Sets, Relations and Probability

Cambridge IA Formal Methods

Owen Griffiths & Alexander Roberts

Version 13.04.2020

Chapter 3 & Bibliography added (13.04.2020)
Introduction & Acknowledgements added (13.04.2020)
Solutions added (31.03.2020)
Slight rewording of exercise 2.5.1 (31.03.2020)
Footnote 1 added on p. 2 (31.03.2020)
Corrections made to exercises 1.6.4, 1.8.3 & 1.9.3 (03.03.2020)

Contents

Introduction

This textbook is for the IA Formal Methods lectures on Sets, Relations and Probability. It covers the elements of set theory (including relations and functions) and probability theory.

The book contains a number of exercises, which are optional although recommended. The exercises are designed for the reader to check that they understand the concepts and ideas which have been introduced in the relevant section. Please note that the exercises will not be discussed in the additional logic classes run by colleges, although the book does contain a solutions booklet.

If you spot any errors, typographical or otherwise, please send them to Alex Roberts at ajr207@cam.ac.uk.

Acknowledgements

This textbook draws on parts of Arif Ahmed's handouts for his Formal Methods lectures, which are available on his website. In particular, the chapter on probability draws substantially on Arif's last four handouts. Be warned that there are some small differences between Arif's handouts and this textbook. The most significant difference is that in this textbook relations are identified with their extensions for reasons of convenience.

Thanks are due to Annina Loets for proof-reading some exercises and to Oxana Zhigalova for bringing some typos to our attention.

1 Sets

We speak of collections all the time: this university consists of its colleges, faculties, etc; this textbook consists of its pages; the philosophy faculty library consists of its books. We speak of collections, aggregates, classes, *sets*. This part of the course will first be concerned with what these sets are. As well as being an important philosophical notion in itself, what purposes have sets been put to?

The first motivation is mathematical. At the end of the 17^{th} century, Newton and Leibniz discovered the calculus (analysis). This is the study of real-valued functions of a real variable. Real numbers are rather different to the other sorts of number – they are *infinite* objects. And, in its original form, the calculus made use of the notion of 'infinitesimal'. Generally, the notion of the infinite was becoming more central to mathematics. One crucial use of set theory is in taming the infinite.

Set theory has also come to serve a crucial, *foundational* role for mathematics. We would very much like our mathematical theories – of arithmetic, geometry, and so on – to be *consistent*. Virtually every mathematical theory of interest can be shown to be consistent *if* set theory is. So the consistency of set theory is of considerable philosophical interest.

There are other, more mundane, motivations. First, sets can provide a neat way of discussing semantics and metalogic. For example, what does the predicate 'is red' refer to? We might think *redness*, but what sort of object could that be? Rather, some have preferred to say that it refers to the *set* of red things. Similarly for metalogic: we might want to discuss the *soundness* and *completeness* of a proof system for a logic. How can we do this? One helpful way is to compare the *set* of logical truths to the *set* of theorems and ask if they are the same.

1.1 Collections

Intuitively, a set of things is just a collection of things considered as a single object, for example the set of coins in my pocket. However, unlike the coins in my pocket, the set of those coins in my pocket is not a physical object. For the purposes of this course, we shall treat it—and indeed all sets—as abstract objects.

You can think of a set as an abstract shopping bag into which you can put all sorts of things. These can be concrete things like Jack and Jill. Or they can be abstract entities like the natural numbers from 0 to 10. They can also be completely unrelated like the number 2, the concept of humility and the Queen's Diamond Jubilee.

We will use a specific notation for describing sets. If there are the individuals Jack and Jill, then the set containing exactly Jack and Jill is written as $\{\text{Jack}, \text{Jill}\}$. We shall also say that Jack and Jill are *members* of $\{\text{Jack}, \text{Jill}\}$.

In fact, we will use a special notation to express membership claims. The claim that Jack is a member of $\{\text{Jack}, \text{Jill}\}$ will be written as Jack $\in \{\text{Jack}, \text{Jill}\}$. More generally, when x is any individual and S is any set, the expression $x \in S$ is to be understood as the claim that x is a member of S. If x is any individual and y is *not* a set, then the claim $x \in y$ will be ill-formed.

The *order* of the objects in a set is insignificant. So, for example: $\{\text{Jack}, \text{Jill}\} = \{\text{Jill}, \text{Jack}\}$.

Sets can also be members of sets. So the set $\{\{\text{Jack}, \text{Jill}\}, \text{Jane}\}$ is a set. It has two members: the object Jane and the set $\{\text{Jack}, \text{Jill}\}$. The objects Jack and Jill are members of $\{\text{Jack}, \text{Jill}\}$: $\text{Jack} \in \{\text{Jack}, \text{Jill}\}$ and $\text{Jill} \in \{\text{Jack}, \text{Jill}\}$. But they are not members of $\{\{\text{Jack}, \text{Jill}\}, \text{Jane}\}$: $\text{Jack} \notin \{\{\text{Jack}, \text{Jill}\}, \text{Jane}\}$ and $\text{Jill} \notin \{\{\text{Jack}, \text{Jill}\}, \text{Jane}\}$.

1.2 Empty Set

I said that a set of things is just a collection of things, but there is one set which is a collection of no things whatsoever. We call this the *empty set* and denote as either $\{\}$ or $\emptyset$. (In some texts, you might see it referred to as the null set.)

It will be very mathematically convenient to have the empty set at our disposal, even if it is slightly at odds with our ordinary notion of a collection. For instance, I can't claim to have a collection of priceless artworks if I have none whatsoever.

Exercises for 1.1-1.2

1. Are the following claims true, false, or ill-formed?[1]

   a. $0 \in \{0, 1, 2\}$

   b. $0 \in \{1, 2\}$

   c. $0 \in \{\{0\}\}$

   d. $0 \in 0$

   e. $0 \in \emptyset$

2. Write down the following claims in our notation. Are the they true, false, or ill-formed?

   a. The number 2 is a member of the set containing exactly Jack and Jill.

   b. The empty set is a member of the number 2.

   c. The number 2 is a member of the set containing exactly the empty set and the number 2.

1.3 Extensionality

Above, I presumed that there is a unique set: *the* empty set. What justifies this presumption?

Our notion of a set is governed by an important principle: **the principle of extensionality**.

Extensionality. If a set S and a set S' have the exact same members, then they are identical.

This principle justifies my presumption that there is a unique empty set. I can use **Extensionality** to reason as follows. Suppose S is an empty set and S' is also an empty set. Then, clearly they have exactly the same members, namely none. Therefore, by **Extensionality** they are one and the same set. Thus, the empty set is unique.

[1] Note: in some textbooks, you might find that numbers like 2 are *defined* to be certain sets. That is not the approach taken in what follows.

I can also write a version of **Extensionality** in the language of first-order logic with identity which has been enriched with our predicate of set membership '$\in$' and the predicate 'Set', understood informally as the predicate 'is a set'.

Extensionality. $\forall x \forall y((\mathsf{Set}(x) \wedge \mathsf{Set}(y)) \rightarrow (\forall z(z \in x \leftrightarrow z \in y) \rightarrow x = y))$

Notice that, more generally, **Extensionality** gives us a means of proving that two sets are identical. We must establish two conditionals (when x and y are sets):

1. $\forall z(z \in x \rightarrow z \in y)$
2. $\forall z(z \in y \rightarrow z \in x)$

1. and 2. allow us to conclude that some object z is a member of x if, and only if, it is a member of y. By **Extensionality**, we can then conclude that $x = y$.

Exercises for 1.3 (* = greater difficulty)

1. Why is the set {The Prime Minister on 10 February 2020, Jeremy Corbyn} identical with the set {Boris Johnson, Jeremy Corbyn}? What does this tell us about how we describe sets?

2. Use **Extensionality** to argue that there is only one set which contains exactly the empty set.

3.* If we enrich the language of first-order logic (with identity) with the predicate '$\in$' and the name '$\emptyset$' for the empty set, then we can define the predicate 'Set'. How?

1.4 Notation for Describing Sets

It is helpful to introduce some more general notation for describing sets instead of just listing their members in between curly brackets. After all, the set of natural numbers (the numbers 0, 1, 2, ...) is infinite, so we can't simply list its members on our page.

Our notation will permit us to define a set in terms of a *condition* which all and only its members satisfy. For example, instead of listing all of the natural numbers, we shall define the set of natural numbers as *the set of all things which meet the condition of being a natural number*.

Notation. The expression '$\{x : x \text{ is a natural number}\}$' denotes the set of natural numbers.

Likewise, $\{x : x \text{ is red}\}$ is the set which contains all and only the red things. To be clear, the notation will often be a matter of convenience: $\{2\} = \{x : x \text{ is an even prime number }\}$. These are two ways of naming the very same set. Nothing hinges on which one we choose. Generally, we'll write $\{x : \phi(x)\}$ for *the* set of ϕs. Indeed, I can now redescribe some of the sets we considered previously. For example, {Jack, Jill} $= \{x : x = \text{Jack} \vee x = \text{Jill}\}$.

As an aside, when a set contains one member, like the set {Jack} does, we shall call it a *singleton*. Clearly, everything is the one and only member of its singleton. For example, the singleton of Donald Glover is {Donald Glover}. Importantly, Donald Glover $\neq$ {Donald Glover}. Why? By Leibniz's Law, since Donald Glover writes sitcoms and sets can't write sitcoms!

Sets also have singletons, so the singleton of {Donald Glover} is {{Donald Glover}}. Again, {Donald Glover} $\neq$ {{Donald Glover}}. Why? They don't have the same members: Donald Glover is a member of the former but not of the latter. By **Extensionality**, therefore, these are different sets. Again, nothing odd is going on here: continuing the shopping bag analogy, a singleton is a shopping bag with just one thing in it. We could then put this bag inside another, and we'd end up with a different bag, extensionally speaking.

Exercises for 1.4

1. Using our notation, write down the claim that the number zero is not a member of the set of all humans.

2. Using our notation, redescribe the set $\{\emptyset\}$.

3. Using our notation, redescribe $\emptyset$.

4. Write out in our formal language the claim that everything is the one and only member of its singleton.

1.5 Subsets and Supersets

Some sets bear interesting relations to another. For example, every member of {Jack} is also a member of {Jack, Jill}. When this occurs, we say that {Jack} is a *subset* of {Jack, Jill}.

More precisely, a set S is a *subset* of a set S' iff every member of S is a member of S'. When S is a subset of S', we write $S \subseteq S'$. We can also define this relation in our formal language.

Definition. (Subset) $S \subseteq S' := \forall x(x \in S \to x \in S')$

It is absolutely crucial that the subset relation is not confused with the membership relation. Membership is the relation that holds between an *object* (possibly a set) and a set. Subsethood is a relation that holds between two sets. It holds between sets S and S' just when every member of S is also a member of S'. For example, the set of footballers F is a subset of the set of sportspeople S, since every footballer is also a sportsperson: $F \subseteq S$. Notice also that every set is a subset of itself. However, when $S \subseteq S'$ and $S \neq S'$, S is a *proper subset* of S'.

Definition. (Proper Subset) $S \subset S' := S \subseteq S' \wedge S \neq S'$

For example, since there are also sportspeople who are not footballers, such as snooker players, the set of footballers F is a proper subset of the set of sportspeople S.

For another example, consider the set of all triangles A and the set of all trilateral shapes L. Here, $A \subseteq L$ but of course $L \subseteq A$. When this holds, we can conclude $A = L$ by **Extensionality**. This is of course what we'd expect. We've said that one test for whether two sets S and S' are identical is whether it is the case that every member of S is a member of S' and *vice versa*. We now know this equivalent to saying that $S \subseteq S'$ and $S' \subseteq S$.

One helpful fact to bear in mind is that $\emptyset$ is a subset of every set. This is because it is vacuously true that every member of $\emptyset$ is a member of any given set as $\emptyset$ has no members at all.

It is also helpful to have a term for the converse of the subset relation: the superset relation.

Definition. (Superset) $S \supseteq S' := S' \subseteq S$

In other words, when $S \supseteq S'$ every member of S' is a member of S.

Worked Example. Show that if $S \subset S'$ then it is not the case that $S' \subseteq S$.

Strategy. To prove this conditional, I will use a technique known as conditional proof. This is an analogue of the natural deduction introduction rule for the conditional, except in natural language. It allows me to suppose the antecedent of a conditional, prove its consequent as a consequence, and then conclude that the conditional is true.

Proof. First, let us *suppose* that $S \subset S'$. By the definition of proper subset, it follows that $S \subseteq S' \wedge S \neq S'$. In other words, every member of S belongs to S' but, since $S \neq S'$, there is some member of S' which does not belong to S. As a result, S' is not a subset of S. Therefore, by our informal rule of conditional proof, if $S \subset S'$ then it is not the case that $S' \subseteq S$. □

Exercises for 1.5

1. Define the proper superset relation.

2. Show that if $S_1 \subseteq S_2$ and $S_2 \subseteq S_3$, then $S_1 \subseteq S_3$.

3. Let $A = \{0, 1\}$ and $B = \{0, 1, 2, 3\}$. Are the following claims true, false, of ill-formed?
   a. $A \subseteq B$
   b. $A \subset B$
   c. $A \in B$
   d. $0 \in A$
   e. $0 \subseteq B$

1.6 Union and Intersection

There are various algebraic operations on sets that are crucial to understand.

Consider the sets {Jack} and {Jill}. One natural thought is that there should a set which results from pooling together exactly the members of each set. Indeed, there is: {Jack, Jill}. The set {Jack, Jill} is the *union* of the sets {Jack} and {Jill}, and we shall denote it as follows: {Jack} $\cup$ {Jill}. Venn diagrams provide a helpful way of thinking about this:

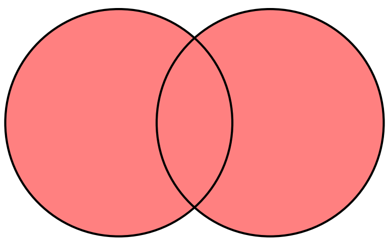

Definition. (Union) $S \cup S' := \{x : x \in S \vee x \in S'\}$

For another helpful example, consider the sets $\{0, 1, 2\}$ and $\{2, 3, 4\}$. $\{0, 1, 2\} \cup \{2, 3, 4\} = \{0, 1, 2, 2, 3, 4\} = \{0, 1, 2, 3, 4\}$. The definition of union delivers the first identity; **Extensionality** delivers the second. There is an obvious analogy between *union* and (inclusive) *disjunction*. To check that you understand the notion of union, it might help to convince yourself of the following fact: every set S is a subset of its union with any set S' (i.e. $S \cup S'$).

There is another algebraic operation on sets that is important to understand too. Consider the sets {Jack, Jill} and {Jill, Jane}. One natural thought is that there should a set which results from taking the members which are common to these sets. Indeed, there is: {Jill}. The set {Jill} is the *intersection* of the sets {Jack, Jill} and {Jill, Jane}, and we shall denote as follows {Jack, Jill} $\cap$ {Jill, Jane}. Again, Venn diagrams are helpful for visualising this operation.

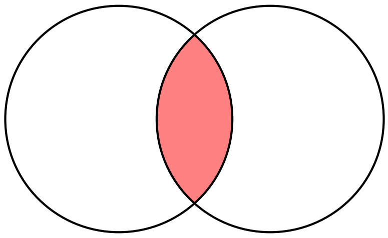

Definition. (Intersection) $S \cap S' := \{x : x \in S \wedge x \in S'\}$

For example, if A is the set of all animals and C is the set of all chimpanzees, the set $A \cap C$ is the set of all things that are both animals and chimpanzees, namely the set of all chimpanzees. There is an obvious analogy between intersection and *conjunction*. To check that you understand the notion of intersection, it might help to convince yourself of the following fact: the intersection of any sets S and S' is a subset of S (and of S').

Worked Example. Show that $((S \cap S') \cup S) = S$.

It's important that you set out your proofs of certain set-theoretic facts in an elegant manner. To demonstrate a properly presented proof, I'll set out a model answer for this example. I'll begin by sketching the idea of the proof to you. This will not be the official proof, but it'll contain all the core ideas needed to present the proof properly.

Sketch. x is a member of $((S \cap S') \cup S)$ iff either x is a member $(S \cap S')$ or x is a member of S. If the latter, then clearly x is a member of S. If the former, then x is a member of S and S': so it must be a member of S.

Proof.
$$\begin{aligned} x \in ((S \cap S') \cup S) &\leftrightarrow x \in (S \cap S') \vee x \in S && \text{Def. } \cup \\ &\leftrightarrow (x \in S \wedge x \in S') \vee x \in S && \text{Def. } \cap \\ &\leftrightarrow x \in S && \text{PL} \end{aligned}$$

The final step is justified by propositional logic (PL): clearly, the statement $((A \wedge B) \vee A) \leftrightarrow A$ is a tautology of propositional logic. □

When some sets have *no* members in common, we say that they are *disjoint*. In fact, we can define this notion in terms of intersection.

Definition. (Disjointness) S is disjoint from $S' := (S \cap S') = \emptyset$

For example, if A is the set of all animals and O is the set of all oranges, then A and O are disjoint, since $A \cap O = \emptyset$.

Exercises for 1.6

1. Show that $S \cup S$ is just S.

2. Show that $S \cap \emptyset$ is just $\emptyset$.

3. Show that if S and S' are disjoint and S' is non-empty, then $S \subset (S \cup S')$.

4. Show that Show that $(S_1 \cup (S_2 \cap S_3)) = ((S_1 \cup S_2) \cap (S_1 \cup S_3))$

1.7 Complement

Consider the sets $\{\text{Jack}, \text{Jill}\}$ and $\{\text{Jack}\}$. One natural thought is that there should be a set which results from removing the only member of $\{\text{Jack}\}$ from $\{\text{Jack}, \text{Jill}\}$. Indeed, there is: $\{\text{Jill}\}$. The set $\{\text{Jill}\}$ is the *complement* of $\{\text{Jack}\}$ in $\{\text{Jack}, \text{Jill}\}$.

Definition. (Complement) $S - S' := \{x : x \in S \wedge x \notin S'\}$

As before, Venn diagrams are helpful for visualising this operation.

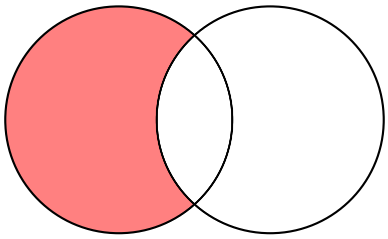

With the notions of union, intersection, and complement, one can now prove various identities between different sets.

Worked Example. Show that $((S \cap S') \cup (S - S')) = S$.

Proof.

$$\begin{aligned} x \in ((S \cap S') \cup (S - S')) &\leftrightarrow x \in (S \cap S') \vee x \in (S - S') && \text{Def. } \cup \\ &\leftrightarrow (x \in S \wedge x \in S') \vee x \in (S - S') && \text{Def. } \cap \\ &\leftrightarrow (x \in S \wedge x \in S') \vee (x \in S \wedge x \notin S') && \text{Def. } - \\ &\leftrightarrow x \in S && \text{PL} \end{aligned}$$

The last step is justified by the tautology of propositional logic $((A \wedge B) \vee (A \wedge \neg B)) \leftrightarrow A$. □

Exercises for 1.7

1. Show that $(S - (S \cap S')) = S - S'$.

2. Show that $((S \cup S') - (S \cap S')) = ((S - S') \cup (S' - S))$.

3. Show that $((S \cup S') - (S \cap S')) = (S \cup S')$ when S and S' are disjoint.

1.8 Power sets

We have defined the subset relation on sets. We've also seen that every set has at least one subset: itself. However, every non-empty set will have multiple distinct subsets. For example, any singleton set, such as {Jack}, has two subsets: {Jack} and $\emptyset$. One natural thought is that for any given set S, there should be a set which includes all and only the subsets of S. We shall call this the *power set* of S, and we shall denote it as $\mathcal{P}(S)$.

Definition. (Power set) $\mathcal{P}(S) := \{x : x \subseteq S\}$

For example, consider the set $\{a, b, c\}$. Its subsets are: $\emptyset, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}$, $\{a, b, c\}$. The set of all of these subsets, the power set of $\{a, b, c\}$, is therefore:

$$\mathcal{P}(\{a, b, c\}) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}, \{a, b, c\}\}.$$

Tip 1. Sometimes you might be asked to calculate the power set of a given set. To double-check you have calculated the power set properly, it is helpful to bear in mind that if a set has n members, then its power set has 2^n members. For example, the set $\{a, b, c\}$ has 3 members, so (as you can verify) $\mathcal{P}(\{a, b, c\})$ has $2^3 = 8$ members.

Tip 2. Since the empty set is a subset of every set, it is a *member* of every power set. Likewise, since every set is a subset of itself, the power set of a given set S always includes S itself as a member.

Worked Example. Show that if $\mathcal{P}(S) = \mathcal{P}(S')$ then $S = S'$.

Strategy. To prove this, I'll use a technique called *proof by contraposition*. We know that in propositional logic the conditional $p \to q$ is logically equivalent to $\neg q \to \neg p$. To introduce some terminology, the conditional $\neg q \to \neg p$ is the *contrapositive* of $p \to q$. Proof by contraposition is when a conditional is proved by proving its contrapositive.

Proof. In the example at hand, I am asked to prove a conditional: if $\mathcal{P}(S) = \mathcal{P}(S')$ then $S = S'$. I'll do so by just proving the contrapositive of this conditional: if $S \neq S'$, then $\mathcal{P}(S) \neq \mathcal{P}(S')$. So suppose that $S \neq S'$. By **Extensionality**, there is something x in one set but not the other. Thus, $\{x\}$ will be in the power set of one set and not in the power set of the other set. Hence, $\mathcal{P}(S) \neq \mathcal{P}(S')$.

Exercises for 1.8

1. Is any set its own power set?

2. Define $\mathcal{P}(S) \cup S$ using our notation for describing sets.

3. Show that $\mathcal{P}(S \cap (S - S)) \subseteq \mathcal{P}(S')$.

1.9 Pairs and Products

We've said that *order* doesn't matter to sets. Consider two-membered sets. We could call these *unordered* pairs, since the order of the members doesn't matter: $\{a, b\} = \{b, a\}$. Sometimes we do care about order, though. For example, if I tell you that Jane and Ian finished first and second in a race, respectively, I don't simply mean that the members of {Jane, Ian} finished first and second: I mean that Jane won and Ian was second.

To respect this, we introduce the idea of an *ordered* pair. The ordered pair $\langle a, b\rangle \neq \langle b, a\rangle$. For ordered pairs, $\langle a, b\rangle = \langle c, d\rangle$ if, and only if, $a = c$ and $b = d$. In fact, we shall just define ordered pairs in terms of sets.

Definition. (Ordered pair) $\langle a, b\rangle = \{\{a\}, \{a, b\}\}$

Notice that $\langle a, a\rangle \neq \{a, a\}$ since $\langle a, a\rangle = \{\{a\}, \{a, a\}\}$. For notational convenience, however, we shall stipulate that the symbol '$\langle a\rangle$' is just alternative notation for '$\{a\}$'.

Now that we have the definition of an ordered pair, we can of course go further. If we want to order three objects, we can use an *ordered triple*, $\langle x, y, z\rangle$. For four objects, we can use an *ordered quadruple*, $\langle w, x, y, z\rangle$. Generally, if we want to order n-many objects, we can use an *ordered n-tuple*, $\langle x_1, ..., x_n\rangle$. In fact, we shall just define ordered in terms of ordered pairs.

Definition. (Ordered n-tuple) $\langle x_1, ..., x_n\rangle = \langle\langle x_1, ..., x_{n-1}\rangle, x_n\rangle$

If S and S' are sets, then it is natural to think that there should be a set of all the ordered pairs whose first component is taken from S and whose second component is taken from S'. We shall call this the *Cartesian product* of S and S', and shall denote it as $S \times S'$.

Definition. (Cartesian Product) $S \times S' = \{\langle x, y\rangle : x \in S \land y \in S'\}$

Worked Example. Let $A = \{0, 1\}$ and $B = \{1, c, d\}$. Calculate their Cartesian product.

Answer. $A \times B = \{\langle 0, 1\rangle, \langle 0, c\rangle, \langle 0, d\rangle, \langle 1, 1\rangle, \langle 1, c\rangle, \langle 1, d\rangle\}$ □

We can also talk about the Cartesian product of a set and itself. In this case, we say that the Cartesian product of A and itself is A^2.

Exercises for 1.9

1. Let $A = \{a\}$, $B = \{b, c\}$, and $C = \{d\}$. Calculate the following.
   i. $A \times B$
   ii. $\mathcal{P}(A) \times B$
   iii. $\mathcal{P}(A \times C)$

2. Calculate $\emptyset^2$.

3. Define the set $(A \times B) \cap (A \times A)$. Assuming A and B are non-empty, under which conditions is $(A \times B) \cap (A \times A)$ non-empty?

1.10 Cardinality

Intuitively, we can understand the idea that a given set has a *size*. It is natural to measure the size of a set in terms of the number of members it has. To introduce some terminology, we shall call this the *cardinality* of a set. For example, if $S = \{a, b\}$ then the cardinality of S is 2, as it has two members: a and b. It will be helpful to have some notation for expressing the cardinality of a set, so when S is a set $|S|$ is the cardinality of S, which will always be a number.

Definition. (Cardinality) The cardinality of a set is the number of members it has.

Of course, there are different sets with one and the same cardinality. For example, the set $\{\text{Boris}, \text{Sajid}\}$ is different from the set $\{\text{Boris}, \text{Rishi}\}$, but the cardinality of both is 2. In terms of our notation: $|\{\text{Boris}, \text{Sajid}\}| = |\{\text{Boris}, \text{Rishi}\}| = 2$.

Since some of the sets we will consider in this course will contain numbers as members themselves, it's important not to get confused between membership and cardinality. For instance, consider the set which contains exactly the number 2, $\{2\}$. The cardinality of this set (i.e. $|\{2\}|$) is the number 1, since $\{2\}$ is a singleton set: it contains exactly one member. So, $|\{2\}| \notin \{2\}$.

One of the powerful features of contemporary set theory is that it permits not just finite sets, but also sets with infinite cardinality: sets which contain infinitely many members. In fact, we've already seen one example of an infinite set already: the set of natural numbers (recall that the natural numbers are the numbers $0, 1, 2, ...$). One of the fascinating facts that one can *prove* in contemporary set theory is that there are different infinite sets with different cardinalities: in intuitive terms, this may be understood as the claim that some infinite sets are bigger than other infinite sets.

Exercises for 1.10

1. Let $A = \{1, 2\}$. Calculate the following.

   i. $|A - \{2\}|$

   ii. $|A \times (A - \{2\})|$

   iii. $|\mathcal{P}(A)|$

2. When n is a non-zero natural number, show that $|\{1, ..., n\}| \in \{1, ..., n\}$. Deduce that there's a particular non-zero natural number such that the cardinality of its singleton set is a member of its singleton.

1.11 Comprehension & Russell's Paradox

So far, we've used lots of different sets in the examples and exercises. We've used the set of natural numbers, the set of the current Prime Minister and the Leader of the Opposition, and the set containing only the set containing Donald Glover. This prompts a natural question: for any things, is there a set containing exactly them?

We might try to make this question a little more precise: if we specify any condition whatsoever, is there a set containing exactly the things meeting that condition? For example, we have simply assumed there is a set of things meeting the condition *being a natural number*. But what about the condition *being a left shoe of Donald Trump or a right shoe of Bertrand Russell*, or any condition whatsoever? More generally, for any condition ϕ which we can express, is there a set $\{x : \phi(x)\}$?

Here, ϕ is also allowed to be a property of sets, such as *having 4 or more members*. The sorts of things with members are sets, so this set will include the set of members of the Beatles, the set of natural numbers, the set of pages in *Principia Mathematica*, the set of Oscar-winning films, and so on.

In our formal language, we can state this 'naïve' principle that for any condition we can express there is a set consisting of exactly those things which meet that condition.

Naïve Comprehension. $\exists y \forall x(x \in y \leftrightarrow \phi(x))$

Think about both directions of this principle: it tells us that there is a set y which contains only things that meet condition ϕ (the left-to-right direction), and contains every thing which meets condition ϕ (the right-to-left direction).

The reason why this principle is known as ***Naïve*** **Comprehension** is because the guiding idea, that for any condition there is a set of things meeting exactly that condition, is just that: *naïve*. As it turns out, the principle of **Naïve Comprehension** is *inconsistent*, as the following famous argument shows.

Argument. Consider the condition of *being non-self-membered*. By **Naïve Comprehension**, this defines the set of all sets that are not members of themselves. That is, it defines the set $S = \{x : x \notin x\}$. Now, is S a member of itself? Suppose that it is, i.e. $S \in S$. Then by the membership condition for S, S is not a member of itself, i.e. $S \notin S$. Now suppose that S is not a member of itself, i.e. $S \notin S$. If $S \notin S$ then S meets the condition for being a member of itself: *being non-self-membered*. So, $S \in S$. However, we've now argued using basic propositional logic and the existence of S that $S \in S \leftrightarrow S \notin S$, which is an obvious contradiction. So there is no set of all sets which are not members of themselves.

This argument is due to Bertrand Russell and it is known as **Russell's Paradox**. Here's a fictional analogy to help you understand the style of reasoning used in the argument:

> In a village, there is a barber who shaves all and only those people in the village who do not shave themselves. Does the barber shave themselves? Well, suppose that he does. Then, the barber must be a person who does not shave himself, since he shaves all and only those people in the village who do not shave themselves. So, suppose that the barber doesn't shave himself. Well, then the barber meets the condition of being exactly one of those people who the barber shaves! So, using the description of the fictional scenario, we've argued using basic propositional logic that the barber shaves himself if and only if he doesn't. Hence the fictional scenario we've described must be logically inconsistent.

1.12 Separation & The Universal Set

Given that **Naïve Comprehension** is inconsistent, we can't use it in our theory of sets. However, philosophers, logicians and mathematics have proposed a weaker principle for set formation which can be used to articulate a consistent theory of sets. We will not be making heavy use of this principle in this course, but understanding it will help explain one important fact about sets.

One natural way to weaken **Naïve Comprehension**, one might think, is to add in an extra condition for set formation. Instead of thinking, naïvely, that for any condition we can express there is a set containing exactly those things meeting that condition, a different thought is that for any condition we can express and *any set we already know to exist*, there is a *subset* of that set which contains exactly those things meeting that condition.

For example, since the set $S = \{\text{Boris}, \text{Rishi}\}$ exists, then there is the set of all and only the members of S which are currently Prime Minister, i.e. $\{\text{Boris}\}$. Intuitively, the thought is that given any set and any condition, we can *separate* the items of that set which meet that condition into a set themselves. For another example, take the set of natural numbers $\{0, 1, 2, 3, ...\}$ and the condition of being an odd number. This condition can be used to separate that set into the set of odd numbers $\{1, 3, ...\}$. In our formal language, we can state this principle of separation.

Separation. $\forall z \exists y \forall x (x \in y \leftrightarrow (x \in z \wedge \phi(x)))$

In contrast to **Naïve Comprehension**, **Separation** requires the things which meet condition ϕ to form the set y to *already belong* to a set z. Intuitively, then, the principle just separates the ϕs in z into a possibly new set y.

Here is the crucial fact about sets and **Separation** you need to know for this course: if there is a set of absolutely everything then **Separation** is equivalent to **Naïve Comprehension**. In other words, if there is a universal set, then **Separation** collapses into **Naïve Comprehension**. The reason for this straightforward: if there is a set of everything, then separating the ϕs of that set into a set will just amount to forming the set of all and only the ϕs.

Due to this crucial fact, in standard set theory **there is no universal set**. That is there is no set which contains absolutely everything there is. This will be important to keep in mind, as it will explain why the universal set is absent in what follows.

1.13 References & Further Reading

For further reading about Russell's Paradox and other exercises, see the first two chapters of Enderton, H. (1977) *Elements of Set Theory*. London: Elsevier. Note that some of the material in those chapters (and definitely the remainder of the book) goes beyond the level of this course. See also the first chapter of Steinhart, E. (2009) *More Precisely*. Ontario: Broadview.

2 Relations

Many philosophical debates concern particular relations between individuals, events, situations, or acts. For example, ethicists might debate whether one situation is *better than* another situation, morally speaking. Metaphysicians might debate whether one event *causally depends* on another. Similarly, epistemologists might debate whether one way of revising one's beliefs in reaction to new evidence is *more rational than* an alternative way of doing so.

When philosophers are interested in such delicate questions about relations between certain things, they usually identify certain structural features of the relations that are pertinent to the discussion. For example, consider the following structural features of some of the above relations.

1. If situation s_1 is better than situation s_2, and situation s_2 is better than situation s_3, then s_1 is better than s_3.
2. If effect e causally depends on cause c, then c does not causally depend on e.
3. No event causally depends on itself.

Very often in philosophy it is useful to get clear on which structural features of a relation are agreed on by all participants to the dispute, and which, if any, are being contested. In this section of the course, we will use our set theory to provide a framework for discussing relations and properties of relations more rigorously.

2.1 Properties and Extensions

You are familiar with the concept of a *predicate* from first-order logic.

'_____ is a dog'

'_____ is a philosopher'

One-place predicates like these express *properties*. For instance, '_____ is a dog' expresses the property of *being a dog*. Crucially, properties have *extensions*: the set of all and only the things which possess that property. For example, the extension of *being a dog* is the set of all and only the dogs.

Sometimes, there might even be different properties which are *coextensive*: they share the same extension. For example, the properties of *being a naturally featherless biped* and *being human* have the exact same extension: they apply to precisely the same objects. But it is natural to think that they are different properties: after all, surely it is possible for there to be a naturally featherless biped which is not human.

Nevertheless, in what follows **we shall identify properties with their extensions**. Thus, since the set $\{x : x \text{ is a dog}\}$ is the extension of the property *being a dog*, we shall identify *being a dog* with the set $\{x : x \text{ is a dog}\}$. As a consequence, in our system two properties are identical when they are coextensive. For example, the properties of *having three angles* and of *having three sides* will be treated as identical: since $\{x : x \text{ has three angles}\} = \{x : x \text{ has three sides}\}$.

There are very deep philosophical questions about how to individuate properties, if not by the criterion of extensionality. However, it will be mathematically very convenient (and standard) to work under the idealisation that properties are extensional.

Recall that predicates can have more than one gap, like in the following examples:

> '$\_\_\_\_\__1$ loves $\_\_\_\_\__2$'
>
> '$\_\_\_\_\__1$ is taller than $\_\_\_\_\__2$'
>
> '$\_\_\_\_\__1$ is between $\_\_\_\_\__2$ and $\_\_\_\_\__3$'
>
> 'the distance from $\_\_\_\_\__1$ to $\_\_\_\_\__2$ equals the distance from $\_\_\_\_\__3$ to $\_\_\_\_\__4$'

These polyadic predicates express *relations*. For example, '$\_\_\_\_\__1$ loves $\_\_\_\_\__2$' expresses the relation of *loving* and '$\_\_\_\_\__1$ is between $\_\_\_\_\__2$ and $\_\_\_\_\__3$' expresses the relation of *betweenness*.

Just as properties have extensions, so do relations. However, the extension of a relation can't be a set of objects because relations don't apply to objects individually. Instead, two-place relations like *loving* apply to *pairs* of objects. For instance, the *loving* relation applies to the pair of objects, Romeo and Juliet. One might thus think that the extension of a relation like *loving* is just a set of pairs, that is of two-membered sets of objects.

Unfortunately, however, there is the possibility of unrequited love. It might be that Jack loves Jill, but Jill does not love Jack back in return. But if we simply treated the extension of *loving* as a set which includes the pair $\{\text{Jill}, \text{Jack}\}$ it would also include the pair $\{\text{Jack}, \text{Jill}\} = \{\text{Jill}, \text{Jack}\}$. So the extension of *loving* should not consist of merely unordered pairs.

For another example, we want the relation *shorter than* to be distinct from the relation *taller than*, but the same pairs of objects would fall under both according to the suggested approach. For example, Kylie Minogue is shorter than Nick Cave so the set {Minogue, Cave} would fall under the first. But the set {Cave, Minogue} would fall under the second. The problem, of course, is that {Minogue, Cave} = {Cave, Minogue}.

The solution is to use *ordered pairs* as the members of extensions of relations like *loving* and *shorter than*. Remember that we denote these special types of pairs using angular brackets. So, for example, the ordered pair of Minogue and Cave will be written as follows ⟨Minogue, Cave⟩. Remember also that the main principle governing ordered pairs is the following.

Order Principle. If $x \neq y$, then ⟨x, y⟩ ≠ ⟨y, x⟩

As a consequence of this principle, ⟨Minogue, Cave⟩ ≠ ⟨Cave, Minogue⟩. So, we can take the extensions of relations like *loving* and *shorter than* to be sets of ordered pairs. In particular, the extension of *shorter than* is $\{\langle x, y\rangle : x \text{ is shorter than } y\}$. Likewise, the extension of *loving* is $\{\langle x, y\rangle : x \text{ loves } y\}$. Like with properties, **we shall identify relations with their extensions**.

Just as properties like *unicorn* can have empty extensions, so can relations, like *shorter and taller than*. We call these *empty relations*. In fact, there is only one of them, by our assumption of extensionality.

The relations of *loving* and *shorter than* are *binary* relations: they apply to pairs of objects. More carefully, we define the notion of a binary relation as follows.

Definition. (Binary relation) A set is a *binary relation* iff it contains only ordered pairs.

Tip. Be clear on the fact that a binary relation is a *set* of ordered pairs, and *not* itself an ordered pair.

In addition to binary relations, there are also *ternary* relations, like *betweenness*, and also *quarternary* relations like the notion of *equidistance* used in the example above. More generally, for any finite n (where $n \geq 2$) there are n-ary relations: relations which apply to sequences of n objects. Remember that we call such sequences n*-tuples*. For instance, the sequence $\langle 1, 2, 3 \rangle$ is a 3-tuple, and the sequence $\langle \text{Jack}, \text{Jill} \rangle$ is a 2-tuple whilst also being an ordered pair. Moreover, just like ordered pairs, the order of n-tuples matter. So, for example, the ordered 3-tuple $\langle 1, 2, 3 \rangle$ is different from the ordered 3-tuple $\langle 2, 1, 3 \rangle$.

Definition. (n-ary relation) A set is an n*-ary relation* (for $n \geq 2$) iff it contains only ordered n-tuples.

Notice that according to this definition the empty relation, whose extension is just the empty set, is an n-ary relation for any $n \geq 2$: it is vacuously true that everything which the empty set contains is an ordered n-tuple, since it doesn't contain anything at all.

Notation. Since we are identifying properties and relations with their extensions, we can introduce some useful notation. For instance, when L is the loving relation, we may write $\langle \text{Romeo}, \text{Juliet} \rangle \in L$ for the claim that Romeo loves Juliet.

Exercises for 2.1

1. If $x, y \in S$, then specify a set S' in terms of S such that $\langle x, y \rangle \in S'$.

2.2 Properties of Binary Relations on Sets

We now turn to defining structural properties of relations. We shall do so in a two-step manner. First, we shall define what it is for a relation to have a certain structural property *on a set of objects*. Then, we shall define what it is for a relation to have a structural property *simpliciter*.

First of all, we define what it is for a relation to be *reflexive* on a set. Intuitively, this occurs when every member of that set bears the relation to itself. For example, the relation *is the same human as* is reflexive on the set of all humans since every human is the same human as themselves. Likewise, the relation x *is the same height as* y is reflexive on the set of living people, for clearly every living person is the same height as themselves.

Definition. (Reflexivity on S) A binary relation R is *reflexive* on a set S iff for every $x \in S$, $\langle x, x \rangle \in R$

Notice immediately that this definition makes use of our convention of writing $\langle x, x \rangle \in R$ for the claim that x bears R to x.

In our formal language of set theory, I can write the statement that a relation R is reflexive on a given set S as follows: $\forall x(x \in S \rightarrow Rxx)$.

Next, we define what it is for a relation to be *irreflexive* on a set. Intuitively, this occurs when every member of that set does *not* bear the relation to itself. For example, the relation *is a different human to* is irreflexive on the set of all humans since every human is not distinct from themselves. Likewise, the relation x *is taller than* y is irreflexive on the set of living people, for clearly no living person is taller than themselves. Finally, in one of our motivating examples, we saw that some philosophers think that *causal dependence* is irreflexive on the set of events, since no event causally depends on itself.

Definition. (Irreflexivity on S) A binary relation R is *irreflexive* on a set S iff for every $x \in S$, $\langle x, x\rangle \notin R$

In our formal language of set theory, I can write the statement that a relation R is irreflexive on a given set S as follows: $\forall x(x \in S \rightarrow \neg Rxx)$. This should allow you to appreciate that a relation's being irreflexive on a given set is not always the same as its not being reflexive on that set. A relation can be not reflexive when just one element of the set does not bear that relation to itself, yet that need not require all elements of the set not to bear the relation to themselves.

Can a relation be *both* reflexive and irreflexive on a given set? Only if the set is empty: then it will be vacuously true that the relation is both reflexive and irreflexive on that set.

The next property of relations we will consider is *symmetry*. The relation *is married to* is symmetric on the set of all people since, if someone x is married to someone y, then y is also married to x. Likewise, the relation *is the same human as* is symmetric on the set of all humans since if a human x is the same human as y then of course y is the same human as x.

Definition. (Symmetry on S) A binary relation R is *symmetric* on a set S iff for every $x, y \in S$, if $\langle x, y\rangle \in R$ then $\langle y, x\rangle \in R$.

In our formal language of set theory, I can write the statement that a relation R is symmetric on a given set S as follows: $\forall x \forall y((x \in S \wedge y \in S) \rightarrow (Rxy \rightarrow Ryx))$.

In one of our motivating examples, we saw that some philosophers think that if effect e causally depends on cause c, then c does not causally depend on e. Similarly, if one situation is better, morally speaking, than another situation, then the latter cannot be better than the former. Such cases are examples of an *asymmetric* relation on a given set.

Definition. (Asymmetry on S) A binary relation R is *asymmetric* on a set S iff for every $x, y \in S$, if $\langle x, y\rangle \in R$ then $\langle y, x\rangle \notin R$.

In our formal language of set theory, I can write the statement that a relation R is asymmetric on a given set S as follows: $\forall x \forall y((x \in S \wedge y \in S) \rightarrow (Rxy \rightarrow \neg Ryx))$. Like reflexivity and irreflexivity, it is important to appreciate that a relation can be not symmetric on a given set without being asymmetric on that set. For instance, if there is just one case of unrequited love amongst humans then *loving* is not symmetric on the set of humans. Yet if all other cases of love are reciprocal then *loving* is not asymmetric on the set of humans either.

There is also a related property of a relation being *anti-symmetric* on a given set. Intuitively, a relation is anti-symmetric on a given set when different elements of the set do not bear the relation to one another (even though each element might bear the relation to itself). For example, *greater than or equal to* is anti-symmetric on the set of natural numbers, since no two different natural numbers are greater than or equal to one another. But of course each natural number is greater than or equal to itself.

Definition. (Anti-symmetry on S) A binary relation R is *anti-symmetric* on a set S iff for every $x, y \in S$, if $x \neq y$, then if $\langle x, y\rangle \in R$ then $\langle y, x\rangle \notin R$.

In our formal language of set theory, I can write the statement that a relation R is anti-symmetric on a given set S as follows: $\forall x \forall y((x \in S \wedge y \in S) \rightarrow (x \neq y \rightarrow (Rxy \rightarrow \neg Ryx)))$. It might also help to think about this condition in contrapositive terms. If R is anti-symmetric on S then when elements of S bear R to one another they are identical. In our formal language: $\forall x \forall y((x \in S \wedge y \in S) \rightarrow (Rxy \rightarrow Ryx) \rightarrow x = y))$.

The above considerations bring out an important point. Since symmetry, asymmetry and anti-symmetry are all conditional, the empty relation will have all three properties vacuously on any set. Note, however, that the empty relation isn't vacuously *reflexive* on just any set. (Although it is vacuously reflexive on the empty set, like all relations are.) Reflexivity on a set demands that every object-object pair $\langle x, x\rangle$ is in the extension of the given relation when x is an element of the set.

Worked Example. Show that any binary relation which is asymmetric on a given set is irreflexive on that set.

Answer. Let R be a binary relation which is asymmetric on a given set S. For every $x, y \in S$, if $\langle x, y\rangle \in R$ then $\langle y, x\rangle \notin R$. Thus, it cannot be the case that there is some $z \in S$ such that $\langle z, z\rangle \in R$, for were that the case then $\langle z, z\rangle \notin R$ by the asymmetry of R on S. So R is irreflexive on S too. □

Next, we define the property of *transitivity* for a given relation on a set. Intuitively, a relation is transitive on a set when it 'holds across chains'. For example, the relation *greater than* is transitive on the set of natural numbers because when n is greater than m and m is greater than k, n is also greater than k. Likewise, one of our motivating examples involved transitivity: if situation s_1 is better than situation s_2, and situation s_2 is better than situation s_3, then s_1 is better than s_3.

Definition. (Transitivity on S) A binary relation R is *transitive* on a set S iff for every $x, y, z \in S$, if $\langle x, y\rangle \in R$ and $\langle y, z\rangle \in R$, then $\langle x, z\rangle \in R$.

In our formal language of set theory, I can write the statement that a relation R is transitive on a given set S as follows: $\forall x \forall y \forall z((x \in S \wedge y \in S \wedge z \in S) \rightarrow ((Rxy \wedge Ryz) \rightarrow Rxz))$.

Be careful not to be deceived by the metaphor of transitivity 'holding across chains'. By our definition, *being the same human as* is transitive on the set of all humans since if human x is human y and human y is human z then $x = z$. However, in this 'chain' there is no 'intermediary link' between x and z: x just is z.

For a philosophical example of transitivity on a set, remember Derek Parfit's distinction between psychological *continuity* and *connectedness* in 'Personal Identity' (1971). Let the set again be all people. Psychological connectedness is the holding of a range of psychological states such as memories in common. It is not transitive on the set of humans, since this close psychological connection fails to hold down chains: I will have quite different psychological states to myself as a child. Psychological continuity, however, is transitive on the set of humans since it is the overlapping chain of psychological connectedness. I *am* psychologically continuous with myself as a child: hence, in Parfit's language, that child is *surviving* as me.

Tip. In questions on relations, you might be asked to provide examples of relations that have some of the above properties on a given set. You should get in the habit of using sets of natural numbers and relations on those sets in your answers. For example, suppose you are asked to provide an example of a transitive relation on a given set. Then your answer might be:

$S = \{x : x \text{ is a natural number}\}$

$R = \{\langle x, y\rangle : x \text{ is less than } y\}$

Do not assume that the person who will be reading your answer to the question will have any specialist knowledge about, for example, sets of football teams or sets of Love Island contestants and relations amongst them.

We have introduced the following properties of relations on a set: reflexivity, symmetry and transitivity. When a relation has all three of these properties on a set we call it an *equivalence relation* on that set.

Definition. (Equivalence relation on S) A binary relation R is an *equivalence relation* on a set S iff it is reflexive, symmetric and transitive on S.

Consider the relation *x is exactly as tall as y* on the set of living people. It is clearly reflexive on that set, since everyone is exactly as tall as themselves. It is also symmetric on that set: if x is exactly as tall as y, then y is exactly as tall as x. And it is transitive on that set: if x is exactly as tall as y, and y as z, then x is exactly as tall as z. Hence *x is exactly as tall as y* is an equivalence relation the set of living people. Other examples on the same set are:

- *x was born in the same town as y*
- *x is exactly as old as y*
- x *is the same person as* y

In contrast, *x is married to y* is not an equivalence relation on that set. This is because it fails to be reflexive on that set, since no one is married to themselves. The relation *x is a brother of y* on the same set is also not an equivalence relation, since it fails to be symmetric on that set: x may be a brother of y and y be a *sister* of x. Finally, *x was born or died in the same town as y* on the set of all people ever is not an equivalence relation on that set because it fails to be transitive on that set: x and y may have only their place of birth in common, y and z only their place of death in common, and x and z neither place in common.

Equivalence relations on a set have the interesting property of *partitioning* a set. A partition is a way of dividing up the members of a set into sectors such that every member belongs to exactly one sector. The division is *exclusive* (nothing belongs to more than one sector) and *exhaustive* (every member belongs to at least one sector).

As with many of the notions under discussion, set theory provides a helpful way of defining partitions.

Definition. (Partition on S) A partition P on S is a set of subsets of S such that every member of S belongs to exactly one member of P.

As before, we can write this in our formal language for set theory:

P is a *partition* on S iff $\forall x(x \in S \to \exists y(y \in P \wedge x \in y \wedge \forall z((x \in z \wedge z \in P) \to y = z)))$

The sectors into which a partition carves a set are *equivalence classes* of the relation that does the partitioning.

Consider the relation *x was born in the same country as y* on the set of living people. This is an equivalence relation on that set because it is reflexive, symmetric and transitive on the set. It therefore partitions the set into equivalence classes. In this case, one class will include all of the people born in the UK, one all the people born in France, and so on for every country.

We have said that *x is exactly as tall as y* is an equivalence relation on the set of living people. This relation partitions the set into equivalence classes of people of the same height. There is a class of people who are exactly 5 feet tall, one of people 6 feet tall, and so on. For another example, consider the relation of *being the same person as*. We have said that this is an equivalence relation on the set of of living people so it must partition the members of that set into equivalence classes. What are these classes? Each contains only one member, since no person is identical with anything but themselves.

Finally, I will define an important type of binary relation, which is used throughout philosophy, logic and mathematics. The intuitive idea is that binary relation is a *function* on a given set when it produces a unique output for a given input. For example, the relation of *having as one's first child* is a function on the set of all people who ever existed because if y is the first born child of x then no other individual z is the first born child of x too. (Note that this doesn't require every person who ever existed to have a first born child.)

Definition. (Function on S) A binary relation is a *function* on a set S iff for every $x, y, z \in S$, if $\langle x, y \rangle \in R$ and $\langle x, z \rangle \in R$, then $y = z$.

In our formal language, I can write the statement that a relation R is a function on a given set S as follows: $\forall x \forall y \forall z((x \in S \wedge y \in S \wedge z \in S) \to ((Rxy \wedge Rxz) \to y = z))$.

For another example of a function on a given set, it might help to appreciate that *being the same human as* on the set of all humans is a function, since each human is identical with only one human: themselves.

2.3 Properties of Relations: Reference List

A binary relation R is:

reflexive on a set S iff for every $x \in S$, $\langle x, x \rangle \in R$.

irreflexive on a set S iff for every $x \in S$, $\langle x, x \rangle \notin R$.

symmetric on a set S iff for every $x, y \in S$, if $\langle x, y \rangle \in R$ then $\langle y, x \rangle \in R$.

asymmetric on a set S iff for every $x, y \in S$, if $\langle x, y \rangle \in R$ then $\langle y, x \rangle \notin R$.

anti-symmetric on a set S iff for every $x, y \in S$, if $x \neq y$, then if $\langle x, y \rangle \in R$ then $\langle y, x \rangle \notin R$.

transitive on a set S iff for every $x, y, z \in S$, if $\langle x, y \rangle \in R$ and $\langle y, z \rangle \in R$, then $\langle x, z \rangle \in R$.

an *equivalence relation* on a set S iff it is reflexive, symmetric and transitive on S.

a *function* on a set S iff for every $x, y, z \in S$, if $\langle x, y \rangle \in R$ and $\langle x, z \rangle \in R$, then $y = z$.

Exercises for 2.2-2.3

1. Consider any set S and any binary relation which is not reflexive on S. Specify a condition on S which guarantees that the relation is also irreflexive on S.

2. Let $S \neq \emptyset$. Show that $\emptyset$ is not reflexive on S.

3. Let $S = \{1, 2, 3\}$ and let $R = \{\langle 1, 2 \rangle, \langle 2, 3 \rangle, \langle 3, 2 \rangle\}$.

   a. Is R reflexive on S?

   b. Is R symmetric on S?

   c. Is R anti-symmetric on S?

   d. Define a relation R' such that $R \subseteq R'$ and R' is transitive on S. Must R' be an equivalence relation on S?

   e. Is $R - \{\langle 3, 2 \rangle\}$ anti-symmetric on S?

4. Show that every relation which is asymmetric on a set is also anti-symmetric on that set.

5. A relation is *anti-transitive* on a set S iff for every $x, y, z \in S$, if $\langle x, y \rangle \in R$ and $\langle y, z \rangle \in R$, then $\langle x, z \rangle \notin R$.

   a. Write out the definition of anti-transitivity on a set in our formal language of set theory.

   b. Can a relation be both reflexive and anti-transitive on a single non-empty set? Substantiate your answer.

   c. Provide two examples of a pair of a relation and a set such that the relation is both transitive and anti-transitive on a that set.

6. Let S be the set of natural numbers. Let $\langle x, y \rangle \in R$ iff $x + 1 = y$. Show that R is a function on S.

2.4 Properties of Binary Relations

It is straightforward to see that a relation might be symmetric on a certain set but not symmetric on another set. For example, consider the following relation and sets.

$S = \{0, 1\}$

$S' = \{0, 1, 2\}$

$R = \{\langle 0, 1\rangle, \langle 1, 0\rangle, \langle 1, 2\rangle\}$

The relation R is symmetric on S, since for all $x, y \in S$ if $\langle x, y\rangle \in R$ then $\langle y, x\rangle \in R$. However, $2 \in S'$ and $\langle 2, 1\rangle \notin R$, so R is not symmetric on S'.

This raises a question: is there a notion of being symmetric *simpliciter* which binary relations might possess? It turns out that there is a very simple way to define symmetry *simpliciter*: being symmetric on *all* sets. In fact, we can define other absolute properties of binary relations by similar means.

A binary relation R is:

symmetric iff it is symmetric on every set.

asymmetric iff it is asymmetric on every set.

anti-symmetric iff it is anti-symmetric on every set.

transitive iff it is transitive on every set.

an *equivalence relation* iff it is an equivalence relation on every set.

a *function* iff it is a function on every set.

There is one glaring omission from this list, namely the property of being reflexive. However, there is an important reason for this omission. If a binary relation is reflexive on every set, then its extension must consist of all ordered pairs of the form $\langle x, x\rangle$ for every individual x there is. Yet this is tantamount to claiming that there is a universal set, which we have banned.[2] So we will not define a property of binary relations, being reflexive *simpliciter*. Likewise, in our system there is not a relation of identity *simpliciter*. This is slightly unintuitive, but Naïve Comprehension was intuitive and it generated a logical disaster.

Tip. You should observe that I have not defined a notion of 'being a binary relation on a set'. I have just defined the notion of a binary relation (a set of only ordered pairs), and various properties which a binary relation can have on a set.

2.5 Properties of Functions

It will help to introduce some properties of functions themselves. Before we do so, however, we first require some terminology for talking about functions.

[2] If you come to study some more set theory, you will see that it is very simple to prove that there is a universal set if there is a set of all ordered pairs of the form $\langle x, x\rangle$ for every individual x.

Intuitively, a function is a relation 'from' given inputs, which all belong to certain set, 'to' various respective outputs, which all belong to a set too. We can articulate these intuitive notions more carefully.

Definition. (Domain) If F is a function, then the *domain of* F is the set of all x such that there is some y for which $\langle x, y\rangle \in F$. We denote this set $dom(F)$.

Definition. (Range) If F is a function, then the *range of* F is the set of all y such that there is some x for which $\langle x, y\rangle \in F$. We denote this set $ran(F)$.

Notation. If F is a function and x is in its domain, then $F(x)$ denotes the unique thing y such that $\langle x, y\rangle \in F$. When $F(x) = y$, we say that F *maps* x *to* y.

For example, *the number of coins in the jacket pocket of* is a function on the set of all humans. Its domain is the set of all people whose jacket has a pocket, and its range is some subset of the natural numbers. For instance, lets suppose that Jeff Bezos is the person with the most coins in his jacket pocket, namely 1,000. Then the range of the function is the set of natural numbers from 0 to 1,000.

Definition. (From/Into) F is a function *from* S *into* S' iff F is a function whose domain is S and whose range is a subset of S'.

Notation. We may abbreviate the claim that F is a function from S into S' with the following: $F : S \to S'$.

It is important to notice that when $F : S \to S'$ the range of F need only be a *subset* of S'. The reason for this choice is that it will allow us to define certain properties of functions in an elegant way, as you're about to see.

Functions themselves can possess interesting properties. For example, a function F from S into S' might have the special property of missing out no member of S'. The following diagram will help to visualise this property.

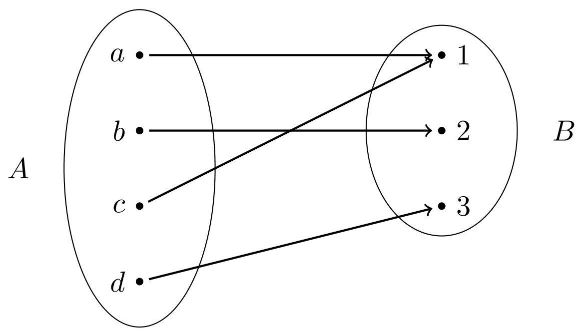

Definition. (Surjection) A function F is a *surjection* from S into S' iff for all $y \in S'$ there is some $x \in s$ such that $F(x) = y$.

Notation. If F is a surjection from S into S' we say that F is a function from S *onto* S'.

Another special property which a function can have is mapping no two different members of its domain to one and the same element of its range. Again, a diagram will help to visualise this.

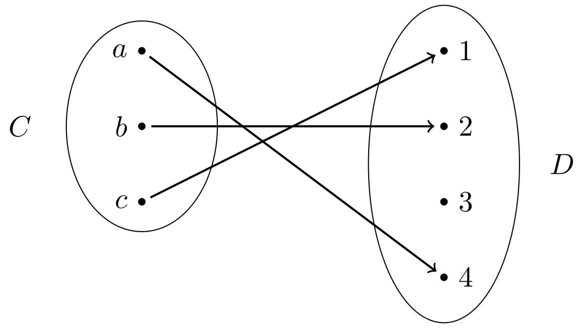

Definition. (Injection) A function F is an *injection* from S into S' iff if $F(x) = F(y)$ then $x = y$.

Now, observe that in the first diagram the function is not an injection into B. Moreover, in the second diagram the function is not a surjection into D. However, there are functions which are both injections and surjections into a given set. In fact, when a function from S into S' has this property, it puts the sets S and S' into a *one-one correspondence*. A final diagram will help to visualise when this occurs

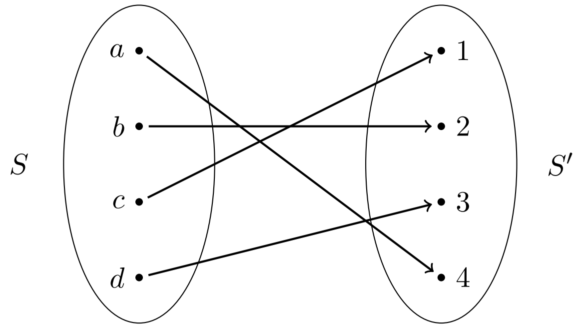

Definition. (One-One Correspondence) A function F is a *one-one correspondence* between S and S' iff it is a surjection and an injection into S'.

To check that you are following the definitions here, it might be helpful to think through why the term 'one-one correspondence' is used. The thought is that when F is a surjection and an injection from S into S', it is a unique pairing between members of S and S'. First of all, F is a surjection into S' so for every $y \in S'$ there must be some $x \in S$ for which $F(x) = y$. But F is also an injection into S', so x must be the only member of S which F maps to y. Thus, F pairs every $y \in S'$ with a unique $x \in S$.

Notation. Sometimes, you might see a one-one correspondence between certain sets referred to as 'bijection' between those sets. However, I will avoid this term in order to reduce the amount of terminology.

It is important to know that one-one correspondences between sets are connected intimately with the notion of cardinality. When sets stand in a one-one correspondence to one another, they have exactly as many members as one another, since those members are uniquely paired up by the correspondence. Thus, one way of establishing that a pair of sets have the same cardinality is to specify a one-one correspondence between them. For example, the set $\{0, 1, 2\}$ can be put into one-one correspondence with the set $\{3, 4, 5\}$ as follows:

$$\begin{array}{ccc} 0 & 1 & 2 \\ \downarrow & \downarrow & \downarrow \\ 3 & 4 & 5 \end{array}$$

In this diagram, the arrow represents the function $\{\langle 0,3\rangle, \langle 1,4\rangle, \langle 2,5\rangle\}$. This function is a unique pairing between the members of the sets $\{0,1,2\}$ and $\{3,4,5\}$, so they have the same cardinality.

There are other more interesting examples of one-one correspondences between pairs of sets. For instance, I can argue that the set of even natural numbers has the same cardinality as the set of odd natural numbers by specifying a one-one correspondence between them as follows:

$$\begin{array}{ccccc} 0 & 2 & 4 & 6 & \cdots \\ \downarrow & \downarrow & \downarrow & \downarrow & \\ 1 & 3 & 5 & 7 & \cdots \end{array}$$

In this diagram, the arrow represents the function F from the set of even natural numbers into the set of odd natural numbers such that $F(x) = x+1$. F is a one-one correspondence between these sets because for every odd natural number n there is some even number which F uniquely maps to n, namely $n-1$.

Once you understand this connection between cardinality and one-one correspondences, you are in a position to appreciate an odd fact about infinite sets: infinite sets stand in one-one correspondences with some of their proper subsets. For example, one can establish that the set of natural numbers stands in a one-one correspondence with the set of *even* natural numbers, which is a proper subset of the set of natural numbers. To see this, observe the following one-one correspondence between the natural numbers and the even natural numbers.

$$\begin{array}{ccccc} 0 & 1 & 2 & 3 & \cdots \\ \downarrow & \downarrow & \downarrow & \downarrow & \\ 0 & 2 & 4 & 6 & \cdots \end{array}$$

In this diagram, the arrow represents the one-one correspondence F from the set of natural numbers into the set of even natural numbers such that $F(x) = 2x$. (You are asked to verify that this is a one-one correspondence between these sets in exercise 4.) Thus, the set of natural numbers stands in a one-one correspondence to one of its proper subsets. As a result, the set of natural numbers has the same cardinality as one of its proper subsets! This might take some getting used to, but it is one of the many odd facts about infinite sets. Indeed, the German mathematician Richard Dedekind proposed that a set is infinite exactly when it stands in a one-one correspondence to one of its proper subsets.

Exercises for 2.5

1. In our formal language, write '$Func^S(R)$' to abbreviate the formula that R is a function on S, i.e. $\forall x \forall y \forall z((x \in S \wedge y \in S \wedge z \in S) \rightarrow ((Rxy \wedge Rxz) \rightarrow y = z))$. Using this abbreviation write out the following claims in the language:

a. R is function on S which is a surjection into S'.

b. R is function on S which is an injection into S'.

2. Let $F : S \to S'$ be such that S' just is the range of F.

a. Must F be an injection into S'?

b. Must F be a surjection into S'?

Substantiate each answer with either a counterexample or a proof.

3*. Let $F_1 : S \to S'$ and $F_2 : S \to S'$ with the range of F_1 disjoint from the range of F_2.

a. If F_1 and F_2 are both injections into S', must $F_1 \cup F_2$ also be an injection into S'?

b. Can either F_1 or F_2 be a surjection into S'?

4. Verify that the function $F(x) = 2x$ from the set of natural numbers into the set of even natural numbers is indeed a one-one correspondence between these sets.

2.6 Converses & Ancestrals

Some relations bear interesting relations to other relations. The first example is when one relation is the *converse* of another. This idea is very intuitive: if R is a relation then the converse of R is the relation that holds between x and y just if R holds between y and x. For instance, the converse of *being taller than* is *being smaller than* since whenever x is taller than y, y is smaller than x.

Definition. (Converse) When R is a binary relation, a relation R' is its *converse* iff: $\langle x, y\rangle \in R$ iff $\langle y, x\rangle \in R'$.

You might be tempted to think that the converse of a relation is in some sense its *opposite*. However, be careful: every symmetric relation is its own converse, yet nothing is the opposite of itself! For example, since the relation *being a sibling of* is symmetric x is a sibling of y iff y is a sibling of x.

Another interesting relation between relations is exhibited by the pair of relations *is a parent of* and *is an ancestor of*. Clearly, the relation *is a parent of* is not transitive since Prince Charles is a parent of Prince Harry and Queen Elizabeth II is a parent of Prince Charles, but Elizabeth is not a parent of Harry. Yet she is his ancestor. There is a certain sense, then, in which the relation *is an ancestor of* is the result of making *is a parent of* transitive by filling in the missing links: it chains together instances of *is a parent of*. Due to this, we shall say that *is an ancestor of* is the *ancestral* of *is a parent of*.

Definition. (Ancestral) When R is a binary relation, a relation R' is its *ancestral* iff: $\langle x, y\rangle \in R'$ iff there are some $z_1, ..., z_n$ such that $\langle x, z_1\rangle \in R$ and , ... , and $\langle z_n, y\rangle \in R$.

To link this once again to personal identity, we know that the memory relation *x remembers being y* on the set of living people is not transitive, since I may not remember being a time slice of myself as a child. But personal identity is transitive on that set, so *x and y are the same person* cannot be the memory relation. But take the ancestral of the memory relation and the claim is more plausible.

To verify that you are following the definitions, it might help to convince yourself that every transitive relation is its own ancestral.

Worked Example. Show that if R is transitive its converse must be transitive too.

Answer. Let R' be the converse of R, and suppose that R is transitive. Suppose that $\langle x, y\rangle, \langle y, z\rangle \in R'$. We wish to show that $\langle x, z\rangle \in R'$ too. However, since R' is the converse of R, $\langle y, x\rangle, \langle z, y\rangle \in R$. Yet we are supposing that R is transitive, so $\langle z, x\rangle \in R$ too. Again, since R' is the converse of R, $\langle x, z\rangle \in R'$. $\square$

Exercises for 2.6

1. Choose an example of a function F from a set S into a set S' such that its converse is not a function from S' into S.

2. Let F be an injection from S into S'. Show that its converse is a function from S' into S.

2.7 Complexity

So far, the relations we have considered have all been intuitive and natural-looking. This needn't be the case, however. Recall that a *predicate* is defined very broadly: *any* deletion of names from a sentence and replacement with gaps is a predicate. Hence predicates can look very unnatural, e.g. 'x is the lecturer for Formal Logic in 2019/20 and y is the lecturer for Sets, Relations and Probability in 2019/20'. That expresses a relation that holds only between the pair of Owen and Alex. It's hard to imagine such a relation being of any real interest, but it's a relation nevertheless.

The generous definition also entails that predicates can have logical complexity, e.g. 'x is Donald Glover $\leftrightarrow$ y is Ariana Grande'. That's a perfectly reasonable predicate, expressing a perfectly reasonable relation, even though it contains logical complexity.

When logical complexity is involved, it can be difficult (and certainly unintuitive) to judge its properties on a given set. Let's work through this example on the set of currently living people. My first tip is not to attempt to have intuitions about the case (at least not until you're very comfortable with these) but instead to write down the long sentence that must be true if the relation is to have the property in question.

Let's start with reflexivity. The following sentence has to be true on the domain of currently living people if the relation is reflexive:

$\forall x$(x is Donald Glover $\leftrightarrow$ x is Ariana Grande)

Is that sentence true? Well, it certainly holds of all of us, since none of us are Donald Glover and none of us are Ariana Grande. How about Donald Glover himself? It is *false* of him, since he is Donald Glover but he is not Ariana Grande. By analogous reasoning, it is false of Ariana Grande. So the universal has false instances. So the sentence is false. The relation is *not* reflexive.

What about symmetry? To be symmetric on the same set, this sentence must be true:

1. $\forall x \forall y$((x is Donald Glover $\leftrightarrow$ y is Ariana Grande) $\rightarrow$ (y is Donald Glover $\leftrightarrow$ x is Ariana Grande))

Is this true? It's a conditional so it's only false when the antecedent is true and the consequent false. Hence it is only false when, for some value of x and y:

2. x is Donald Glover $\leftrightarrow$ y is Ariana Grande

3. $\neg$(y is Donald Glover $\leftrightarrow$ x is Ariana Grande)

There are only two ways that 2. can be true:

4. $x =$Donald Glover and $y =$Ariana Grande

5. $x \neq$ Donald Glover and $y \neq$ Ariana Grande

Let's start with 4. If that's the case, then 3 is false. Why? Because both sides of the biconditional are false, so the biconditional itself is true, so the whole sentence is false. So 4 is not a counterexample to 1.

We are left with 5. We need x to be anyone but Donald Glover and y anyone but Ariana Grande. Let's say they are Owen (not Donald Glover) and Alex (not Ariana Grande). But then both sides of the biconditional in 3 are false again, so the biconditional is true again, and its negation false. No counterexample.

But we haven't exhausted the possibilities in 5, since there are many ways of failing to be Donald Glover and Ariana Grande. In particular, let $x =$ Ariana Grande and $y =$ Owen. Now the biconditional in 3 is *false*, since its left-hand side is false (Owen $\neq$ Donald Glover) but its right-hand side is *true* (Ariana Grande = Ariana Grande). So its negation is true. This is a valuation that makes 2 (the antecedent of the conditional in 1) true but 3 (the consequent of the conditional in 1) false. So 1 has a false instance, so 1 is false. The relation is *not* symmetric.

That's a slow working-through of how we would decide whether a complex relation like this has some property or other. It's difficult to have intuitions about, so you should work through something like this process. I leave it as an exercise to judge whether the same relation possesses the other properties discussed, on the same set.

Exercises for 2.7

1. On the set of all people ever, determine whether the relation *x is Frege or y is Russell* is reflexive, whether it is symmetric and whether it is transitive.

2. Say that a binary relation R is *Euclidean* on a set just if for every $x, y, z \in S$, if $\langle x, y\rangle \in S$ and $\langle x, z\rangle \in S$, then $\langle y, z\rangle \in S$. On the set of living people, is the relation *x loves y only if y loves x*:

a. reflexive?

b. symmetric?

c. transitive?

d. Euclidean?

2.8 References & Further Reading

For further reading and other exercises, see the third chapter of Enderton, H. (1977) *Elements of Set Theory*. London: Elsevier. See also the second chapter of Steinhart, E. (2009) *More Precisely*. Ontario: Broadview.

3 Probability

Probabilistic reasoning is common in ordinary life. Before leaving the house, one might estimate the likelihood of it raining given that there are grey clouds in the sky. One might then try to calculate the likelihood of the tube being overcrowded given that one of the lines is closed down.

Probabilistic reasoning is also rife in philosophy. In the philosophy of science, we discuss the extent to which theories are confirmed probabilistically. In metaphysics, probabilistic accounts of causation are popular. In the philosophy of language, there are probabilistic accounts of conditionals. Probability is central to quantum mechanics, philosophy of biology, philosophy of statistics and philosophy of social science. Probability is also a central component of decision theory, game theory and their applications to e.g. moral and political philosophy. In the philosophy of mind, various mental states are often represented as probabilities (degrees of belief). And, of course, there is the philosophy of probability itself, which we'll touch on later. In this part of the course, we will focus on the probability calculus but also have something to say of its philosophical interpretation.

3.1 Fallacies

To highlight the importance of knowing how to reason with probability, we will first consider four common fallacies which people make when reasoning probabilistically. The account of probability developed in the following sections will offer further insight into why these patterns of reasoning are fallacious.

(i) **Gambler's Fallacy**

In the Budget Casino Nevada, gamblers bet on the simple game of *coin toss* in which a fair coin is flipped and bets are placed on it landing either heads or tails. On one evening, the coin is flipped twenty times in a row and each time it comes up heads. Gamblers rush to place high bets on the twenty-first flip landing tails, for they think it is extremely unlikely that the coin will land heads for a twenty-first time in a row.

Which fallacy do you think the gamblers are making?

Intuitively, the issue is that the outcome of the twenty-first coin flip is completely *independent* of the outcomes of previous coin flips. The outcomes of those previous coin flips don't constrain the likelihood of the coin landing heads in the twenty-first flip.

(ii) **Conjunction Fallacy**

The famous psychologists Tversky and Kahneman gathered data on responses to the following puzzle.

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. Which is more probable?

1. Linda is a bank teller.

2. Linda is a bank teller and is active in the feminist movement.

A majority of people responded with option 2. Perhaps their line of reasoning was that Linda is very likely to be active in the feminist movement, so it's more likely that she is both a bank teller and active in the movement.

However, this would be a *fallacious* line of reasoning. What do you think is fallacious about it?

The issue is that the probability of a conjunction of events occurring is always no greater than the probability of one of them occurring.

(iii) **Inverse Fallacy**

A true story: I once knew a probability theorist who went to the doctor to have a test for Tuberculosis. Before the having the test, the doctor explained that the test was highly accurate. In response, the probability theorist asked 'do you mean that

1. the probability that the test shows positive given that I have Tuberculosis is high?, or

2. the probability that I have Tuberculosis given that the test shows shows positive is high?'

In response, and much to the probability theorist's dismay, the doctor asked 'aren't they the same thing?'

What is fallacious about the doctor's identification of 1 and 2 in this scenario?

The issue can be explained in terms of false positives and false negatives. Perhaps the test is highly accurate in the sense that it results in very few false negatives: given that someone has Tuberculosis, it is highly likely that the test will show positive. However, whilst this might be the case, perhaps the test records quite a few false positives: there is a significant number of cases in which the test shows positive but the person does not indeed have Tuberculosis.

(iv) **Base Rate Fallacy**

Suppose that 85% of the taxis in Smallville are green and 15% are blue, and there was an accident involving a taxi. A witness reports that the taxi was blue. Let's say the witness correctly identifies the colour of taxis 80% of the time and incorrectly identifies them 20% of the time. How likely is it that the taxi was blue?

One might think that it is very likely that the taxi was blue, for the witness is very reliable. However, this would be fallacious. Why do you think that is the case?

This fallacy is tricker to spot than the others we've considered so far, but the issue is that we are ignoring the low *base rate* or initial probability that we're dealing with a blue taxi. The point is that, in spite of the reliability of our witness, the proportion of blue taxis in Smallville is only 15%. This means that the proportion of green taxis incorrectly identified as blue remains greater than the proportion of blue taxis correctly identified as blue. So, given the low base rate of blue taxis, it is still more likely that our very reliable witness misidentified a green taxi than correctly identified a blue taxi. It's just so unlikely that they saw a blue car at all.

Although it helps to have an informal understanding of the nature of these fallacies, it would be better to have a formal framework for reasoning probabilistically which will guarantee that one avoids making these fallacies in one's own reasoning.

3.2 Outcomes, Events & Fields

We will be assigning probabilities to *sets of outcomes*. There are other candidates in the literature, such as propositions and sentences, but sets of outcomes seem the most natural for various reasons.

Let's say that the probability of the Conservatives winning the next election is 40%. It seems natural to regard a Conservative victory as a set of outcomes since it includes the case where they win by 1 seat, by 100 seats, etc, and these are clearly different states of affairs. How finely the states of affairs are individuated will normally be settled by context.

Definition. (Outcome space) The set of possible outcomes in which we are interested will be called the *outcome space*, V. In other places, you will find this referred to as the *reference set* or the *sample space*. In this course, we shall only consider finite outcome spaces.

For example, if we throw a fair six-sided die, there are 6 possible outcomes. We might describe the outcome space as follows: $V = \{1, 2, 3, 4, 5, 6\}$. There are 6 possible outcomes but there are more than 6 *sets* of outcomes, which we will call *events*.

For example, we may want to know the probability of rolling an even number, or of rolling either a 1 or a 5. The first event would be represented as $\{2, 4, 6\}$ and the second as $\{1, 5\}$. How many events are there? The following definition answers this question.

Definition. (Field) For outcome space V, the events will be the members of $\mathcal{P}(V)$. The set of events $\mathcal{P}(V)$ will be called the *field*, F_V, of V.

Thus, for an outcome space with n members, we know that there will be 2^n events.

Notation. When the outcome space V is clear from the context, we shall sometimes write F_V as F.

Fact. The field contains the *certain* event, $\{1, 2, 3, 4, 5, 6\}$ and the impossible event, $\emptyset$. It is also *closed* under the operations of intersection, union and complement in the following senses.

1. If $X \in F_V$ and $Y \in F_V$, then $X \cap Y \in F_V$.
2. If $X \in F_V$ and $Y \in F_V$, then $X \cup Y \in F_V$.
3. If $X \in F_V$, $(V - X) \in F_V$.

Notation. When V is clear from the context, we shall write $\overline{X}$ for $V - X$. So, the final condition may be written as: if $X \in F_V$, $\overline{X} \in F_V$. Some texts will write X^* for $V - X$; we are not using that notation.

Exercises for 3.2

1. Let V be an outcome space. Show that the field F_V of V is indeed closed under the operations of intersection, union and complement, as stated in the above fact.

2. Let V be an outcome space. A *meadow* of sets M_V on V is defined as follows.

   1. $V \in M_V$
   2. If $X \in M_V$ and $Y \in M_V$, then $X \cup Y \in M_V$.
   3. If $X \in F_V$, $\overline{X} \in M_V$

Show that any meadow of sets on V is closed under intersection.

3.3 Probability Functions & Kolmogorov Axioms

The next step is to define a probability function on a field of events which will allow us to assign numbers to members of that field.

Intuitively, the probability function will assign the certain event the number 1, and every event will be assigned some non-negative real number.[3] In addition, for incompatible events like a fair coin landing heads and the same coin landing tails (in the same flip), the probability of one-or-the-other happening is simply the sum of their probabilities, i.e. the sum of the probability of the coin landing heads and the probability of the coin landing tails.

Definition. (Probability function) Let V be an outcome space. A *probability function on* F_V is a function Pr which assigns non-negative real numbers to members of F and obeys the following *Kolmogorov axioms* for any $X, Y \in F_V$:

Axiom 1. $Pr(V) = 1$

Axiom 2. $Pr(X) \geq 0$

Axiom 3. If $X \cap Y = \emptyset$, then $Pr(X \cup Y) = Pr(X) + Pr(Y)$

With just these constraints on probability functions, we are already in a position to establish a helpful fact about how they behave.

Worked Example. Show that $Pr(X) + Pr(\overline{X}) = 1$.

Answer.

$$\begin{aligned}
Pr(X) + Pr(\overline{X}) &= Pr(X) + Pr(V - X) && \text{Def. } \overline{X} \\
&= Pr(X \cup (V - X)) && \text{Axiom 3} \\
&= Pr(V) && \text{Defs. } \cup \text{ \& } - \\
&= 1 && \text{Axiom 1}
\end{aligned}$$

□

[3] The real numbers include numbers like $0, 1, \frac{1}{3}, 0.3333\ldots$, and so on.

Exercises for 3.3

1. Show that $Pr(\emptyset) = 0$.

2. Let V be an outcome space and let $X \in F_V$.

   i. Show that $Pr(\overline{X}) = 1 - Pr(X)$.

   ii. Conclude that $0 \leq Pr(X) \leq 1$.

3. Show that $Pr(A \cap B) \leq Pr(A)$. How does this fact connect to the conjunction fallacy?

3.4 Interpretation of Probability: Principle of Indifference

The Kolmogorov Axioms provide us with general structural constraints on probability functions. However, we wish to know how to assign probabilities to events. Our method will be first to assign probabilities to outcomes by assigning probabilities to their singleton sets. In simple cases, we already know how to do this. For example, with a fair six-sided die, the probability of the event that it lands 6 after being rolled is, of course, $\frac{1}{6}$. But what justifies this? Can we articulate the underlying principle which will guide our probability assignments in more complex cases? These are questions which pertain to the interpretation of our working notion of probability.

We might be tempted to substantiate our assignment of probability be appealing to the physical symmetry of the die. After all, it is a fair, evenly weighted die and so of course it will come up 6 on 1 in 6 rolls.

Indeed, we might try to articulate the general *principle of indifference* which underlies this idea this as follows: whenever there is no evidence favouring one event over another, each should be assigned the same probability as the others. In the words of John Maynard Keynes:

> The Principle of Indifference asserts that if there is no known reason for predicating of our subject one rather than another of several alternatives, then relatively to such knowledge the assertions of each of these alternatives have an equal probability. Thus equal probabilities must be assigned to each of several arguments, if there is an absence of positive ground for assigning unequal ones.
>
> Keynes, J. M. (1921) *A Treatise on Probability*, p. 42

This principle is at the heart of the *classical* interpretation of probability.

Nevertheless, we will **not** be justifying assignments of probability with the principle of indifference, for **the principle is inconsistent**. This inconsistency arises because the principle is overly sensitive to how events are described.

To appreciate this point, suppose I tell you only that a car travelled 100 miles at an average speed of between 50 and 100mph. What is the probability that the average speed was between 75 and 100mph? You may be tempted to say 0.5, because the ranges 50-75mph and 75-100mph are equally possible. But now I reformulate the problem: a car took between 1 and 2 hours to travel 100 miles. This is intuitively the same situation. What is the probability that the journey took between 1 hour and $1\frac{1}{3}$ hours? You may be tempted to say $\frac{1}{3}$ since the ranges $1 - 1\frac{1}{3}, 1\frac{1}{3} - 1\frac{2}{3}, 1\frac{2}{3} - 2$ hours are equally likely.

This is a problem: I have asked equivalent questions and the principle of indifference has delivered incompatible answers. This interpretation seems to make probability far too *sensitive to description*.

3.5 Interpretation of Probability: Frequentism

Instead of using the principle of indifference, our assignment of probabilities to events is due to their *frequencies*: the probability of rolling a 6 is $\frac{1}{6}$ because it has in fact come up 6 about one-sixth of the time.

This is still not an ideal interpretation of probability. There are difficult questions about how to assign probabilities to event types which have never had tokens. Moreover, there are the usual Humean worries about induction. We know that past experience does not make the future *certain* so it also doesn't make it *likely*. The procedure is, then, for the usual Humean reasons, irrational.[4]

Nevertheless, unlike the principle of indifference, it is a consistent way of assigning probabilities to events, and it will suffice for current purposes.

To return to the die example, we will say that $Pr(\{6\}) = \frac{1}{6}$. (Notice that we assign a probability to a *singleton* since the function is defined on the field of outcomes.) Similarly, for a standard deck of cards (52 different cards; 13 in each of 4 suits), the probability of drawing the king of clubs is $\frac{1}{52}$. In other words, $Pr(\{KC\}) = \frac{1}{52}$.

We can also work out the probability of other events. What is the probability that a card drawn from a standard deck of cards is either the ace of hearts or an even spade? We need to consider the relevant subset of V. Let's call it $T = \{2S, 4S, 6S, 8S, 10S, AH\}$. We want to know $Pr(T)$. From the third Kolmogorov axiom:

$$\begin{aligned} Pr(T) &= Pr(\{2S\}) + Pr(\{4S\}) + Pr(\{6S\}) + Pr(\{8S\}) + Pr(\{10S\}) + Pr(\{AH\}) \\ &= \frac{6}{52} \\ &= \frac{3}{26} \end{aligned}$$

It might help to note that $\frac{|T|}{|V|} = \frac{6}{52}$, where V is the outcome space.

For another example, let U be the event that an ace is drawn, i.e. $U = \{AH, AS, AC, AD\}$. What's the probability that an ace is drawn and its either the ace of hearts or an even spade? This event is just $(U \cap T) = \{AH\}$, so $Pr(U \cap T) = \frac{1}{52}$.

3.6 Conditional Probability

Next, we define the notion of *conditional probability*: the probability of one event occurring, given that another event occurs. For example, perhaps we want to calculate the probability of it raining in an hour given that the sky is grey. Or perhaps we want to

[4]See Hájek (2009) discussion of the philosophical issues surrounding even quite sophisticated forms of frequentism.

calculate the probability of a student getting a First given that they performed extremely well in their supervisions.

We define the conditional probability of A conditional on B in terms of the probability of both A and B occurring over the probability of B occurring. For example, suppose we are interested in probability of having drawn the ace of spades (A) given that an ace has been drawn (B). We know the probability that an ace is drawn is $\frac{4}{52}$. We also know that the probability of drawing both an ace and the ace of spades is just the probability of drawing the ace of spades, which is $\frac{1}{52}$. So, the probability of drawing the ace of spades conditional on drawing an ace is $\frac{1}{4}$. Intuitively, this is correct: *given* that I have drawn an ace, there are four cards I might have drawn, and the ace of spades is one of them.

Definition. (Conditional Probability)

$$Pr(A|B) = \frac{Pr(A \cap B)}{Pr(B)} \text{ providing } Pr(B) \neq 0$$

We pronounce '$Pr(A|B)$' as 'the probability of A *given* B'. The conditional probability of A given B remains undefined when $Pr(B) = 0$.

Worked Example. What is the probability that a fair six-sided die comes up 2, given that it comes up even?

Answer. Let $V = \{1, 2, 3, 4, 5, 6\}$. The outcomes where the die is even are $E = \{2, 4, 6\}$. The outcome that the die is 2 is $T = \{2\}$. So, applying the formula:

$$Pr(T|E) = \frac{Pr(T \cap E)}{Pr(E)}$$

But note that $T \cap E = T$ because $T \subseteq E$. Now $Pr(T \cap E) = Pr(T) = \frac{1}{6}$. And $Pr(E) = \frac{1}{2}$. So:

$$Pr(T|E) = \frac{\frac{1}{6}}{\frac{1}{2}} = \frac{1}{3}$$

□

This definition also puts us in a position to diagnose the issue with the inverse fallacy more explicitly. In certain cases, we can manipulate the definition of conditional probability to show that the probability of A given B is different from the probability of B given A.

For example, we've just worked out the probability that a fair six-sided die comes up 2, given that it comes up even. Yet we can also use the definition of conditional probability to show that the probability that it comes up even, given that it comes up 2 is 1: it's certain to happen.

$$Pr(E|T) = \frac{Pr(E \cap T)}{Pr(T)} = \frac{Pr(T)}{Pr(T)} = 1$$

Thus, $Pr(T|E) \neq Pr(E|T)$.

Finally, notice that we now have a helpful way to calculate the probabilities of intersections of events, by a simple rearrangement of the definition of conditional probability.

$$Pr(A \cap B) = Pr(A|B) \times Pr(B)$$

This gives us a means of calculating $Pr(A \cap B)$ (intuitively, A *and* B) when we have $Pr(A|B)$ and $Pr(B)$.

Exercises for 3.6

1. Jessica is dealt a single card from a standard deck.

   a. What's the probability that it's a 7 given that it's a spade?

   b. What's the probability that it's a spade given that it's a 7?

   c. What's the probability that it's a spade given that it's not a 7?

   d. What's the probability that it's neither a spade nor a club given that it's not a 7?

2. The Cambridge Scrabble Team has 10 members. The Team is divided in two: there are six A-team members and four B-team members. Four A-team members are also in the Cambridge Cluedo Club, as is one B-team member. A member of the Cambridge Scrabble Team is selected at random.

   a. What's the probability that they are in the Cluedo Club?

   b. What's the probability that they are in the Cluedo Club given that they are an A-team member?

   c. What's the probability that they are an A-team member given that they are in the Cluedo Club?

   d. What's the probability that they are in the Cluedo Club given they are in the A-team and the B-team?

3.7 Probabilistic Independence & Dependence

Intuitively, some events are *probabilistically independent* of other events. For example, suppose I roll a pair of fair dice at once. The outcomes of those rolls are probabilistically independent: the probability of either occurring does not depend on the probability of the other occurring. To appreciate this, notice that the probability of one die showing 4 is $\frac{1}{4}$, and also the probability of that die showing 4 given that the other shows 6 is still $\frac{1}{4}$: it makes no difference, probabilistically speaking.

This suggests a way of defining of probabilistic independence in terms of conditional probability.

Definition. (Probabilistic Independence)

A is *probabilistically independent* of B with respect to Pr iff $Pr(A|B) = P(A)$

This relation is symmetric, which you are asked to show in an exercise below.

Events are defined to be *probabilistically dependent* (with respect to a probability function) when they are not probabilistically independent (with respect to that function). For example, suppose I roll a single fair die. The event of its showing 4 is probabilistically dependent on whether it shows an even number: the probability that it shows 4 given that it shows an even number is higher than the probability that it shows 4.

Philosophically, it is interesting to note that an event might depend probabilistically on another event even though it is not caused by it. For example, the event of the fair die showing 4 depends probabilistically on whether it shows an even number, as we've just seen. Yet there is no sense in which the die showing an even number *caused* the die to show 4. These events occur simultaneously, so how could one cause the other?

Our definition of probabilistic independence has a helpful consequence. Namely, from it we can derive that when A and B are probabilistically independent, the probability of them both occurring is just the probability of one occurring by the probability of the other occurring.

Fact. If A is probabilistically independent of B with respect to Pr, then $Pr(A \cap B) = Pr(A) \times Pr(B)$.

A simple argument establishes this fact. Suppose A is probabilistically independent of B with respect to Pr, then we may reason as follows.

$$\begin{aligned} Pr(A \cap B) &= Pr(A|B) \times Pr(B) && \text{Def. } Pr(A|B) \\ &= Pr(A) \times Pr(B) && \text{Supposition} \end{aligned}$$

No such argument is available when A and B are probabilistically dependent.

Worked Example. Exactly 50% of the population are female and 30% of females have blue eyes. What is the probability that a randomly selected person is female *and* blue-eyed?

Answer. Let F be the event of selecting a female and B the event of selecting someone with blue eyes. We are given that $Pr(F) = 0.5$ and $Pr(B|F) = 0.3$. Rearranging the conditional probability formula and substituting:

$$Pr(B \cap F) = Pr(B|F) \times Pr(F) = 0.3 \times 0.5 = 0.15$$

□

We're also now in a position to appreciate the error embodied in the gambler's fallacy. In the initial case with the twenty-one coins flips, the successive coin flips are probabilistically independent of one another. Thus, the probability that the coin will lands heads on the twenty-first flip given that it has landed heads on the previous twenty flips is just the probability that the coin will lands heads on the twenty-first flip.

Exercises for 3.7

1. Show that probabilistic independence is symmetric.

2. Thinking back to the gambler's fallacy, show that the probability of flipping twenty-one heads in a row equals the probability of twenty heads in a row and then tails.

3.8 Bayes' Theorem

In this section, we'll use the probability calculus to derive two versions of a powerful theorem: Bayes' Theorem.

3.8.1 First & Second Versions

We begin with an example. Polly is a political scientist interested in the reasons for which people voted to leave the European Union. As we all know *ad nauseam*, 52% of the electorate voted to leave and 48% voted to remain. However, Polly conducts a highly accurate survey and comes to two conclusions. First, 59% of the electorate cited reasons pertaining to constitutional sovereignty as their main motivation for voting the way they did. Second, 21 out of 26 people who voted to leave cited reasons pertaining to constitutional sovereignty as their main motivation the way they did. Polly asks the following question.

Polly's Question: Suppose I randomly select a member of the electorate. What's the probability that, given the person cites reasons pertaining to constitutional sovereignty as their main motivation for the way they did, they voted to leave?

A powerful theorem of the probability calculus known as Bayes' Theorem allows one to find the answer to Polly's question. It allows one to calculate the probability of a hypothesis given some evidence on the basis of the probabilities of the evidence, the hypothesis, and the evidence given the hypothesis.

Theorem. (Bayes' Theorem, First Version)

$$Pr(H|E) = \frac{Pr(E|H) \times Pr(H)}{Pr(E)}$$

Proof.

$$\begin{aligned} Pr(H|E) &= \frac{Pr(H \cap E)}{Pr(E)} && \text{Def. } Pr(H|E) \\ &= \frac{Pr(E \cap H)}{Pr(E)} && \text{Def. } \cap \\ &= \frac{Pr(E|H) \times Pr(H)}{Pr(E)} && \text{Rearranging Def. } Pr(E|H) \end{aligned}$$

□

Using this version of Bayes' Theorem, Polly can answer her question. First, let's state the facts which Polly has already established.

H := the randomly selected person voted to leave (the 'hypothesis')

E := the randomly selected person cites reasons pertaining to constitutional sovereignty as their main motivation for the way they did (the 'evidence')

$Pr(H) = \frac{52}{100}$

$Pr(E) = \frac{59}{100}$

$Pr(E|H) = \frac{21}{26}$

We can then plug these data into Bayes' Theorem to calculate $Pr(H|E)$.

$$Pr(H|E) = \frac{\frac{21}{26} \times \frac{52}{100}}{\frac{59}{100}} = \frac{42}{59}$$

Polly knew the probabilities of the hypothesis, the relevant piece of evidence, and of that evidence given that hypothesis, so she could calculate the probability of that hypothesis *given* that evidence.

There is an interesting upshot: if some hypothesis predicts an observation that was very unlikely in advance, then actually making that observation is very strong evidence for the hypothesis. For example, suppose I hold some bizarre conspiracy theory H according to which the end of planet earth will occur this week. And suppose that on the basis of H I predict E, that there will be widespread earthquakes all over the world tomorrow. Today, one might think that $Pr(H)$ and $Pr(E)$ are very low. One might also think that $Pr(E|H)$ is very high, for argument's sake 1. This is because, although one doesn't think my bizarre theory has any plausibility, we're happy to allow that *given* it's correct, then there will be widespread earthquakes. In light of all this, *given* one observes widespread earthquakes tomorrow, then the likelihood of the hypothesis rises enormously (by a factor of $\frac{1}{Pr(E)}$).

The first version of Bayes' Theorem allows us to calculate $Pr(H|E)$ on the basis of knowledge of $Pr(E|H)$, $Pr(H)$ and $Pr(E)$. Sometimes, however, we may be unaware of $Pr(E)$. Nevertheless, there's a neat trick for available for figuring out $Pr(E)$ on the basis of other available information. This trick will lead us to the second version of Bayes' Theorem.

To see this, let's tweak the case of Polly so that she does not know the probability that a randomly selected member of the electorate cites reasons pertaining to constitutional sovereignty as their main motivation for voting the way they did, i.e. $Pr(E)$. However, let's suppose that Polly *does* know the probability that the person's main reason pertained to constitutional sovereignty given that they voted to remain, i.e. $Pr(E|\overline{H})$:

$Pr(E|\overline{H}) = \frac{17}{48}$

As I mentioned at the beginning, Polly already knows $Pr(\overline{H}) = \frac{48}{100}$. Now, Polly can use what she knows to calculate $Pr(E)$ using the following fact.

Fact. $Pr(E) = [Pr(E|H) \times Pr(H)] + [Pr(E|\overline{H}) \times Pr(\overline{H})]$

Proof. First, notice the following two set-theoretic facts, whose proofs are left as exercises to the reader (see exercise 3.8.1).

$$E = (E \cap H) \cup (E \cap \overline{H})$$

$$(E \cap H) \cap (E \cap \overline{H}) = \emptyset$$

Hence, by Axiom 3 of the probability calculus:

$$Pr(E) = Pr((E \cap H) \cup (E \cap \overline{H})) = Pr(E \cap H) + Pr(E \cap \overline{H})$$

Thus, by the formula for conditional probability:

$$Pr(E) = [Pr(E|H) \times Pr(H)] + [Pr(E|\overline{H}) \times Pr(\overline{H})]$$

□

Using this fact, we obtain the second version of Bayes' Theorem.

Theorem. (Bayes' Theorem, Second Version)

$$Pr(H|E) = \frac{Pr(E|H) \times Pr(H)}{[Pr(E|H) \times Pr(H)] + [Pr(E|\overline{H}) \times Pr(\overline{H})]}$$

Polly can now use this version of Bayes' Theorem to find $Pr(H|E)$. Plugging in the pieces of information which Polly knows, we obtain:

$$Pr(H|E) = \frac{\frac{21}{26} \times \frac{52}{100}}{(\frac{21}{26} \times \frac{52}{100}) + (\frac{17}{48} \times \frac{48}{100})} = \frac{42}{59}$$

Let's consider another example using this version of Bayes' Theorem.

Worked Example. There are two boxes, A and B. A contains 2 red balls and 8 yellow balls. B contains 7 red balls and 3 yellow balls. You randomly select a box (you don't know which) with equal probability and randomly select a ball from the box – it's red. What is the probability that you have box A?

Answer. Let's first label the following sets:

$A :=$ Box A was chosen

$R :=$ A red ball was picked

We want $Pr(A|R)$, which is:

$$\frac{Pr(R|A) \times Pr(A)}{[Pr(R|A) \times Pr(A)] + [Pr(R|\overline{A}) \times Pr(\overline{A})]}$$

We know the following values:

$Pr(R|A) = \frac{2}{10}$

$Pr(R|\overline{A}) = \frac{7}{10}$

$Pr(A) = \frac{1}{2}$

$Pr(\overline{A}) = \frac{1}{2}$

From all this, we can conclude the following:

$$Pr(A|R) = \frac{\frac{2}{10} \times \frac{1}{2}}{(\frac{2}{10} \times \frac{1}{2}) + (\frac{7}{10} \times \frac{1}{2})} = \frac{2}{9}$$

□

3.8.2 Miracles

The most famous application of Bayes' Theorem is Hume's treatment of miracles. Suppose that a witness reports that a man walked on water. Let's call this T, for 'testimony'. And let's call the walking-on-water M, for 'miracle'.

To calculate the probability that a miracle occurred given that the witness reported it, $Pr(M|T)$, we need $Pr(M)$, $Pr(T|M)$ and $Pr(T|\overline{M})$.

Let's use '1M' as shorthand for '1,000,000', and say that $Pr(M) = \frac{1}{1M}$, so a miracle is highly unlikely. Let's also say that $Pr(T|M) = \frac{9}{10}$ and $Pr(T|\overline{M}) = \frac{1}{100}$, so the testimony is highly reliable. Now we can apply Bayes' Theorem:

$$Pr(M|T) = \frac{\frac{0.9}{1M}}{\frac{0.9}{1M} + \frac{0.01 \times 999{,}999}{1M}} \approx \frac{1}{10,000}$$

In spite of the reliability of the witness, it is still *extremely* unlikely that the miracle occurred.

Let's give the last word to Hume:

> The plain consequence is (and it is a general maxim worthy of attention), "That no testimony is sufficient to establish a miracle, unless the testimony be of such a kind, that its falsehood would be more miraculous, than the fact, which it endeavours to establish; and even in that case there is a mutual destruction of arguments, and the superior only gives us an assurance suitable to that degree of force, which remains, after deducting the inferior."
>
> Hume, 'Of Miracles' pp. 115-116

3.8.3 Base Rate Fallacy

We can also apply Bayes' Theorem to diagnose the error present in the base rate fallacy. Consider our initial example of the taxis in Smallville. Recall that 85% of the taxis in Smallville are green, 15% are blue, and there was an accident involving a taxi. A highly reliable witness, who correctly identifies the colour of taxis 80% of the time and incorrectly identifies them 20% of the time reports that the taxi was blue. We asked the question: how likely is it that the taxi was blue?

Let's start to address the question more carefully than before by labelling some sets:

B := The taxi was blue

W := The witness said the taxi was blue

And we have the following values: $Pr(W|B) = 80\%$; $Pr(B) = 15\%$; $Pr(W|\overline{B}) = 20\%$; $Pr(\overline{B}) = 85\%$. Now:

$$\begin{aligned} Pr(B|W) &= \frac{Pr(W|B) \times Pr(B)}{[Pr(W|B) \times Pr(B)] + [Pr(W|\overline{B}) \times Pr(\overline{B})]} \\ &= \frac{0.8 \times 0.15}{(0.8 \times 0.15) + (0.2 \times 0.85)} \approx 41\% \end{aligned}$$

So the witness is probably wrong. But this may surprise us: after all, the witness is 80% reliable. What's gone wrong? We are ignoring the low *base rate* or initial probability that we're dealing with a blue taxi. To ignore this is to commit a *base rate fallacy*.

3.8.4 The Monty Hall puzzle

The Monty Hall puzzle is a famous problem to do with probability based on the American game show *Let's Make a Deal*, hosted by Monty Hall.[5] You're on a game show and can

[5] The original puzzle is due to the statistician Steve Selvin (1975).

choose one of three doors. Behind one of the doors is a 1975 Lincoln Continental, and behind the other two are goats. You want the car. You choose, for argument's sake, door 1. You then ask Monty Hall to tell you which of the remaining doors does not have the car behind it (or choose one at random if neither does). Monty Hall tells you that the car is *not* behind door 3, and opens it to reveal a goat. Monty Hall then offers you the option to switch to door 2. Question: should you switch?

Let H be the event that you initially picked the correct door. Let E be the event that Monty Hall opens door 3. Thus, $Pr(H) = \frac{1}{3}$ and $Pr(E|H) = \frac{1}{2}$. Notice also that:

$$Pr(E) = [Pr(E|H) \times Pr(H)] + [Pr(E|\overline{H}) \times Pr(\overline{H})] = \frac{1}{2}$$

Now, plugging all this into Bayes' Theorem:

$$Pr(H|E) = \frac{Pr(E|H) \times Pr(H)}{Pr(E)} = \frac{\frac{1}{2} \times \frac{1}{3}}{\frac{1}{2}} = \frac{1}{3}$$

So you should be twice as confident that the car is behind door 2. You should switch! (Assuming that you don't really want a goat.)

This is all highly counterintuitive. What's going on? When you made your original choice, your chance of being correct (the car is behind door 1) was clearly $\frac{1}{3}$. So your chance of being incorrect (the car is behind door 2 or 3) was $\frac{2}{3}$. But, Monty having eliminated door 3, the $\frac{2}{3}$ probability is now, if you like, all behind door 2.

Exercises for 3.8

1. Establish the following two set-theoretic facts.

$$E = (E \cap H) \cup (E \cap \overline{H})$$

$$(E \cap H) \cap (E \cap \overline{H}) = \emptyset$$

2. On Levin's farm there are 1500 crops: 900 food crops and 600 oil crops. 30 food crops and 60 oil crops have a nutrient deficiency. Levin is a diligent farmer and devises a chemical test to discern whether a crop is lacking in nutrient deficiency. The test, however, is not particularly reliable. When applied to a nutrient deficient crop, the test is positive 60% of the time. And when applied to a nutrient non-deficient crop it tests positive 10% of the time. A crop from Levin's farm is randomly selected. Calculate probability that:

a. the crop is nutrient deficient given that it is an oil crop.

b. the crop is an oil crop given that it is nutrient deficient.

c. the crop is not nutrient deficient given that it is a food crop.

d. the crop is nutrient deficient given that the test is positive.

e. the crop is not nutrient deficient given that the test is positive.

3.9 Repeated Trials

So far, we have only considered single die rolls or card draws, but we may well be interested in *repetition* of these. The approach we take to such cases will depend on whether the trials are probabilistically independent or not. The former is somewhat simpler, so let's start there.

3.9.1 Independent Trials

Suppose you toss a coin twice and you want to know the probability of tossing at least one heads. These are clearly independent events, since each coin toss has no bearing on the other.

To find out the probability of tossing at least one heads, we can break the process down into two steps.

Step 1. (Determine the outcome space) The first step with questions of this sort is to work out the outcome space. The outcome space of each toss considered individually is $V = \{H, T\}$. So, there are four outcomes for the entire trial.

Outcome 1: both heads

Outcome 2: first heads, second tails

Outcome 3: first tails, second head

Outcome 4: both tails

In other words, the outcome space for both tosses considered jointly is the set:

$$\{\langle H, H\rangle, \langle H, T\rangle, \langle T, H\rangle, \langle T, T\rangle\}$$

It is important to note that this set is just V^2, the Cartesian product of V with itself. Indeed, this suggests a more general fact, which is helpful to keep in mind.

Fact. Let V be an outcome space for some trial (e.g. rolling a die, or flipping a coin). The outcome space for that trial and an independent repetition of that trial is $V^2 = \{\langle x, y\rangle : x \in V \land y \in V\}$. Order is of course crucial here, so that we can distinguish events that are intuitively distinct.

Examples of independent repetitions of trials include flipping a fair coin two times successively, or drawing a card from a standard deck once, *replacing it*, and then making another draw from the pack.

Step 2. (Determine the relevant formula) The second step is to determine the relevant formula. To do this, it is helpful to consider the set whose probability we are after. So, recall that we are after the probability of tossing at least one heads. This event is the following set:

$$\{\langle H, H\rangle\} \cup \{\langle H, T\rangle\} \cup \{\langle T, H\rangle\}$$

This set has an important feature which one might notice immediately by considering the formula above: it is the union of incompatible (i.e. disjoint) events: they each have only one member and it is unique in each case. This is extremely helpful to notice, since it allows us to apply Axiom 3 of the probability calculus to figure out the probability of the entire event.

First, recall Axiom 3.

Axiom 3. If $X \cap Y = \emptyset$, then $Pr(X \cup Y) = Pr(X) + Pr(Y)$

Next, we use this to calculate the probability of the event:

$$Pr(\{\langle H, H\rangle\}) + Pr(\{\langle H, T\rangle\}) + Pr(\{\langle T, H\rangle\}) = \frac{1}{4} + \frac{1}{4} + \frac{1}{4} = \frac{3}{4}$$

Tip. As a rule of thumb, in questions about repeated independent trials which do **not** involve conditional probability, typically the answer will be available by a straightforward application of Axiom 3.

3.9.2 Dependent Trials

An example of a repeated trial where we have probabilistic *dependence* is repeated card draws from a deck *without replacement*. The probability of the second draw will depend on the first, since a card has been removed from the deck. Intuitively, the first draw *makes a difference* to the second, since the deck is reduced. Nothing like this was the case with the coin tosses.

Let's take a specific example of a repeated trial with dependence. Suppose we want to know the probability that at least one of our two card draws without replacement is an ace. We can break the process down into steps like before.

Step 1. (Determine the outcome space) The outcome space, V, for a single card draw has 52 members. In the case of two card draws without replacement, the outcome space cannot be V^2, since we cannot remove the same card twice. So, we first need to figure out what the outcome space is.

Think about a simulation of the case. If you draw a card from the deck and don't replace it, then one thing you can never do is draw that card again. So, regardless of which card you draw first, there will never be an outcome in which you draw it in the second draw. Thus, the outcome space is just the result of removing all of those 'double-draws' from V^2, which we know can be described as follows:

$$V^2 - \{\langle x, x\rangle : x \in V\}$$

This set has 52×51 members, each of whose members has equal probability.

Step 2. (Determine the relevant formula) Recall that we want to know the probability that at least one of our two card draws without replacement is an ace. It will first help to label some events to stop the notation getting too cumbersome:

A The first card is an ace and the second is not.

B The second card is an ace and the first is not.

C Both cards are aces.

We want $Pr(A \cup B \cup C)$. Now, note that the intersection of any pair of these sets is empty: every member of A has an ace in the first place, removing it from B, and a non-ace in second, removing it from C. Similar reasoning shows that all the sets are disjoint. So using Axiom 3 again:

$$Pr(A \cup B \cup C) = Pr(A) + Pr(B) + Pr(C)$$

What are these probabilities?

A has 4×48 members, so:

$$Pr(A) = \frac{4 \times 48}{52 \times 51}$$

By parity of reasoning, B has 4×48 members, so:

$$Pr(B) = \frac{4 \times 48}{52 \times 51}$$

The set C has 4×3 members, so:

$$Pr(C) = \frac{4 \times 3}{52 \times 51}$$

Consequently, our answer is:

$$Pr(A \cup B \cup C) = \frac{(4 \times 48) + (4 \times 48) + (4 \times 3)}{52 \times 51} = \frac{99}{13 \times 51} = \frac{33}{13 \times 17} = \frac{33}{221}$$

Tip. Often it is useful to keep the numbers in these unexpanded forms at first, since they will just cancel.

Exercises for 3.9

1. Suppose you draw two cards from a deck *with replacement*, what's the probability of:

   a. drawing two clubs?

   b. drawing a club first and a non-club second?

   c. drawing at least one club?

2. Alexei Ivanovich likes to play roulette, in which there is a wheel with thirty-seven numbers 0-36.

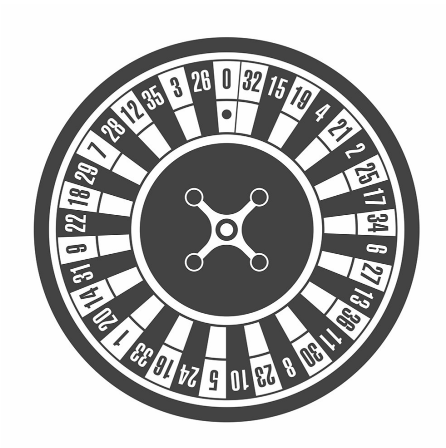

The wheel is spun, a ball is rolled onto it, and when the wheel stops the ball lands on a number. Alexei plays for three spins.

a. What's the probability of Alexei getting:

   i. 32, a 3, and a 21 in that order?

   ii. a 32, a 3, and a 21

   iii. three 1s?

b. Thinking back to the gambler's fallacy, show that the probability of flipping twenty-one heads in a row equals the probability of twenty heads in a row and then tails.

3. In *Thunderball*, one-hundred balls labelled '1'-'100' are placed into the Thundersphere and whirled around. Two draws are made at random in succession without replacement. What's the probability of:

a. drawing exactly two balls from 1-50?

b. drawing one ball from 1-50 and then two a ball from 51-100?

c. drawing at least one ball from 51-100?

3.10 Conditional Probability in Repeated Trials

In this section, we consider how to combine repeated trials with conditional probability. We begin with independent trials and then move on to the more complex dependent ones.

3.10.1 Conditional Probability in Independent Trials

Let's again take coin tosses as our example. Let's say you toss a fair coin three times and you want to know the probability that they are all heads given that at least one of them is.

Let V be the outcome space for a coin toss, the total outcome space will be:

$V^3 = \{\langle H,H,H\rangle, \langle H,H,T\rangle, \langle H,T,H\rangle, \langle H,T,T\rangle, \langle T,H,H\rangle, \langle T,H,T\rangle, \langle T,T,H\rangle, \langle T,T,T\rangle\}$

These 8 outcomes all have equal probability: $\frac{1}{8}$.

We want to calculate:

$Pr(\{\langle H,H,H\rangle\}|\{\langle H,H,H\rangle, \langle H,H,T\rangle, \langle H,T,H\rangle, \langle H,T,T\rangle, \langle T,H,H\rangle, \langle T,H,T\rangle, \langle T,T,H\rangle\})$

From the formula for conditional probability, this value is:

$$\frac{Pr(\{\langle H,H,H\rangle\} \cap \{\langle H,H,H\rangle, \langle H,H,T\rangle, \langle H,T,H\rangle, \langle H,T,T\rangle, \langle T,H,H\rangle, \langle T,H,T\rangle, \langle T,T,H\rangle\})}{Pr(\{\langle H,H,H\rangle, \langle H,H,T\rangle, \langle H,T,H\rangle, \langle H,T,T\rangle, \langle T,H,H\rangle, \langle T,H,T\rangle, \langle T,T,H\rangle\})}$$

Notice again that the top quantity reduces to $Pr(\{\langle H,H,H\rangle\}) = \frac{1}{8}$. And this needs to be divided by $\frac{7}{8}$.

$$\frac{\frac{1}{8}}{\frac{7}{8}} = \frac{1}{7}$$

Suppose we now want to calculate the probability that all three tosses are heads, given that the *first* is heads. In this case:

$$Pr(\{\langle H,H,H\rangle\}|\{\langle H,H,H\rangle,\langle H,H,T\rangle,\langle H,T,H\rangle,\langle H,T,T\rangle\}) = \frac{\frac{1}{8}}{\frac{4}{8}} = \frac{1}{4}$$

These results should give us pause. We have shown that the probability of tossing three heads given that *at least one* is heads is $\frac{1}{7}$. And the probability of tossing three heads given that *the first one* is heads is $\frac{1}{4}$. So the latter is almost twice as likely as the first. But that seems counterintuitive: if we know that at least one coin toss is a heads, how could it matter that the *first* one is heads?

Intuitively, the reason is that there are fewer outcomes in which you toss a head first than there are outcomes in which you just toss a head in any of the three throws. To see this, look at the calculations we made. In both cases, the numerator is the same. So our explanation for the difference must be in the denominator, which in the first case is $\frac{7}{8}$ and in the second $\frac{4}{8}$. In the second, our value is almost half. Why is this? Remember what this is measuring: in the first case, the number of outcomes with at least one head; in the second, the number of outcomes with a head in first place. There are fewer of these (4 rather than 7) so they are less likely to occur.

Moral. The lesson is that conditional probabilities can often yield surprising results. It is crucial that you apply the formula for conditional probability and don't just offer 'intuitive' reasoning, which is normally an utter disaster.

3.10.2 Conditional Probability in Dependent Trials

Finally, let's consider conditional probabilities and *dependent* trials. And again, let's take card draws without replacement as our example. Suppose two cards are drawn from a pack without replacement. Let's calculate the probability that both are aces given that one is an ace. Again, let's label some events:

$A :=$ Both cards are aces.

$B :=$ At least one of the cards is an ace.

From our earlier example, we know that $Pr(A) = \frac{4\times 3}{52\times 51} = \frac{1}{221}$. Similarly, we know that $Pr(B) = \frac{33}{221}$. We also know that $A \cap B = A$ so:

$$Pr(A|B) = \frac{Pr(A\cap B)}{Pr(B)} = \frac{Pr(A)}{Pr(B)} = \frac{\frac{1}{221}}{\frac{33}{221}} = \frac{1}{33}$$

Tip. It's always worth looking at previous results when answering these questions. Often, earlier results will save time in answering later questions.

Now let's calculate the probability that they are both aces given that one of them is *the ace of spades*. Let:

$C :=$ One of the cards is the ace of spades.

$C_1 :=$ The first card is the ace of spades (and the second is not). [51 outcomes]

$C_2 :=$ The second is the ace of spades (and the first is not). [51 outcomes]

$A \cap C_1$:= The first is the ace of spades; the second is another ace. [3 outcomes]

$A \cap C_2$:= The second is the ace of spades; the first is another ace. [3 outcomes]

We want $Pr(A|C)$ so we need $Pr(A \cap C)$ and $Pr(C)$. These are:

$$Pr(C) = Pr(C_1) + Pr(C_2) = \frac{102}{52 \times 51}$$

$$Pr(A \cap C) = Pr(A \cap C_1) + Pr(A \cap C_2) = \frac{6}{52 \times 51}$$

$$Pr(A|C) = \frac{6}{102} = \frac{1}{17}$$

Again, we might be surprised by these results. The probability that both cards are aces given that one is *the ace of spades* is almost double the probability that both are aces given that *at least one is an ace*. Why should this extra information be relevant? Again, conditional probabilities can be surprising.

Exercises for 3.10

1. Consider again the game of roulette. Suppose you play for two spins. What's the probability of:

a. getting a 7 given that both spins show a number between 1-7?

b. getting 7 twice given that you get at least one 7?

c. getting 7 twice given that you get 7 second?

d. getting an 8 given that one spin shows a number between 1-7?

2. Consider again the game of Thunderball. What's the probability of:

a. drawing exactly two balls from 1-50 given that you draw at least one from 1-50?

b. drawing exactly two balls from 1-50 given that you draw a 7?

3.11 Bayesian Epistemology

3.11.1 Subjective Probability

We have been working on the assumption that probabilities measure *frequencies*. We might call such probabilities *objective*. Consider quantum mechanics: certain events at this level are absolutely unpredictable. For example, a radium atom may decay in a given interval or it may not. Decaying and non-decaying atoms cannot be distinguished: all we can say is that any given atom has a certain probability of decaying. A radium atom has a 50% probability of decaying over an interval of 1,602 years – the *half-life* of the atom. This probability seems, in some sense, to be 'in the world'.

On the other hand, we might take probabilities as measuring our confidence in a particular outcome. Such probabilities would be *subjective*.

For example, let's say that you leave your home with your umbrella and sunglasses. Do you think that it will rain? Well, you can't be *certain* or you wouldn't take your sunglasses.

But then you can't be *certain* that it will be sunny or you wouldn't take your umbrella. Rather, you have a *degree of confidence*, or *credence*, that it will rain, and a degree of confidence that it will be sunny, and we can represent these using probabilities. Of course, for most mental states this is the norm, since we aren't certain of all that much.

The *Bayesian* claim is that, after you learn some new piece of evidence E, your new confidence in any hypothesis is your old $Pr(H|E)$. This is referred to as the *rule of conditionalization*:

Definition. (Rule of Conditionalization)

$$Pr_{new}(H) = Pr_{old}(H|E)$$

In other words, you should be constantly *updating* your confidence in hypotheses in light of new evidence, using conditional probability.

How can we judge your degree of belief that some event will occur? An obvious suggestion is observation of your *behaviour*. Let's say you go to the pub, you buy two pints and you take a seat at a table for two, checking your watch often. This all suggests that your degree of belief that someone will be joining you is high. At the extreme end of this thought, we could *identify* your belief with this behaviour. That would be a strong form of *behaviourism* about the mental.

That's all quite imprecise, so let's focus in on a subset of your behaviours: your willingness to accept betting odds. This approach to subjective probability was first offered by Frank Ramsey in his 'Truth and Probability' (1926). The approach is an intuitive one: to make your degrees of belief amenable to probabilistic treatment, we need to assign numerical values to your beliefs. And betting is an area where people often attach numerical values to their beliefs.

To understand this approach, it will be helpful to first have a reminder about betting odds, for those of you who haven't had a flutter for a while. Our standard example will be the odds on offer at various bookies' for Glastonbury 2021 headliners.

Skybet are offering 5/1 (pronounced '5 to 1') that Paul McCartney will headline. What does this mean? If you bet £1 and you win, you will win £6 ($5 \times$ £1 + your initial stake of £1 back). Generally, if the odds are A/B and you win, you'll win £$\frac{A+B}{B}$ for every £1 bet. The higher this figure – $\frac{A+B}{B}$ – the *better* the odds. If another bookies' were to offer, say, 4/1 on McCartney's headlining, they would be offering *worse* odds.

McCartney's odds of 5/1 are sometimes expressed as '5 to 1 *against*', since the amount you would win is greater than the amount staked. Intuitively, he is *unlikely* to headline. I see that the odds for The Cure headlining are 1/2. Here, the second number is larger than the first. Odds such as this are sometimes expressed as '2 to 1 *on*' since the amount you would win is less than the amount staked (here, for every £1 staked, you win 50p). Intuitively, The Cure are *likely* to headline.

The credence that you have in event E is determined by the *worst* odds that you would accept, at least for small stakes. Let's say I believe that Lady Gaga will headline Glastonbury. How confident am I? About 80%. If your credence in E is $\frac{B}{A+B}$, then the worst odds you will be prepared to accept on E are A/B. So the worst odds I would accept on Lady Gaga headlining are 1/4.

If a bookies' were offering 1/10 on Lady Gaga headlining, I wouldn't take the bet, since those are *worse* odds than 1/4. 1/10 represents a credence of about 90%, and my credence isn't that high. If another bookies' were offering 1/3 on Lady Gaga headlining, which represents a credence of 75%, I would take the bet, since my credence is higher.

This all makes intuitive sense: if you are *certain* that E will happen, you will accept *any* odds A/B (however small A is, as long as $A > 0$). And if you are certain that E *won't* happen, then you won't accept any odds A/B (however small B is, as long as $B > 0$).

3.11.2 Dutch Books

If we interpret the probability calculus subjectively, using betting odds, then the Kolmogorov axioms have a rational compulsion. It turns out that, if your reasoning fails to conform to any one of them, you will accept a bet that you are guaranteed to lose. To see this we need to first discuss the notion of a *Dutch book*, which we'll introduce by example.

Let's say it's the final of Rupaul's Drag Race and one of three queens will win: Alaska, Roxxy and Jinkx. You discover that Ladbroke's is offering 3/1 on Alaska, 3/1 on Roxxy and 3/1 on Jinkx. What should you do? You should take every penny you have and put $\frac{1}{3}$ of your money on each: you are *bound* to win. Why? Let's say you start with £3 and put £1 on each. If Alaska wins, you'll get £4 back; if Roxxy wins, you'll get £4 back; if Jinkx wins, you'll get £4 back. And one of them will win so, in any case, you end up with more than you started.

The bookies is bound to lose. Of course, no actual bookies would set up such odds and guarantee that they lose money. (Although, apparently, such situations can arise in international sport.) If you accept betting odds that guarantee that you'll lose money, we say that you're open to a *Dutch book*. In this example, the bookies is open to a Dutch book. But of course bettors themselves can also be open to Dutch books, if they accept the wrong odds.

If our credences fail to conform to the axioms of probability theory, we leave ourselves open to Dutch books. Let's say I have a credence of 60% that it will snow tomorrow. And I also have a credence of 60% that it won't snow tomorrow. This plainly violates the probabilisitic theorem we proved that $Pr(X) + Pr(\overline{X}) = 1$. And I am open to a Dutch book: 60% credence shows that I'd accept odds of 2/3 that it will snow tomorrow and odds of 2/3 that it won't.

Let's say that I start out with £6 and I put £3 on each outcome. It will either snow or not snow, so I'll lose one of the bets and win the other. The one I lose will cost me £3. The one I win will give me £5. So I end up with £5 – overall, I've lost £1. A Dutch book can be made against you if, and only if, your credences fail to conform to the axioms of probability.

This is all very exciting but it's important to note some of the idealisations being made here. We are assuming that, for any person S at any time t and for any outcome E, there will be a number between 0 and 1 that represents S's credence in E at t. There are obvious counterexamples to be had: some outcomes we obviously haven't entertained. But, even in the cases we have considered, is it plausible that we can quantify credences in this way? Or are the idealisations involved so great that they shed little light on actual human behaviour? Questions for another time but, if you're interested, I recommend starting with Ramsey's 'Truth and Probability'.

3.11.3 Expected Utility

In principle, any action can be construed as a sort of gamble. Suppose you decided (wisely) to revise the material on functions before your exams. So presumably your credence that the material on functions will be on the exam was greater than your belief it wouldn't be. Your credences in the exam happening were presumably also very high. You weren't *certain*: the exam could have been cancelled and you could have missed the message, but that's unlikely.

In general, when someone performs an action, the *expected utility* of that action is greater than other actions. A very famous example of expected utility is Pascal's wager. Let's say that your credence in the existence of God is 1%. If you are correct that God doesn't exist, you'll take some pleasure in being correct. Let's say this has positive utility 10. If God exists, though, and you don't believe in Her, you'll suffer eternal damnation. Let's say this has negative utility 1M.

How about if you *do* believe in God? Well, if She exists, you'll enjoy eternal happiness. Let's say this has positive utility 1M. And if you believe and She doesn't exists, you'll take some displeasure in being wrong, say negative utility 10.

We can helpfully display these data in a table:

	B	$\overline{B}$
G	+1M	-1M
$\overline{G}$	-10	+10

And we've said that $Pr(G) = 0.01$ and $Pr(\overline{G}) = 0.99$. The expected utility of believing in God is $(1M \times 0.01) + (-10 \times 0.99) = 9,990.1$. The expected utility of not believing is $(-1M \times 0.01) + (10 \times 0.99) = -9990.1$. So you should believe in God.

3.12 References & Further Reading

Hájek, A. (2009) 'Fifteen Arguments Against Hypothetical Frequentism', in *Erkenntnis* 70: 211-235.

Hume, D. ((2000) 1748) 'Of Miracles', in *An Enquiry Concerning Human Understanding*. Oxford: Oxford University Press.

Keynes, J. M. (1921) *A Treatise on Probability*. London: MacMillan.

Ramsey, F. P. (1926) 'Truth and Probability', in Ramsey, F. P. (1978) *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*. Mellor, D.H. (Ed.) London: Routledge & Kegan Paul.

Selvin, S. (1975) 'A Problem in Probability (Letter to the Editor)', in *The American Statistician* 29: 67.

For further reading on the philosophy of probability see Rowbottom, D. P. (2015) *Probability*. Polity Press.

4 Solutions

4.1 Exercises for 1.1-1.2

1. Are the following claims true, false, or ill-formed?

a. $0 \in \{0, 1, 2\}$ True

b. $0 \in \{1, 2\}$ False

c. $0 \in \{\{0\}\}$ False

d. $0 \in 0$ Ill-formed

e. $0 \in \emptyset$ False

2. Write down the following claims in our notation. Are the they true, false, or ill-formed?

a. The number 2 is a member of the set containing exactly Jack and Jill.

$2 \in \{\text{Jack}, \text{Jill}\}$, False

b. The empty set is a member of the number 2.

$\emptyset \in 2$, Ill-formed

c. The number 2 is a member of the set containing exactly the empty set and the number 2.

$2 \in \{\emptyset, 2\}$, True

4.2 Exercises for 1.3

1. Why is the set {The Prime Minister on 10 February 2020, Jeremy Corbyn} identical with the set {Boris Johnson, Jeremy Corbyn}? What does this tell us about how we describe sets?

Boris Johnson is The Prime Minister on 10 February 2020, so by **Extensionality** these sets are one and the same. Thus, one and the same set may be described in different ways.

2. Use **Extensionality** to argue that there is only one set which contains exactly the empty set.

If S and S' both contain exactly the empty set, then they have exactly the same members. After all, there is only one empty set, by **Extensionality**. So, again by **Extensionality**, S and S' are one and the same.

3.* If we enrich the language of first-order logic (with identity) with the predicate '$\in$' and the name '$\emptyset$' for the empty set, then we can define the predicate 'Set'. How?

$$Set(x) \;:=\; x = \emptyset \vee \exists y(y \in x)$$

4.3 Exercises for 1.4

1. Using our notation, write down the claim that the number zero is not a member of the set of all humans.

$0 \notin \{x : x \text{ is human}\}$

2. Using our notation, redescribe the set $\{\emptyset\}$.

$\{x : x = \emptyset\}$

3. Using our notation, redescribe $\emptyset$.

$\{x : x \neq x\}$

4. Write out in our formal language the claim that everything is the one and only member of its singleton.

$\forall x(x \in \{x\} \wedge \forall z(z \in \{x\} \rightarrow x = z))$

4.4 Exercises for 1.5

1. Define the proper superset relation.

$S \supset S' \;:=\; S \supseteq S' \wedge S' \not\supseteq S$

2. Show that if $S_1 \subseteq S_2$ and $S_2 \subseteq S_3$, then $S_1 \subseteq S_3$.

Suppose that $S_1 \subseteq S_2$ and $S_2 \subseteq S_3$. For arbitrary x, if $x \in S_1$ then $x \in S_2$ since $S_1 \subseteq S_2$. Yet if $x \in S_2$ then $x \in S_3$ since $S_2 \subseteq S_3$. Thus, if $x \in S_1$ then $x \in S_3$, so $S_1 \subseteq S_3$.

3. Let $A = \{0, 1\}$ and $B = \{0, 1, 2, 3\}$. Are the following claims true, false, of ill-formed?

a. $A \subseteq B$ True

b. $A \subset B$ True

c. $A \in B$ False

d. $0 \in A$ True

e. $0 \subseteq B$ Ill-formed

4.5 Exercises for 1.6

1. Show that $S \cup S$ is just S.

$$
\begin{aligned}
x \in (S \cup S) &\leftrightarrow x \in S \vee x \in S && \text{Def. } \cup \\
&\leftrightarrow x \in S && \text{PL}
\end{aligned}
$$

Like the examples in the lecture notes, the line justified by **PL** uses a tautology of propositional logic: $(A \vee A) \leftrightarrow A$.

2. Show that $S \cap \emptyset$ is just $\emptyset$.

$(S \cap \emptyset) = \{x : x \in S \wedge x \in \emptyset\} = \emptyset$

3. Show that if S and S' are disjoint and S' is non-empty, then $S \subset (S \cup S')$.

Clearly, $S \subseteq (S \cup S')$. To show that $S \subset (S \cup S')$, notice that S and S' are disjoint and S' is non-empty, so there is at least one thing $x \in S'$ such that $x \notin S$. Thus, $x \in (S \cup S')$, but $x \notin S$. Hence, $S \subset (S \cup S')$.

4. Show that $(S_1 \cup (S_2 \cap S_3)) = ((S_1 \cup S_2) \cap (S_1 \cup S_3))$

$x \in (S_1 \cup (S_2 \cap S_3))$	iff $x \in S_1$ or $x \in (S_2 \cap S_3)$	Def. $\cup$
	iff $x \in S_1$ or ($x \in S_2$ and $x \in S_3$)	Def. $\cap$
	iff ($x \in S_1$ or $x \in S_2$) and ($x \in S_1$ or $x \in S_3$)	**PL**
	iff $x \in (S_1 \cup S_2)$ and $x \in (S_1 \cup S_3)$	Def. $\cup$
	iff $x \in ((S_1 \cup S_2) \cap (S_1 \cup S_3))$	Def. $\cap$

Like the examples in the lecture notes, the line justified by **PL** uses a tautology of propositional logic: $(A \vee (B \wedge C)) \leftrightarrow ((A \vee B) \wedge (A \vee C))$.

4.6 Exercises for 1.7

1. Show that $(S - (S \cap S')) = S - S'$.

$x \in (S - (S \cap S'))$	iff $x \in S$ and $x \notin (S \cap S')$	Def. $-$
	iff $x \in S$ and not: ($x \in S$ and $x \in S'$)	Def. $\cap$
	iff $x \in S$ and $x \notin S'$	**PL**
	iff $x \in (S - S')$	Def. $-$

2. Show that $((S \cup S') - (S \cap S')) = ((S - S') \cup (S' - S))$.

$x \in ((S \cup S') - (S \cap S'))$	iff $x \in (S \cup S') \wedge x \notin (S \cap S')$	Def. $-$
	iff $(x \in S \vee x \in S') \wedge \neg(x \in S \wedge x \in S')$	Defs. $\cup, \cap$
	iff $(x \in S \wedge x \notin S') \vee (x \in S' \wedge x \notin S)$	**PL**
	iff $(x \in (S - S')) \vee (x \in (S' - S))$	Def. $-$
	iff $x \in ((S - S') \cup (S' - S))$	Def. $\cup$

Like the examples in the lecture notes, the line justified by **PL** uses a tautology of propositional logic: $((A \wedge \neg B) \vee (B \wedge \neg A)) \leftrightarrow ((A \vee B) \wedge \neg(A \wedge B))$.

3. Show that $((S \cup S') - (S \cap S')) = (S \cup S')$ when S and S' are disjoint.

Since S and S' are disjoint, $(S - S') = S$ and $(S' - S) = S'$. Thus, $((S - S') \cup (S' - S)) = (S \cup S')$, and so by Exercise 1.7.2, $((S \cup S') - (S \cap S')) = (S \cup S')$.

4.7 Exercises for 1.8

1. Is any set its own power set?

No: the cardinality of a power set is always greater than the cardinality of the set in question.

2. Define $\mathcal{P}(S) \cup S$ using our notation for describing sets.

$\{x : x \subseteq S \vee x \in S\}$

3. Show that $\mathcal{P}(S \cap (S - S)) \subseteq \mathcal{P}(S')$.

If $x \in \mathcal{P}(S \cap (S - S))$ then $x \in \mathcal{P}(S \cap \emptyset)$	Def. $-$
then $x \in \mathcal{P}(\emptyset)$	Def. $\cap$
then $x = \emptyset$	Def. $\mathcal{P}$
then $x \in \mathcal{P}(S')$, for any set S'	Def. $\mathcal{P}$

4.8 Exercises for 1.9

1. Let $A = \{a\}$, $B = \{b, c\}$, and $C = \{d\}$. Calculate the following.

i. $A \times B$

$\{\langle a, b\rangle, \langle a, c\rangle\}$

ii. $\mathcal{P}(A) \times B$

$\{\langle \{a\}, b\rangle, \langle \{a\}, c\rangle, \langle \emptyset, b\rangle, \langle \emptyset, c\rangle\}$

iii. $\mathcal{P}(A \times C)$

$\{\emptyset, \{\langle a, d\rangle\}\}$

2. Calculate $\emptyset^2$.

$\emptyset^2 = \emptyset \times \emptyset = \{\langle x, y\rangle : x \in \emptyset \wedge y \in \emptyset\} = \emptyset$

3. Define the set $(A \times B) \cap (A \times A)$. Assuming A and B are non-empty, under which conditions is $(A \times B) \cap (A \times A)$ non-empty?

$\{\langle x, y\rangle : x \in A \wedge y \in (A \cap B)\}$. It is non-empty when A and B are not disjoint.

4.9 Exercises for 1.10

1. Let $A = \{1, 2\}$. Calculate the following.

i. $|A - \{2\}|$

$|A - \{2\}| = 1$

ii. $|A \times (A - \{2\})|$

$|A \times (A - \{2\})| = 1$

iii. $|\mathcal{P}(A)|$

$|\mathcal{P}(A)| = 4$

4.10 Exercises for 2.1

1. If $x, y \in S$, then specify a set S' in terms of S such that $\langle x, y\rangle \in S'$.

$\{\langle x, y\rangle : x, y \in S\}$

4.11 Exercises for 2.2-2.3

1. Consider any set S and any binary relation which is not reflexive on S. Specify a condition on S which guarantees that the relation is also irreflexive on S.

S is a singleton set.

2. Let $S \neq \emptyset$. Show that $\emptyset$ is not reflexive on S.

Since $S \neq \emptyset$, there is some $x \in S$. Yet of course $\langle x, x\rangle \notin \emptyset$.

3. Let $S = \{1, 2, 3\}$ and let $R = \{\langle 1, 2\rangle, \langle 2, 3\rangle, \langle 3, 2\rangle\}$.

a. Is R reflexive on S? No

b. Is R symmetric on S? No

c. Is R anti-symmetric on S? No

d. Define a relation R' such that $R \subseteq R'$ and R' is transitive on S. Must R' be an equivalence relation on S? No. Consider: $R \cup \{\langle 2, 2\rangle, \langle 3, 3\rangle, \langle 1, 3\rangle\}$

e. Is $R - \{\langle 3, 2\rangle\}$ anti-symmetric on S? Yes

4. Show that every relation which is asymmetric on a set is also anti-symmetric on that set.

A relation R is asymmetric on a set S iff for **any** $x, y \in S$, if $\langle x, y\rangle \in R$ then $\langle y, x\rangle \notin R$. Thus, if $x \neq y$, then it's still the case that if $\langle x, y\rangle \in R$ then $\langle y, x\rangle \notin R$.

5. A relation is *anti-transitive* on a set S iff for every $x, y, z \in S$, if $\langle x, y\rangle \in R$ and $\langle y, z\rangle \in R$, then $\langle x, z\rangle \notin R$.

a. Write out the definition of anti-transitivity on a set in our formal language of set theory.

$\forall x \forall y \forall z(Rxy \wedge Ryz \to \neg Rxz)$

b. Can a relation be both reflexive and anti-transitive on a single non-empty set? Substantiate your answer.

Let S be any non-empty set and R any reflexive relation on S. Since S is non-empty, there is some $d \in S$ such that $\langle d, d\rangle \in R$. Yet that constitutes a counterexample to anti-transitivity (just assign d as the value of all of the variables x, y, z).

c. Provide two examples of a pair of a relation and a set such that the relation is both transitive and anti-transitive on a that set.

The empty relation on a non-empty set. Any relation on empty set.

6. Let S be the set of natural numbers. Let $\langle x, y\rangle \in R$ iff $x + 1 = y$. Show that R is a function on S.

Clearly, for any natural number n there is only one natural number $m = n + 1$. Thus, for each natural number n there is only one natural number m such that $\langle n, m\rangle \in R$. So R is a function on S.

4.12 Exercises for 2.5

1. In our formal language, write '$Func^S(R)$' to abbreviate the formula that R is a function on S, i.e. $\forall x\forall y\forall z((x \in S \wedge y \in S \wedge z \in S) \to ((Rxy \wedge Rxz) \to y = z))$. Using this abbreviation write out the following claims in the language:

a. R is function on S which is a surjection into S'.

$Func^S(R) \wedge \forall x(x \in S' \to \exists y(y \in S \wedge Ryx))$

b. R is function on S which is an injection into S'.

$Func^S(R) \wedge \forall x\forall y((x \in S \wedge y \in S \wedge x \neq y) \to (Rxz_1 \wedge Ryz_2 \to z_1 \neq z_2)))$

2. Let $F : S \to S'$ be such that S' just is the range of F.

a. Must F be an injection into S'?

No. Let F be a function on the set of natural numbers defined as follows:

$$F(n) = \begin{cases} n, & \text{if } n \text{ is even} \\ n-1, & \text{is } n \text{ is odd} \end{cases}$$

$$\begin{tikzcd} 0 \arrow[d] & 1 \arrow[dl] & 2 \arrow[d] & 3 \arrow[dl] & \cdots \\ 0 & & 2 & & \cdots \end{tikzcd}$$

The range of F is the set of even numbers, but F isn't an injection from the natural numbers into the set of even numbers.

b. Must F be a surjection into S'?

Yes, since $S' = ran(F)$.

Substantiate each answer with either a counterexample or a proof.

3. Let $F_1 : S \to S'$ and $F_2 : S \to S'$ with the range of F_1 disjoint from the range of F_2.

a. If F_1 and F_2 are both injections into S', must $F_1 \cup F_2$ also be an injection into S'?

Actually, there's no guarantee that $F_1 \cup F_2$ is even a function into S'. Consider the following counterexample:

$S = \{0\}$

$S' = \{1, 2\}$

$F_1 = \{\langle 0, 1\rangle\}$

$F_2 = \{\langle 0, 2\rangle\}$

$F_1 \cup F_2 = \{\langle 0, 1\rangle, \langle 0, 2\rangle\}$

b. Can either F_1 or F_2 be a surjection into S'?

Yes. Consider the following example:

$S = \{0\}$

$S' = \{1\}$

$F_1 = \{\}$

$F_2 = \{\langle 0, 1\rangle\}$

F_1 is the empty relation (which is an injective function on any set), and $ran(F_2) = S'$. Since F_1 is the empty relation, $ran(F_1)$ is disjoint from $ran(F_2)$.

4. Verify that the function $F(x) = 2x$ from the set of natural numbers into the set of even natural numbers is indeed a one-one correspondence between these sets.

F is an injection, since if $F(x) = F(y)$ then $2x = 2y$ and so $x = y$. Similarly, F is a surjection, since for every even number n there is some natural number m such that $n = 2m$.

4.13 Exercises for 2.6

1. Choose an example of a function F from a set S into a set S' such that its converse is not a function from S' into S.

Consider the following counterexample:

$S = \{0, 1\}$

$S' = \{2\}$

$F = \{\langle 0, 2\rangle, \langle 1, 2\rangle\}$

Since F isn't an injection into S', the converse of F isn't even a function into S.

2. Let F be an injection from S into S'. Show that its converse is a function from S' into S.

Let's write F' for the converse of F. Suppose that $\langle x, y\rangle \in F'$ and $\langle x, z\rangle \in F'$. Then $x = F(y)$ and $x = F(z)$ by the definition of a converse. Yet F is an injection into S, so $y = z$.

4.14 Exercises for 2.7

1. On the set of all people ever, determine whether the relation *x is Frege or y is Russell* is reflexive, whether it is symmetric and whether it is transitive.

It's neither reflexive nor symmetric on that set, but it is transitive on that set.

2. Say that a binary relation R is *Euclidean* on a set just if for every $x, y, z \in S$, if $\langle x, y \rangle \in S$ and $\langle x, z \rangle \in S$, then $\langle y, z \rangle \in S$. On the set of living people, is the relation *x loves y only if y loves x*:

a. reflexive? Yes

b. symmetric? No

c. transitive? No

d. Euclidean? No