PE3101P: Decision and Social Choice

Lecture Notes 2 — Probability

Tomi Francis

20 August 2026

1 Probabilities: the basics

1.1 Sample spaces and events

We begin with a *sample space* Ω, which is just a collection of all of the possible outcomes over which we are uncertain in the given situation.

For example, it could be the possible results of rolling a die, or coin flips, different ways the weather could go, or the decisions somebody else is going to make without your input.

A sample space is the set of all of these outcomes. By a "set", I just mean a collection of other things. I'll have to talk about these quite a lot, so you may want to go and take a look at the *set theory primer* now available on my website if you want to get more familiar with them.

A set of some of the outcomes is called an event. For example, an event could be the die rolling 2 *or* 3, or it could be the weather being rainy *or* overcast. An event can also contain just a single outcome, and in this case although an event is *formally* distinct from an outcome, we can treat them like they're the same thing. (For example, the event that the die rolls 1).

An event can also be that *any* of the possible outcomes happens. In this case, it is equal to the sample space Ω. This is why I called it the *universal event* last week. The sample space and the universal event are the same thing: the set of all of the things that could happen.

The set containing no outcomes at all is also an event, albeit one that can never happen. We denote this by $\varnothing$ and we can call it the empty event.

1.2 Combining events

If P and Q are events, we can write $P \cup Q$ for the event that something in either P *or* Q (or both) happens. This is a set-theoretic union. But, if you like, you can also write $P \vee Q$. Technically this is an abuse of notation (more on that later), but we won't worry about it.

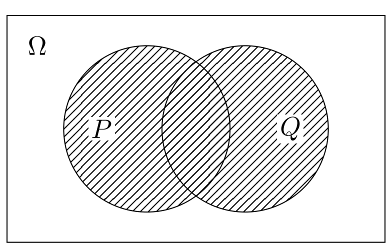

Figure 1: The union $P \cup Q$: everything in P, in Q, or in both.

Check your understanding: unions

(1) If event P is that the die rolls an even number, and event Q is that the die rolls an odd number, what is $P \cup Q$?

(2) If P is that it will rain today, and Q is that it will rain hard today, what is $P \cup Q$?

(3) If P is that my opponent's hand contains the Queen of Spades, and Q is that my opponent's hand contains the Jack, King or Ace of Spades, what is $P \cup Q$?

(4) If P is that it will rain today, and Q is that it will rain today or tomorrow, what is $P \cup Q$?

Similarly, for events P and Q we can write $P \cap Q$ for the event that something in *both* P and Q happens. This is a set-theoretic intersection. Similarly to before, you can write $P \wedge Q$ or $P \& Q$ instead if you like.

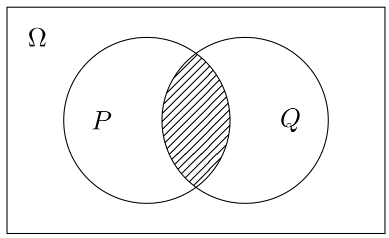

Figure 2: The intersection $P \cap Q$: everything in both P and Q.

Check your understanding: intersections

(1) If P is that the die rolls an even number, and Q is that the die rolls an odd number, what is $P \cap Q$?

(2) If P is that it will rain today, and Q is that it will rain hard today, what is $P \cap Q$?

(3) If P is that the card drawn is a heart, and Q is that it is a King, what is $P \cap Q$?

(4) If P is that the die rolls an even number, and Q is that it rolls higher than 3, what is $P \cap Q$?

Finally, any event P has a complement P^c, which we will often denote by $\neg P$. This is the event that contains all the outcomes which are in the sample space but not in P.

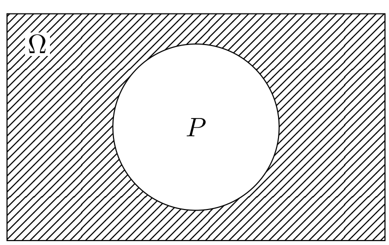

Figure 3: The complement P^c: everything in Ω that is not in P.

Check your understanding: complements

(1) If P is that the die rolls a 6, what is $\neg P$?

(2) If P is that the die rolls an even number, what is $\neg P$?

(3) If P is that it will rain today, what is $\neg P$?

1.3 Set-theoretic vs logical operators

In general, and officially, I take the sample space to be a set of outcomes, and events to be *subsets* (sets which might be smaller) of those outcomes. The proper way to combine these kinds of events to get other events is to use the set theoretic notions of union and intersection. The union of sets X and Y is the set that contains everything in *either* X or Y, and the intersection of X and Y is the set that contains everything in *both* X and Y. (For much more detail on this, see the primer.)

Other authors, including Schwarz in BDRC, prefer to talk about propositions, and to use the logical notions of conjunction ($\wedge$) and disjunction ($\vee$)—I just mean "and" and "or".[1]

Technically, these are different. "And" and "or" are truth-functional conditions which take truth-apt objects like sentences or propositions and return further sentences or propositions whose truth values are grounded in the truth values of their component parts. "Intersection" and "union" are set-theoretic operators which combine two sets to form another set.

They do different things. But they are very tightly related and behave exactly the same way in most contexts, so in this course we will mostly ignore the differences between them. In other words, I'm not going to penalise you if you write $\wedge$ instead of $\cap$.

1.4 On non-measurable sets

Generally in mathematics a *Probability Space* comes with a sample space Ω and a specification of events Σ—a list of all those subsets of Ω which "count" as events. (It also has what's called a *probability measure*, which we're just about to talk about.)

The reason why the specification of events Σ has to exist in general is that *in general, not every subset of a sample space can coherently be assigned a probability, if we want certain intuitively plausible regularities to hold*. The reasons why are quite complicated. For our purposes, we can ignore the problem, and we can generally assume that *any* set of outcomes counts as an event.

1.5 Mutual exclusivity, joint exhaustiveness

Sometimes, two events have no outcome in common at all. When events X and Y are like this, we say that they are *mutually exclusive* or *disjoint*, and the condition for it is that $X \cap Y = \varnothing$, the empty event.

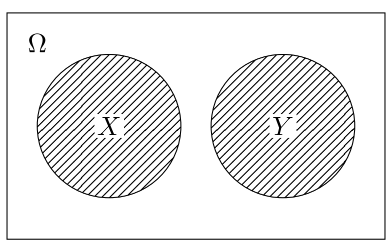

Figure 4: Mutually exclusive (disjoint) events: $X \cap Y = \varnothing$.

Sometimes, two or more events are such that it is guaranteed that at least one will happen. In this case, we say that the events are *jointly exhaustive*, and the condition for it is that $X_1 \cup \cdots \cup X_n = \Omega$, i.e., every outcome is in at least one of the events in question.

[1] Strictly speaking, since Schwarz characterises propositions as sets of possible worlds, these actually amount to exactly the same thing as the set-theoretic notions of intersection and union anyway.

And when a set of events are mutually exclusive and jointly exhaustive, we say that they *partition* the sample space.

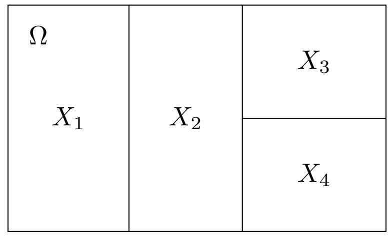

Figure 5: A partition of Ω: the cells are pairwise disjoint *and* jointly exhaustive, so exactly one of them obtains.

Remember from last week: we want to also require that the set of states of nature is mutually exclusive and jointly exhaustive. So, the states of nature partition the sample space for each act.

1.6 Probabilistic independence

Definition (Probabilistic independence)**.** Events A and B are *probabilistically independent* just in case

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

1.7 Probabilistic independence is not mutual exclusivity

These two are easy to run together, but they are actually close to opposites.

Suppose A and B are mutually exclusive, and both have positive credence. Then $A \cap B = \varnothing$, so $\mathrm{Cr}(A \cap B) = 0$, while $\mathrm{Cr}(A) \cdot \mathrm{Cr}(B) > 0$. So they are *not* probabilistically independent.

In fact they are as dependent as two events can be: learning that A happened takes your credence in B straight to zero.

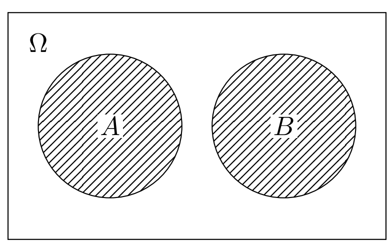

Figure 6: Mutually exclusive: A and B cannot both happen, so learning that A happened rules B out entirely. This is the opposite of probabilistic independence, not an instance of it.

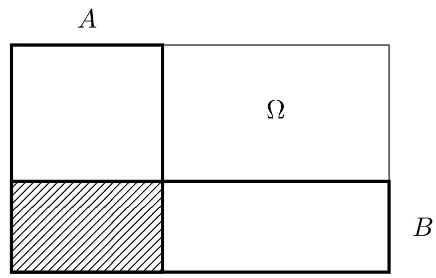

Figure 7: Probabilistically independent: B takes up the same fraction of A as it does of Ω. The shaded cell is $A \cap B$.

2 Probabilism

2.1 Assigning Probabilities

We now turn to assigning probabilities to events, and to what the constraints on doing so are. In particular, we're interested in the structure of credence: the constraints placed by epistemic rationality on how our credences should look.

2.2 Epistemic irrationality

Intuitively, some patterns of credence we can have are incorrect.

Have a think about each of these, and about what is wrong with having credences like this.

Example 1. Joe thinks that there is a 50% chance that it will rain today and a 20% chance that it will not rain today.

Example 2. Biqing thinks that there is a 60% chance she failed the exam she just took, and a 70% chance she failed it badly.

Example 3. Homer thinks that there's a 30% chance he'll drink one beer, a 50% chance he'll drink more than one beer, and a 70% chance he'll drink at least one beer.

2.3 The Probability Axioms (again)

Lecture 1 stated these informally:

(i) **Non-negativity.** Probabilities can't be negative.

(ii) **Normalisation.** The probability of the universal event—of *something* happening—is 1.

(iii) **Additivity.** The probability of *either* of two completely separate events happening is the sum of their probabilities. (E.g., the probability of a rolled die landing 1 *or* 2 is $1/6 + 1/6 = 1/3$.)

We now have names for all the pieces. Ω is the sample space, Σ is the set of events, and $\mathrm{Cr}(\sigma)$ is the credence assigned to an event $\sigma \in \Sigma$. "Completely separate" is disjointness, which we defined above as $X \cap Y = \varnothing$.

(i) **Non-negativity:** $\mathrm{Cr}(\sigma) \geq 0$ for every $\sigma \in \Sigma$.

(ii) **Normalisation:** $\mathrm{Cr}(\Omega) = 1$.

(iii) **Finite additivity:** if $X \cap Y = \varnothing$, then $\mathrm{Cr}(X \cup Y) = \mathrm{Cr}(X) + \mathrm{Cr}(Y)$.

More generally, if the events $\sigma_0, \sigma_1, \ldots$ are pairwise disjoint, then $\mathrm{Cr}(\bigcup_i \sigma_i) = \mathrm{Cr}(\sigma_0) + \mathrm{Cr}(\sigma_1) + \ldots$

2.4 So what went wrong?

We can now say exactly which axiom each of the three cases above breaks.

> **Joe** — *normalisation*, via additivity. Rain and no-rain are disjoint and between them exhaust Ω, so by additivity their credences must sum to $\mathrm{Cr}(\Omega)$, which normalisation fixes at 1. His sum to 0.7.
>
> **Biqing** — *additivity*, together with non-negativity. Failing badly is a *way of* failing, so failing badly is a subset of failing, and failing splits into failing badly and failing-but-not-badly. Additivity makes Cr(fail) the sum of those two, and non-negativity keeps the second one from being negative; so failing badly cannot be more probable than failing. Hers is 0.7 against 0.6.
>
> **Homer** — *finite additivity*. One beer and more than one beer are disjoint, and together they are "at least one beer". So that credence should be $0.3 + 0.5 = 0.8$. His is 0.7.

3 Conditional probabilities and Conditionalisation

3.1 Defining conditional probabilities

Definition (Conditional probability)**.**

$$\mathrm{Cr}(A \mid B) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(B)} \qquad \text{provided } \mathrm{Cr}(B) > 0.$$

In these diagrams areas represent probabilities, and in general probabilities behave very much like areas. When we suppose B, the new area lying in A is $A \cap B$; but the total area over which we are calculating is now the area of B, not the area of Ω, which was 1. That is why, to find how likely A is on the condition that B, we take its new area, $A \cap B$, and divide it by the new total area, B.

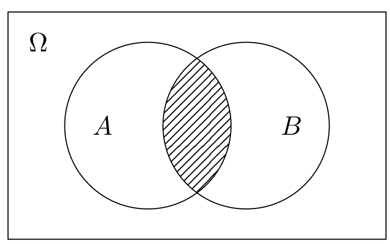

Figure 8: Before conditioning: Ω is everything, and the shaded region $A \cap B$ is measured against all of it.

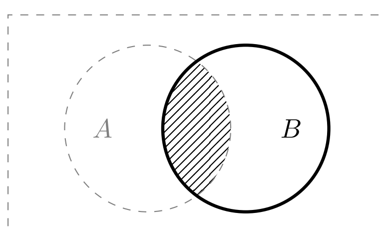

Figure 9: After conditioning on B: everything outside B is discarded, and the same shaded region is now measured against B alone.

So the formula does exactly that: $\mathrm{Cr}(A \cap B)$ measures the overlap, and dividing by $\mathrm{Cr}(B)$ rescales it so that B, rather than Ω, counts as the whole. It is also why we require $\mathrm{Cr}(B) > 0$: if B has no probability, there is nothing to rescale against.

This also gives us a second, equivalent way of saying that two events are *probabilistically independent*, which we met back in section 1. Where $\mathrm{Cr}(B) > 0$,

$$\mathrm{Cr}(A \cap B) = \mathrm{Cr}(A) \cdot \mathrm{Cr}(B) \qquad \text{is the same as} \qquad \mathrm{Cr}(A \mid B) = \mathrm{Cr}(A).$$

The second form is arguably the more intuitive one, since it says directly that learning whether B happened makes no difference at all to how likely A is.

3.2 Conditionalisation

Principle 1 (Conditionalisation)**.** *Suppose your credences are given by* Cr*, and then you learn E and nothing more. Your new credence in any event A should be your old credence in A conditional on E:*

$$\mathrm{Cr}_{new}(A) = \mathrm{Cr}(A \mid E).$$

3.3 A visual explanation of conditionalisation

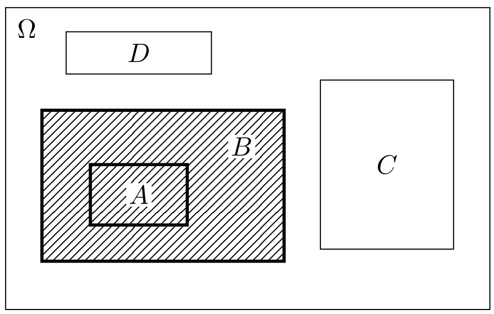

Figure 10: Stage 1. The conditioning event B shaded.

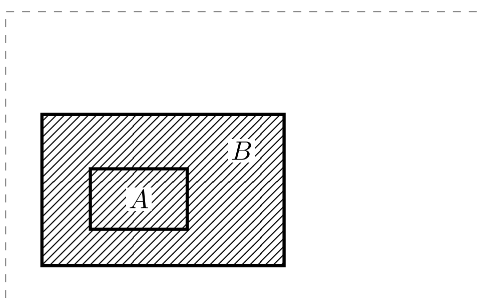

Figure 11: Stage 2. Everything outside B is discarded. Nothing has moved and nothing has changed size yet.

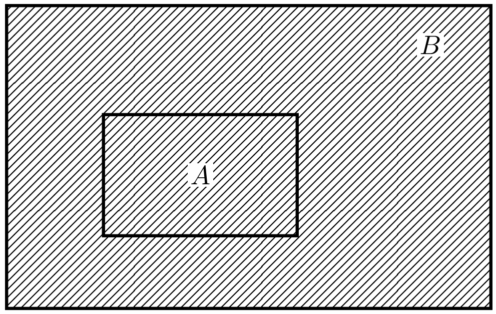

Figure 12: Stage 3. B blown up until it fills the space Ω used to. The scaling is uniform, so every ratio of areas inside B is unchanged.

3.4 The multiplication rule and total probability

Sometimes it is useful to use conditional probabilities to calculate the unconditional probability of an event. Suppose that the sample space Ω can be partitioned into $B_1, B_2, \ldots, B_n$. Then the *law of total probability* states that

$$\mathrm{Cr}(A) = \mathrm{Cr}(B_1)\,\mathrm{Cr}(A \mid B_1) + \mathrm{Cr}(B_2)\,\mathrm{Cr}(A \mid B_2) + \cdots + \mathrm{Cr}(B_n)\,\mathrm{Cr}(A \mid B_n).$$

Why is this true? Note that, since the B_i partition Ω,

$$\mathrm{Cr}(A) = \mathrm{Cr}(A \cap B_1) + \mathrm{Cr}(A \cap B_2) + \cdots + \mathrm{Cr}(A \cap B_n).$$

Then try substituting the definition of conditional probability into the total probability formula.

Exercise 1. Prove the law of total probability.

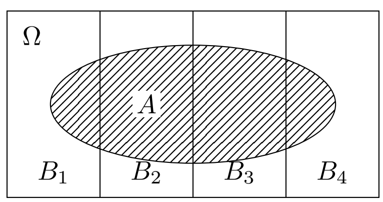

Figure 13: The law of total probability. The B_i partition Ω, so A is cut into the disjoint pieces $A \cap B_i$, and their probabilities sum to $\mathrm{Cr}(A)$.

3.5 Bayes' theorem

What we have, straight from the definition, is a pair of equations:

$$\mathrm{Cr}(A \mid B) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(B)} \qquad \text{and} \qquad \mathrm{Cr}(B \mid A) = \frac{\mathrm{Cr}(A \cap B)}{\mathrm{Cr}(A)}.$$

What we want is a way of getting from one conditional probability to the other:

Principle 2 (Bayes' theorem)**.**

$$\mathrm{Cr}(A \mid B) = \frac{\mathrm{Cr}(B \mid A)\ \mathrm{Cr}(A)}{\mathrm{Cr}(B)} \qquad \textit{provided } \mathrm{Cr}(A), \mathrm{Cr}(B) > 0.$$

Sometimes you know one of these conditional probabilities and want the other, and Bayes' theorem is what gets you across.

Exercise 2. Prove Bayes' theorem.

3.6 Bayes factors

Suppose you are weighing two hypotheses, H_1 and H_2, and some evidence E arrives. The *Bayes factor* is

$$\frac{\mathrm{Cr}(E \mid H_1)}{\mathrm{Cr}(E \mid H_2)}.$$

It measures how much better E is explained by H_1 than by H_2, and it is the whole of what the evidence contributes. That is easiest to see in odds form:

$$\underbrace{\frac{\mathrm{Cr}(H_1 \mid E)}{\mathrm{Cr}(H_2 \mid E)}}_{\text{posterior odds}} = \underbrace{\frac{\mathrm{Cr}(E \mid H_1)}{\mathrm{Cr}(E \mid H_2)}}_{\text{Bayes factor}} \times \underbrace{\frac{\mathrm{Cr}(H_1)}{\mathrm{Cr}(H_2)}}_{\text{prior odds}}$$

So the Bayes factor is exactly the multiplier: it says how far the evidence moves you, and your prior says where you started from. A Bayes factor of 1 means the evidence is equally expected either way, and moves you not at all.

Example 4. Suppose you consider it equally likely that a coin is biased so that it always lands heads (H_1) or is fair (H_2), and you flip it and get heads (E). The probability of its landing heads given that it was biased is 1, and given that it was fair is $^1/_2$, so the Bayes factor is $1/0.5 = 2$. The odds between biased and fair therefore move from $1:1$ to $2:1$ — in other words, a $^2/_3$ chance that the coin is biased and a $^1/_3$ chance that it is fair.

4 Dutch book arguments

4.1 What a Dutch book argument is

We shall now consider some of the best-known arguments for probabilism: that is, for the three Kolmogorov axioms we saw in section 2. These are known as *Dutch book arguments*.

The point of them is to show that agents whose credences do not satisfy these axioms are systematically exploitable: they can end up taking actions which together amount to choosing to give up money, or something else they value, for no gain whatsoever. And, so the argument goes, since agents who are systematically exploitable are irrational, rationality requires that they do not have the sorts of non-probabilistic credences which got them into trouble.

There is also a Dutch book argument for the principle of conditionalisation, though we will not cover that this week. (We may, time permitting, come back to it in week 5.)

4.2 The betting interpretation of credences

In order for any Dutch Book argument to work, we need some link between what an agent believes and what she does. This is because credences on their own are inert: without some principle linking them to an agent's behaviour, we cannot conclude anything about what an agent will do merely from their credences (degrees of belief).

The standard assumption made in these arguments about the link between credences and an agent's behaviour is that credences match up with the agent's betting behaviour.

Write an *$S bet on A* for a deal that pays $S if A is true and nothing otherwise.

Principle 3 (The betting interpretation of credences)**.** *An agent whose credence in A is p will buy an $S bet on A for any price below $pS, and will sell one for any price above $pS.*

So her credence fixes a price, and she is willing to take either side of it. We will assume this for the time being. It is a substantial assumption, and it is probably false.

4.3 The Dutch book theorem

Two of the axioms can be argued for with a *single* bet. These are the easy cases, because nothing in them depends on how a sequence of offers is handled.

Non-negativity. Suppose $\mathrm{Cr}(A) = -0.5$. By the betting interpretation she will sell a \$100 bet on A for any price above $-\$50$ — and a negative price means she pays. So she will happily pay someone \$49 to take the bet on, and then owes them another \$100 if A turns out true.

	A	$\neg A$
Her net gain	$-\$149$	$-\$49$

She loses \$49 whatever happens, from one transaction.

Normalisation. Suppose $\mathrm{Cr}(\Omega) = 0.9$. She will sell a \$100 bet on Ω for \$91. But Ω is guaranteed, so she collects \$91 and pays out \$100.

	Ω (certain)
Her net gain	$-\$9$

And if instead $\mathrm{Cr}(\Omega) = 1.1$, she buys the same bet for \$109 and collects \$100, losing \$9 again.

These cases are easy, because they can be done in a single bet. Finite additivity is not so easy, because it depends on how credences relate to other credences, and so multiple bets need to be taken.

4.4 Finite Additivity Dutch Book for Myopic Choice

Suppose P and Q are mutually exclusive, and that, as an example of a finite additivity violation, our agent has credence 0.2 in P, credence 0.2 in Q, but credence 0.8 in $P \cup Q$. Since P and Q are disjoint, finite additivity requires $\mathrm{Cr}(P \cup Q) = 0.2 + 0.2 = 0.4$.

She is offered three bets in turn, each for a moment and then withdrawn.

Bet 1 She sells a \$100 bet on P for \$21. Her price for that bet is \$20, so selling above it is a gain.

Bet 2 She sells a \$100 bet on Q for \$21. Likewise.

Bet 3 She buys a \$100 bet on $P \cup Q$ for \$79. Her price for that bet is \$80, so buying below it is a gain.

Each bet, taken on its own, is one she wants. So it looks as though she should take all three. But if she does, she is down \$37 however things turn out.

It is worth seeing that happen. Figure 14 plots her net position in the three states — P, Q, and neither — once all three bets have been taken.

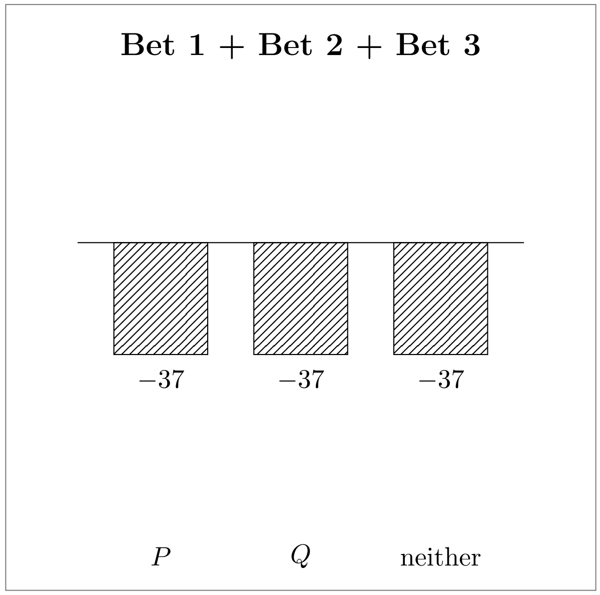

Figure 14: After all three she is down \$37 whatever happens. That is the Dutch book.

4.5 Sophisticated choice

So far we have quietly assumed that she treats each offer as if it were the only decision she will ever face. That is called *myopic* choice. An agent who instead looks ahead, and takes into account what she will do later, is called *sophisticated*. (Much more on this in week 5.)

The temporal structure implied by the argument for additivity above is given by the following decision tree.

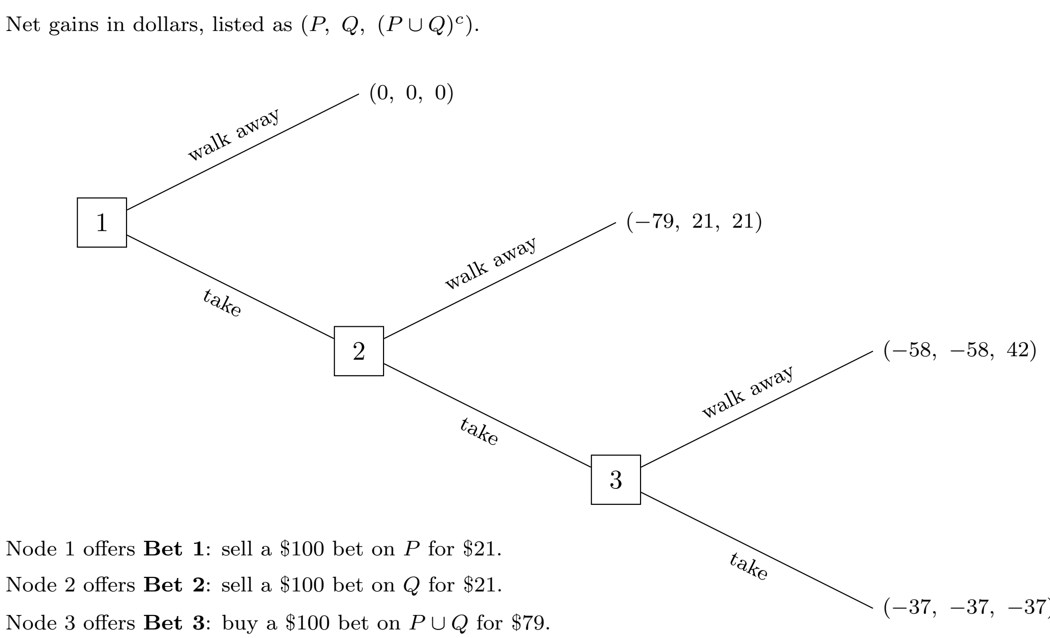

Figure 15: The Dutch book for finite additivity.

The squares are decision nodes — choice points for the agent — and the edges leaving them represent the particular choices open to her. She begins at decision node 1 and may, depending on what she chooses, reach decision node 3 and a final choice. The payoffs are written as three-place vectors: the first place is her gain or loss in dollars if P occurs, the second her gain or loss if Q occurs, and the third her gain or loss if neither of them does.

A sophisticated agent will obviously not take all three bets. For suppose she did. Then she would go on taking bets at node 2 as well, and so at node 1 her choice would really be between

$$(0,\ 0,\ 0) \qquad \text{and} \qquad (-37,\ -37,\ -37),$$

and she would obviously take the first. So she declines at node 1.

It turns out that a sophisticated agent here makes one trade and then walks away, though we will not show this now.

4.6 Can the book be repaired?

It can. The Dutch book for finite additivity can be reconstructed so that it works even against an agent who uses sophisticated choice — which is what she ought to be using. If there is time, we will see how in week 5.

4.7 From the theorem to an argument for probabilism

The Dutch book theorem is a piece of mathematics. To get an argument for probabilism out of it we need some further premises, and the obvious way of supplying them is not very good.

The flat-footed version runs: if your credences are non-probabilistic, a cunning Dutchman might come along and take your money; if they are probabilistic, he cannot; so you had better have probabilistic credences.

On this version of the argument, the reason why we should not have non-probabilistic credences is that avoiding them protects us from actually being exploited in a Dutch book situation. However, suppose that for some reason an agent knew for certain that no Dutch book situation would arise. In that case they would know that having non-probabilistic credences will not hurt them, and the reason to adopt probabilistic credences would vanish. But presumably a proponent of probabilism would not want to say that an agent like this is rational: although they will not be exploited, it is their pattern of credences itself which is supposed to be irrational.

Here is a better version:

(1) If an agent does not bet in line with the betting interpretation for the credences they have, then they are irrational.

(2) If an agent has non-probabilistic credences and bets in line with the betting interpretation, then they are vulnerable to exploitation by a Dutch book.

(3) If an agent is vulnerable to exploitation by a Dutch book, then they are irrational.

$\therefore$ If an agent has non-probabilistic credences, then they are irrational.

Suppose the agent has non-probabilistic credences. Either she bets in line with the betting interpretation or she does not. If she does not, then by (1) she is irrational. And if she does, then by (2) she is vulnerable to a Dutch book, and so by (3) she is irrational. Either way, she is irrational.

4.8 It's not about the money

One might object to Dutch Book arguments in general on the ground that it is not necessarily irrational to lose money. After all, it does not seem that we are rationally required to value money in particular, and if we do not value it, losing it cannot be a sign of irrationality.

Dutch Book arguments, however, do not need to be about money. Anything quantitative that can, in principle, be a prize in a bet, will do the job just fine. (The bets do not need to be practical or realistic, just possible in principle.) So long as there is any quantity the agent cares about like that, and they have non-probabilistic credences, and there is a suitable link to their betting behaviour (whether that comes via the betting interpretation or some weaker principle), the agent seems to be vulnerable to a Dutch Book, and thus open to the charge of irrationality.