated on the basis of that supposition.

To complicate things, there are actually two kinds of supposition, and two kinds of conditional credence. The two kinds of supposition correspond to a grammatical distinction between "indicative" and "subjunctive" conditionals. Compare the following statements.

(1) If Shakespeare didn't write *Hamlet*, then someone else did.

(2) If Shakespeare hadn't written *Hamlet*, then someone else would have.

The first of these (an indicative conditional) is highly plausible: we know that someone wrote *Hamlet*; if it wasn't Shakespeare then it must have been someone else. The second statement (a subjunctive conditional), is plausibly false: if Shakespeare hadn't written *Hamlet*, it is unlikely that somebody else would have stepped in to write the very same play.

The two conditionals (1) and (2) relate the same two propositions – the same possible states of the world. To evaluate either statement, we suppose that our world

is a world in which Shakespeare didn't write *Hamlet*. The difference lies in what we hold fixed when we make that supposition. To evaluate (1), we hold fixed our knowledge that *Hamlet* (the play) exists. Not so in (2). To evaluate (2), we bracket everything we know that we take to be a causal consequence of Shakespeare's writing of *Hamlet*.

We will return to the second, subjunctive kind of supposition in section 9. For now, let's focus on the first, indicative kind of supposition. I will write $\mathrm{Cr}(A/B)$ for the (indicative) conditional credence in A on the supposition that B. Again, intuitively this is the agent's credence that A is true *if* (or *given that* or *supposing that*) B is true.

(Some authors use a vertical slash '|'. Either way, the slash is not a connective. $\mathrm{Cr}(A/B)$ is not the agent's credence in a special proposition designated by 'A/B'.)

How are conditional credences related to unconditional credences? The answer is surprisingly simple, and captured by the following formula.

The Ratio Formula

$$\mathrm{Cr}(A/B) = \frac{\mathrm{Cr}(A \wedge B)}{\mathrm{Cr}(B)}, \text{ provided } \mathrm{Cr}(B) > 0.$$

That is, your credence in some proposition A on the (indicative) supposition B equals your unconditional credence in $A \wedge B$ divided by your unconditional credence in B.

To see why this makes sense, it may help to imagine your credence as distributing a certain quantity of "plausibility mass" over the space of possible worlds. When we ask about your credence in A conditional on B, we set aside worlds where B is false. What we want to know is how much of the mass given to B worlds falls on A worlds. In other words, we want to know what fraction of the mass given to B worlds is given to $A \wedge B$ worlds.

People disagree on the status of the Ratio Formula. Some treat it as a definition. On that approach, you can ignore everything I said about what it means to suppose a proposition and simply read '$\mathrm{Cr}(B/A)$' as shorthand for '$\mathrm{Cr}(A \wedge B)/\mathrm{Cr}(A)$'. Others regard conditional beliefs as distinct and genuine mental states and see the Ratio Formula as a fourth axiom of probability. We don't have to adjudicate between these views. What matters is that the Ratio Formula is true, and on this point both sides agree.

The second concept I want to introduce is that of probabilistic independence. We say that propositions *A* and *B* are **(probabilistically) independent** (for the relevant agent at the relevant time) iff $\text{Cr}(A/B) = \text{Cr}(A)$. Intuitively, if *A* and *B* are independent, then it makes no difference to your credence in *A* whether or not you suppose *B*, so your unconditional credence in *A* is equal to your credence in *A* conditional on *B*.

Unlike causal independence, probabilistic independence is a feature of beliefs. Two propositions can be independent for one agent and not for another. That said, there are interesting connections between probabilistic (in)dependence and causal (in)dependence. For example, if an agent knows that two events are causally independent, then the events are often also independent in the agent's degrees of belief. You may want to ponder why that is the case.

Exercise 2.6 †

Assume $\text{Cr}(Snow) = 0.3$, $\text{Cr}(Wind) = 0.6$, and $\text{Cr}(Snow \wedge Wind) = 0.2$. What is $\text{Cr}(Snow/Wind)$? What is $\text{Cr}(Wind/Snow)$?

Exercise 2.7 ††

Using the Ratio Formula and the Equivalence Rule, show that if *A* is (probabilistically) independent of *B*, then *B* is independent of *A* (assuming that $\text{Cr}(A)$ and $\text{Cr}(B)$ are greater than 0).

Exercise 2.8 ††

A fair die will be tossed, and you give equal credence to all six outcomes. Let *Ex* be the proposition that the die lands 1 or 6. Let *Odd* be the proposition that the die lands an odd number (1, 3, or 5), and let *Low* be the proposition that the die lands 1, 2 or 3. Which of the following are true, in your belief state?

  (a) *Ex* is independent of *Odd*.
  (b) *Odd* is independent of *Ex*.
  (c) *Ex* is independent of *Low*.
  (d) *Odd* is independent of *Low*.
  (e) *Ex* is independent of *Odd* ∧ *Low*.

2.5 Some more rules of probability

If you've studied propositional logic, you'll know how to compute the truth-value of arbitrarily complex sentences from the truth-value of their atomic parts. For example, you can figure out that if A and B are true and C is false, then $A \wedge \neg(B \vee \neg(C \vee A))$ is false. Now suppose instead of the truth-value of A, B, and C, I give you their probability. Could you compute the probability of $A \wedge \neg(B \vee \neg(C \vee A))$? The answer is no. In general, while the probability of $\neg A$ is determined by the probability of A (as we know from the Negation Rule), neither the probability of $A \vee B$ nor the probability of $A \wedge B$ is determined by the individual probabilities of A and B.

Let's have a look at conjunctive propositions, $A \wedge B$. By rearranging the Ratio Formula, we get the following:

The Conjunction Rule

$\mathrm{Cr}(A \wedge B) = \mathrm{Cr}(A) \cdot \mathrm{Cr}(B/A)$.

So the probability of a conjunction is the probability of the first conjunct times the probability of the second *conditional on the first*. If you only know the unconditional probabilities of the conjuncts, you can't figure out the probability of the conjunction.

But there's a special case. If A and B are independent, then $\mathrm{Cr}(B/A) = \mathrm{Cr}(B)$. In that case, the probability of the conjunction is the product of the probability of the conjuncts:

The Conjunction Rule for independent propositions

If A and B are independent, then $\mathrm{Cr}(A \wedge B) = \mathrm{Cr}(A) \cdot \mathrm{Cr}(B)$.

Why do we multiply (rather than, say, add) the probabilities in the Conjunction Rules? Suppose we flip two coins. What is the probability that they both land heads? You'd expect the first coin to land heads about half the time; and in half *of those* cases you'd expect the second to also land heads. The result is a half of a half. And half of a half is $1/2$ times $1/2$.

What about disjunctions, $A \vee B$? We know that if A and B are logically incompatible, then $\mathrm{Cr}(A \vee B) = \mathrm{Cr}(A) + \mathrm{Cr}(B)$. What if A and B are not incompatible? In

that case, we have to subtract the probability of the conjunction:

> **The Disjunction Rule**
>
> $\mathrm{Cr}(A \lor B) = \mathrm{Cr}(A) + \mathrm{Cr}(B) - \mathrm{Cr}(A \land B)$.

Again, you can't compute the probability of the disjunction just from the probability of the disjuncts.

Why do we subtract $\mathrm{Cr}(A \land B)$ in the Disjunction Rule? The proposition $A \lor B$ comprises three kinds of worlds: (1) worlds where A is true and B is false, (2) worlds where B is true and A is false, and (3) worlds where A and B are both true. These three sets are disjoint (mutually exclusive). By Additivity, the probability of the disjunction $A \lor B$ equals the probability of $A \land \neg B$ plus the probability of $B \land \neg A$ plus the probability of $A \land B$. Taken together, the worlds in (1) and (3) comprise precisely the A-worlds, and the worlds in (2) and (3) comprise the B-worlds. So if we add together $\mathrm{Cr}(A)$ and $\mathrm{Cr}(B)$, we have effectively double-counted the $A \land B$ worlds. That's why we need to subtract $\mathrm{Cr}(A \land B)$.

> **Exercise 2.9 †**
>
> Show that two propositions A and B with positive probability are independent if and only if $Cr(A \land B) = Cr(A) \cdot Cr(B)$. (Some authors use this as the definition of independence.)

> **Exercise 2.10 ††**
>
> Prove from the Ratio Formula that $\mathrm{Cr}(A \land B \land C) = \mathrm{Cr}(A/B \land C) \cdot \mathrm{Cr}(B/C) \cdot \mathrm{Cr}(C)$. (This is known as the *Chain Rule*, and generalizes to more than three conjuncts.)

> **Exercise 2.11 †**
>
> In 1999, a British woman was convicted of the murder of her two sons, who she claimed died from Sudden Infant Death Syndrome (SIDS). The eminent paediatrician Sir Roy Meadow explained to the jury that 1 in 8500 infants die

from SIDS and hence that the chance of SIDS affecting both sons was 1/8500 · 1/8500 = 1 in 73 million. What is wrong with Sir Meadow's reasoning?

I want to mention two more rules that play a special role in Bayesian accounts. The first goes back to a suggestion by Thomas Bayes published in 1763.

Bayes' Theorem

$$\mathrm{Cr}(A/B) = \frac{\mathrm{Cr}(B/A) \cdot \mathrm{Cr}(A)}{\mathrm{Cr}(B)}$$

Proof: By the Ratio Formula, $\mathrm{Cr}(A/B) = \mathrm{Cr}(A \wedge B)/\mathrm{Cr}(B)$. By the Conjunction Rule, $\mathrm{Cr}(A \wedge B) = \mathrm{Cr}(B/A) \cdot \mathrm{Cr}(A)$. So we can substitute $\mathrm{Cr}(A \wedge B)$ in the Ratio Formula by $\mathrm{Cr}(B/A) \cdot \mathrm{Cr}(A)$, which yields Bayes' Theorem.

Bayes' Theorem relates the conditional probability of A given B to the inverse conditional probability of B given A. Why that might be useful is best illustrated by an example.

Suppose you are unsure whether the die I am about to roll is a regular die or a trick die that has a six printed on all sides. You currently give equal credence to both possibilities. How confident should you be that the die is a trick die *given that it will land six on the next roll*? That is, what is $\mathrm{Cr}(\textit{Trick}/\textit{Six})$? The answer isn't obvious. Bayes' Theorem helps. By Bayes' Theorem,

$$\mathrm{Cr}(\textit{Trick}/\textit{Six}) = \frac{\mathrm{Cr}(\textit{Six}/\textit{Trick}) \cdot \mathrm{Cr}(\textit{Trick})}{\mathrm{Cr}(\textit{Six})}.$$

The numerator on the right is easy. $\mathrm{Cr}(\textit{Six}/\textit{Trick})$ is 1: if the die has a six on all its sides then it is certain that it will land six. We also know that $\mathrm{Cr}(\textit{Trick})$ is $1/2$. But what is $\mathrm{Cr}(\textit{Six})$, your unconditional credence that the die will land six? Here we need one last rule:

The Law of Total Probability

$$\mathrm{Cr}(A) = \mathrm{Cr}(A/B) \cdot \mathrm{Cr}(B) + \mathrm{Cr}(A/\neg B) \cdot \mathrm{Cr}(\neg B).$$

This follows immediately from exercise 2.3 and the Conjunction Rule.

If we apply the Law of Total Probability to Cr(*Six*) in the above application of Bayes' Theorem, we get

$$\mathrm{Cr}(Trick/Six) = \frac{\mathrm{Cr}(Six/Trick) \cdot \mathrm{Cr}(Trick)}{\mathrm{Cr}(Six/Trick) \cdot \mathrm{Cr}(Trick) + \mathrm{Cr}(Six/\neg Trick) \cdot \mathrm{Cr}(\neg Trick)}.$$

It looks scary, but all the terms on the right are easy to figure out. We already know that Cr(*Six*/*Trick*) = 1 and that Cr(*Trick*) = 1/2. Moreover, Cr(*Six*/¬*Trick*) is plausibly 1/6 and Cr(¬*Trick*) is 1/2. Plugging all these values into the formula, we get Cr(*Trick*/*Six*) = 6/7. Your credence in the trick die hypothesis conditional on seeing a six should be 6/7.

Exercise 2.12 †††

A stranger tells you that she has two children. You ask if at least one of them is a boy. The stranger says yes. How confident should you be that the other child is also a boy? (Assume there are only two sexes, which are equally common and independent among siblings.)

Essay Question 2.1

If an agent's degrees of belief satisfy the probability axioms, it seems to follow from Kolmogorov's axiom (ii) that the agent must be certain of every logical truth. Does this mean that our Bayesian model is inapplicable to ordinary agents, who are not logically omniscient? If so, is this a problem? Do you have an idea of how the model could be adjusted to allow for logical non-omniscience?

Sources and Further Reading

There are many good introductions to elementary probability theory. For a slightly more in-depth discussion of the topics we have covered, you may want to consult chapters 3–7 of Ian Hacking, *An Introduction to Probability and Inductive Logic* (2001). (You may find the rest of the book helpful as well.)

The problems infinitely many worlds raise for the Additivity axiom are nicely explained in Brian Skyrms, "Zeno's paradox of measure" (1983).

The topic of the essay question is commonly discussed as the "problem of logical omniscience". See, for example, Zeynep Soysal, "A metalinguistic and computational approach to the problem of mathematical omniscience" (2022) for an interesting recent proposal with pointers to the earlier discussion.