

ECS315 2012/1 Part II Dr.Prapun

5 Probability Foundations

Constructing the mathematical foundations of probability theory has proven to be a long-lasting process of trial and error. *The approach consisting of defining probabilities as relative frequencies in cases of repeatable experiments leads to an unsatisfactory theory.* The frequency view of probability has a long history that goes back to **Aristotle**. It was not until 1933 that the great Russian mathematician A. N. **Kolmogorov** (1903-1987) laid a satisfactory mathematical foundation of probability theory. He did this by taking a number of axioms as his starting point, as had been done in other fields of mathematics. [22, p 223]

We will try to avoid several technical details¹⁴ ¹⁵ in this class. Therefore, the definition given below is not the “complete” definition. Some parts are modified or omitted to make the definition easier to understand.

¹⁴To study formal definition of probability, we start with the *probability space* (Ω, \mathcal{A}, P) . Let Ω be an arbitrary space or set of points ω . Recall (from Definition 1.15) that, viewed probabilistically, a subset of Ω is an *event* and an element ω of Ω is a *sample point*. Each event is a collection of outcomes which are elements of the sample space Ω .

The theory of probability focuses on collections of events, called event *σ -algebras*, typically denoted by \mathcal{A} (or \mathcal{F}), that contain all the events of interest (regarding the random experiment \mathcal{E}) to us, and are such that we have knowledge of their likelihood of occurrence. The probability P itself is defined as a number in the range $[0, 1]$ associated with each event in \mathcal{A} .

¹⁵The class 2^Ω of all subsets can be too large for us to define probability measures with consistency, across all member of the class. (There is no problem when Ω is countable.)

Definition 5.1. Kolmogorov's Axioms for Probability [12]:
 A **probability measure**¹⁶ is a real-valued (set) function¹⁷ that satisfies

P1 Nonnegativity:

$$P(A) \geq 0.$$

P2 Unit normalization:

$$P(\Omega) = 1.$$

P3 Countable additivity or **σ -additivity**: For every countable sequence $(A_n)_{n=1}^{\infty}$ of disjoint events,

$$P\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} P(A_n).$$

- The number $P(A)$ is called the **probability** of the event A
- The entire sample space Ω is called the **sure event** or the **certain event**.
- If an event A satisfies $P(A) = 1$, we say that A is an **almost-sure event**.
- A **support** of P is any set A for which $P(A) = 1$.

From the three axioms¹⁸, we can derive many more properties of probability measure. These properties are useful for calculating probabilities.

¹⁶Technically, probability measure is defined on a σ -algebra \mathcal{A} of Ω . The *triple* (Ω, \mathcal{A}, P) is called a **probability measure space**, or simply a **probability space**

¹⁷A real-valued set function is a function that maps sets to real numbers.

¹⁸Remark: The axioms do not determine probabilities; the probabilities are assigned based on our knowledge of the system under study. (For example, one approach is to base probability assignments on the simple concept of equally likely outcomes.) The axioms enable us to easily calculate the probabilities of some events from knowledge of the probabilities of other events.

5.2. $P(\emptyset) = 0$.

5.3. Finite additivity¹⁹: If A_1, \dots, A_n are disjoint events, then

$$P\left(\bigcup_{i=1}^n A_i\right) = \sum_{i=1}^n P(A_i).$$

Special case when $n = 2$: **Addition rule** (Additivity)

$$\text{If } A \cap B = \emptyset, \text{ then } P(A \cup B) = P(A) + P(B). \quad (5)$$

5.4. The probability of a finite or countable event equals the sum of the probabilities of the outcomes in the event.

(a) In particular, if A is countable, e.g. $A = \{a_1, a_2, \dots\}$, then

$$P(A) = \sum_{n=1}^{\infty} P(\{a_n\}).$$

(b) Similarly, if A is finite, e.g. $A = \{a_1, a_2, \dots, a_{|A|}\}$, then

$$P(A) = \sum_{n=1}^{|A|} P(\{a_n\}).$$

¹⁹It is not possible to go backwards and use finite additivity to derive countable additivity (P3).

Example 5.5. A random experiment can result in one of the outcomes $\{a, b, c, d\}$ with probabilities 0.1, 0.3, 0.5, and 0.1, respectively. Let A denote the event $\{a, b\}$, B the event $\{b, c, d\}$, and C the event $\{d\}$.

5.6. Monotonicity: If $A \subset B$, then $P(A) \leq P(B)$

Example 5.7. Let A be the event to roll a 6 and B the event to roll an even number. Whenever A occurs, B must also occur. However, B can occur without A occurring if you roll 2 or 4.

5.8. If $A \subset B$, then $P(B \setminus A) = P(B) - P(A)$

5.9. $P(A) \in [0, 1]$.

5.10. $P(A \cap B)$ can not exceed $P(A)$ and $P(B)$. In other words, “the composition of two events is always less probable than (or at most equally probable to) each individual event.”

Example 5.11 (Slides). Experiments by psychologists Kahneman and Tversky.

Example 5.12. Let us consider Mrs. Boudreaux and Mrs. Thibodeaux who are chatting over their fence when the new neighbor walks by. He is a man in his sixties with shabby clothes and a distinct smell of cheap whiskey. Mrs.B, who has seen him before, tells Mrs. T that he is a former Louisiana state senator. Mrs. T finds this very hard to believe. “Yes,” says Mrs.B, “he is a former state senator who got into a scandal long ago, had to resign, and started drinking.” “Oh,” says Mrs. T, “that sounds more likely.” “No,” says Mrs. B, “I think you mean less likely.”

Strictly speaking, Mrs. B is right. Consider the following two statements about the shabby man: “He is a former state senator” and “He is a former state senator who got into a scandal long ago, had to resign, and started drinking.” It is tempting to think that the second is more likely because it gives a more exhaustive explanation of the situation at hand. However, this reason is precisely why it is a less likely statement. Note that whenever somebody satisfies the second description, he must also satisfy the first but not vice versa. Thus, the second statement has a lower probability (from Mrs. T's subjective point of view; Mrs. B of course knows who the man is).

This example is a variant of examples presented in the book *Judgment under Uncertainty* [11] by Economics Nobel laureate Daniel Kahneman and co-authors Paul Slovic and Amos Tversky. They show empirically how people often make similar mistakes when they are asked to choose the most probable among a set of statements. It certainly helps to know the rules of probability. A more discomfoting aspect is that the more you explain something in detail, the more likely you are to be wrong. If you want to be credible, be vague. [19, p 11–12]

5.13. Complement Rule:

$$P(A^c) = 1 - P(A).$$

- “The probability that something does not occur can be computed as one minus the probability that it does occur.”
- Named “probability’s Trick Number One” in [10]

5.14. Probability of a union (not necessarily disjoint):

$$P(A \cup B) = P(A) + P(B) - P(A \cap B)$$

- $P(A \cup B) \leq P(A) + P(B)$.
- Approximation: If $P(A) \gg P(B)$ then we may approximate $P(A \cup B)$ by $P(A)$.

Example 5.15 (Slides). Combining error probabilities from various sources in DNA testing

Example 5.16. In his bestseller *Innumeracy*, John Allen Paulos tells the story of how he once heard a local weatherman claim that there was a 50% chance of rain on Saturday and a 50% chance of rain on Sunday and thus a 100% chance of rain during the weekend. Clearly absurd, but what is the error?

Answer: Faulty use of the addition rule (5)!

If we let A denote the event that it rains on Saturday and B the event that it rains on Sunday, in order to use $P(A \cup B) = P(A) + P(B)$, we must first confirm that A and B cannot occur at

the same time ($P(A \cap B) = 0$). More generally, the formula that is always holds regardless of whether $P(A \cap B) = 0$ is given by 5.14:

$$P(A \cup B) = P(A) + P(B) - P(A \cap B).$$

The event “ $A \cap B$ ” describes the case in which it rains both days. To get the probability of rain over the weekend, we now add 50% and 50%, which gives 100%, but we must then subtract the probability that it rains both days. Whatever this is, it is certainly more than 0 so we end up with something less than 100%, just like common sense tells us that we should.

You may wonder what the weatherman would have said if the chances of rain had been 75% each day. [19, p 12]

5.17. Probability of a union of three events:

5.18. Two bounds:

(a) **Subadditivity or Boole’s Inequality:** If A_1, \dots, A_n are events, not necessarily disjoint, then

$$P\left(\bigcup_{i=1}^n A_i\right) \leq \sum_{i=1}^n P(A_i).$$

(b) **σ -subadditivity or countable subadditivity:** If A_1, A_2, \dots is a sequence of measurable sets, not necessarily disjoint, then

$$P\left(\bigcup_{i=1}^{\infty} A_i\right) \leq \sum_{i=1}^{\infty} P(A_i)$$

- This formula is known as the *union bound* in engineering.

5.19. If a (finite) collection $\{B_1, B_2, \dots, B_n\}$ is a partition of Ω , then

$$P(A) = \sum_{i=1}^n P(A \cap B_i)$$

Similarly, if a (countable) collection $\{B_1, B_2, \dots\}$ is a partition of Ω , then

$$P(A) = \sum_{i=1}^{\infty} P(A \cap B_i)$$

5.20. Connection to classical probability theory: Consider an experiment with **finite** sample space $\Omega = \{\omega_1, \omega_2, \dots, \omega_n\}$ in which each outcome ω_i is **equally likely**. Note that $n = |\Omega|$.

We must have

$$P(\{\omega_i\}) = \frac{1}{n}, \quad \forall i.$$

Now, given any event finite²⁰ event A , we can apply 5.4 to get

$$P(A) = \sum_{\omega \in A} P(\{\omega\}) = \sum_{\omega \in A} \frac{1}{n} = \frac{|A|}{n} = \frac{|A|}{|\Omega|}.$$

We can then say that the probability theory we are working on right now is an extension of the classical probability theory. When the conditions/assumptions of classical probability theory are met, then we get back the defining definition of classical probability. The extended part gives us ways to deal with situations where assumptions of classical probability theory are not satisfied.

²⁰In classical probability, the sample space is finite; therefore, any event is also finite.

6 Event-based Independence and Conditional Probability

Example 6.1 (Slides). Diagnostic Tests.

6.1 Event-based Conditional Probability

Definition 6.2. *Conditional Probability*: The conditional probability $P(A|B)$ of event A , given that event $B \neq \emptyset$ occurred, is given by

$$P(A|B) = \frac{P(A \cap B)}{P(B)}. \quad (6)$$

- Some ways to say or express the conditional probability, $P(A|B)$, are:
 - the “probability of A , given B ”
 - the “probability of A , knowing B ”
 - the “probability of A happening, knowing B has already occurred”
- Defined only when $P(B) > 0$.
 - If $P(B) = 0$, then it is illogical to speak of $P(A|B)$; that is $P(A|B)$ is not defined.

6.3. *Interpretation*: Sometimes, we refer to $P(A)$ as

- a **a priori probability** , or
- the **prior probability** of A , or
- the **unconditional probability** of A .

It is sometimes useful to interpret $P(A)$ as our knowledge of the occurrence of event A *before* the experiment takes place. Conditional probability $P(A|B)$ is the **updated probability** of the event A given that we now know that B occurred (but we still do not know which particular outcome in the set B occurred).

The term **a posteriori** is often used for $P(A|B)$ when we refer to $P(A)$ as a priori.

Example 6.4 (Slide). Roll a fair dice

Example 6.5. In diagnostic tests Example 6.1, we learn whether we have the disease from test result. Originally, before taking the test, the probability of having the disease is 0.01%. Being tested positive from the 99%-accurate test *updates* the probability of having the disease to about 1%.

More specifically, let D be the event that the testee has the disease and T_P be the event that the test returns positive result.

- Before taking the test, the probability of having the disease is $P(D) = 0.01\%$.
- Using 99%-accurate test means

$$P(T_P|D) = 0.99 \text{ and } P(T_P^c|D^c) = 0.99.$$

- Our calculation shows that $P(D|T_P) \approx 0.01$.

Note also that although the symbol $P(A|B)$ itself is practical, it phrasing in words can be so unwieldy that in practice, less formal descriptions are used. For example, we refer to “the probability that a tested-positive person has the disease” instead of saying “the conditional probability that a randomly chosen person has the disease given that the test for this person returns positive result.”

6.6. If the occurrence of B does not give you more information about A , then

$$P(A|B) = P(A) \tag{7}$$

and we say that A and B are *independent*.

- Meaning: “learning that event B has occurred does not change the probability that event A occurs.”
- Interpretation: “the occurrence of event A is not contingent on the occurrence (or nonoccurrence) of event B .”

We will soon define “independence”. Property (7) can be regarded as a “practical” definition for independence. However, there are some “technical” issues that we need to deal with when we actually define independence.

6.7. Similar properties to the three probability axioms:

- (a) Nonnegativity: $P(A|B) \geq 0$
- (b) Unit normalization: $P(\Omega|B) = 1$.

In fact, for any event A such that $B \subset A$, we have $P(A|B) = 1$.

This implies

$$P(\Omega|B) = P(B|B) = 1.$$

- (c) Countable additivity: For every countable sequence $(A_n)_{n=1}^{\infty}$ of disjoint events,

$$P\left(\bigcup_{n=1}^{\infty} A_n \mid B\right) = \sum_{n=1}^{\infty} P(A_n|B).$$

- In particular, if $A_1 \perp A_2$,

$$P(A_1 \cup A_2 | B) = P(A_1 | B) + P(A_2 | B)$$

6.8. Properties:

- $P(A|\Omega) = P(A)$
- $P(A^c|B) = 1 - P(A|B)$
- $P(A \cap B|B) = P(A|B)$
- $P(A_1 \cup A_2|B) = P(A_1|B) + P(A_2|B) - P(A_1 \cap A_2|B)$.
- $P(A \cap B) \leq P(A|B)$

6.9. When Ω is finite and all outcomes have equal probabilities,

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{|A \cap B| / |\Omega|}{|B| / |\Omega|} = \frac{|A \cap B|}{|B|}.$$

This formula can be regarded as the classical version of conditional probability.

Example 6.10. Someone has rolled a fair dice twice. You know that one of the rolls turned up a face value of six. The probability that the other roll turned up a six as well is $\frac{1}{11}$ (not $\frac{1}{6}$). [22, Example 8.1, p. 244]

6.11. Probability of compound events

- (a) $P(A \cap B) = P(A)P(B|A)$
- (b) $P(A \cap B \cap C) = P(A \cap B) \times P(C|A \cap B)$
- (c) $P(A \cap B \cap C) = P(A) \times P(B|A) \times P(C|A \cap B)$

When we have many sets intersected in the conditioned part, we often use “,” instead of “ \cap ”.

Example 6.12. Most people reason as follows to find the probability of getting two aces when two cards are selected at random from an ordinary deck of cards:

- (a) The probability of getting an ace on the first card is $4/52$.
- (b) Given that one ace is gone from the deck, the probability of getting an ace on the second card is $3/51$.

(c) The desired probability is therefore

$$\frac{4}{52} \times \frac{3}{51}.$$

[22, p 243]

Example 6.13. You know that roughly 5% of all used cars have been flood-damaged and estimate that 80% of such cars will later develop serious engine problems, whereas only 10% of used cars that are not flood-damaged develop the same problems. Of course, no used car dealer worth his salt would let you know whether your car has been flood damaged, so you must resort to probability calculations. What is the probability that your car will later run into trouble?

You might think about this problem in terms of proportions.

If you solved the problem in this way, congratulations. You have just used the law of total probability.

6.14. Total Probability Theorem: If a (finite or infinitely) countable collection of events $\{B_1, B_2, \dots\}$ is a partition of Ω , then

$$P(A) = \sum_i P(A|B_i)P(B_i). \quad (8)$$

This is a formula for computing the probability of an event that can occur in different ways.

6.15. Special case:

$$P(A) = P(A|B)P(B) + P(A|B^c)P(B^c).$$

This gives exactly the same calculation as what we discussed in Example 6.13.

Example 6.16. Continue from the “Diagnostic Tests” Example 6.1 and Example 6.5.

$$\begin{aligned} P(T_P) &= P(T_P \cap D) + P(T_P \cap D^c) \\ &= P(T_P | D) P(D) + P(T_P | D^c) P(D^c). \end{aligned}$$

For conciseness, we define

$$p_d = P(D)$$

and

$$p_{TE} = P(T_P | D^c) = P(T_P^c | D).$$

Then,

$$P(T_P) = (1 - p_{TE})p_D + p_{TE}(1 - p_D).$$

6.17. Bayes’ Theorem:

(a) Form 1:

$$P(B|A) = P(A|B) \frac{P(B)}{P(A)}.$$

(b) Form 2: If a (finite or infinitely) countable collection of events $\{B_1, B_2, \dots\}$ is a partition of Ω , then

$$P(B_k|A) = P(A|B_k) \frac{P(B_k)}{P(A)} = \frac{P(A|B_k)P(B_k)}{\sum_i P(A|B_i)P(B_i)}.$$

- Extremely useful for making inferences about phenomena that cannot be observed directly.
- Sometimes, these inferences are described as “reasoning about causes when we observe effects”.

Example 6.18. Continue from the “Disease Testing” Examples 6.1, 6.5, and 6.16:

$$\begin{aligned} P(D | T_P) &= \frac{P(D \cap T_P)}{P(T_P)} = \frac{P(T_P | D) P(D)}{P(T_P)} \\ &= \frac{(1 - p_{TE})p_D}{(1 - p_{TE})p_D + p_{TE}(1 - p_D)} \end{aligned}$$

Example 6.19. Medical Diagnostic: Because a new medical procedure has been shown to be effective in the early detection of an illness, a medical screening of the population is proposed. The probability that the test correctly identifies someone with the illness as positive is 0.99, and the probability that the test correctly identifies someone without the illness as negative is 0.95. The incidence of the illness in the general population is 0.0001. You take the test, and the result is positive. What is the probability that you have the illness? [15, Ex. 2-37]

Let D denote the event that you have the illness, and let T_P denote the event that the test signals positive.

Example 6.20. Bayesian networks are used on the Web sites of high-technology manufacturers to allow customers to quickly diagnose problems with products. An oversimplified example is presented here.

A printer manufacturer obtained the following probabilities from a database of test results. Printer failures are associated with three types of problems: hardware, software, and other (such as connectors), with probabilities 0.1, 0.6, and 0.3, respectively. The probability of a printer failure given a hardware problem is 0.9, given a software problem is 0.2, and given any other type of problem is 0.5. If a customer enters the manufacturers Web site to diagnose a printer failure, what is the most likely cause of the problem?

Let the events H , S , and O denote a hardware, software, or other problem, respectively, and let F denote a printer failure.

Example 6.21. Continue from Example 6.4

Example 6.22. In the early 1990s, a leading Swedish tabloid tried to create an uproar with the headline “Your ticket is thrown away!”. This was in reference to the popular Swedish TV show “Bingolotto” where people bought lottery tickets and mailed them to the show. The host then, in live broadcast, drew one ticket from a large mailbag and announced a winner. Some observant reporter noticed that the bag contained only a small fraction of the hundreds of thousands tickets that were mailed. Thus the conclusion: Your ticket has most likely been thrown away!

Let us solve this quickly. Just to have some numbers, let us say that there are a total of $N = 100,000$ tickets and that $n = 1,000$ of them are chosen at random to be in the final drawing. If the drawing was from all tickets, your chance to win would be $1/N = 1/100,000$. The way it is actually done, you need to both survive the first drawing to get your ticket into the bag and then get your ticket drawn from the bag. The probability to get your entry into the bag is $n/N = 1,000/100,000$. The conditional probability to be drawn from the bag, given that your entry is in it, is $1/n = 1/1,000$. Multiply to get $1/N = 1/100,000$ once more. There were no riots in the streets. [19, p 22]

6.23. Chain rule of conditional probability [9, p 58]:

$$P(A \cap B|C) = P(B|C)P(A|B \cap C).$$

Example 6.24. Your teacher tells the class there will be a surprise exam next week. On one day, Monday-Friday, you will be told in the morning that an exam is to be given on that day. You quickly realize that the exam will not be given on Friday; if it was, it would not be a surprise because it is the last possible day to get the exam. Thus, Friday is ruled out, which leaves Monday-Thursday. But then Thursday is impossible also, now having become the last possible day to get the exam. Thursday is ruled out, but then Wednesday becomes impossible, then Tuesday, then Monday, and you conclude: There is no such thing as a surprise exam! But the teacher decides to give the exam on Tuesday, and come Tuesday morning, you are surprised indeed.

This problem, which is often also formulated in terms of surprise fire drills or surprise executions, is known by many names, for example, the “hangman’s paradox” or by serious philosophers as the “prediction paradox.” To resolve it, let’s treat it as a probability problem. Suppose that the day of the exam is chosen randomly among the five days of the week. Now start a new school week. What is the probability that you get the test on Monday? Obviously $1/5$ because this is the probability that Monday is chosen. If the test was not given on Monday, what is the probability that it is given on Tuesday? The probability that Tuesday is chosen to start with is $1/5$, but we are now asking for the conditional probability that the test is given on Tuesday, given that it was not given on Monday. As there are now four days left, this conditional probability is $1/4$. Similarly, the conditional probabilities that the test is given on Wednesday, Thursday, and Friday conditioned on that it has not been given thus far are $1/3$, $1/2$, and 1 , respectively.

We could define the “surprise index” each day as the probability that the test is not given. On Monday, the surprise index is therefore 0.8, on Tuesday it has gone down to 0.75, and it continues to go down as the week proceeds with no test given. On Friday, the surprise index is 0, indicating absolute certainty that the test will be given that day. Thus, it is possible to give a surprise test but not in a way so that you are equally surprised each day, and it is never possible to give it so that you are surprised on Friday. [19, p 23–24]

Example 6.25. Today Bayesian analysis is widely employed throughout science and industry. For instance, models employed to determine car insurance rates include a mathematical function describing, per unit of driving time, your personal probability of having zero, one, or more accidents. Consider, for our purposes, a simplified model that places everyone in one of two categories: high risk, which includes drivers who average at least one accident each year, and low risk, which includes drivers who average less than one.

If, when you apply for insurance, you have a driving record that stretches back twenty years without an accident or one that

goes back twenty years with thirty-seven accidents, the insurance company can be pretty sure which category to place you in. But if you are a new driver, should you be classified as low risk (a kid who obeys the speed limit and volunteers to be the designated driver) or high risk (a kid who races down Main Street swigging from a half-empty \$2 bottle of Boone's Farm apple wine)?

Since the company has no data on you, it might assign you an equal prior probability of being in either group, or it might use what it knows about the general population of new drivers and start you off by guessing that the chances you are a high risk are, say, 1 in 3. In that case the company would model you as a hybrid—one-third high risk and two-thirds low risk—and charge you one-third the price it charges high-risk drivers plus two-thirds the price it charges low-risk drivers.

Then, after a year of observation, the company can employ the new datum to reevaluate its model, adjust the one-third and two-third proportions it previously assigned, and recalculate what it ought to charge. If you have had no accidents, the proportion of low risk and low price it assigns you will increase; if you have had two accidents, it will decrease. The precise size of the adjustment is given by Bayes's theory. In the same manner the insurance company can periodically adjust its assessments in later years to reflect the fact that you were accident-free or that you twice had an accident while driving the wrong way down a one-way street, holding a cell phone with your left hand and a doughnut with your right. That is why insurance companies can give out "good driver" discounts: the absence of accidents elevates the posterior probability that a driver belongs in a low-risk group. [14, p 111-112]

6.2 Event-based Independence

Plenty of random things happen in the world all the time, most of which have nothing to do with one another. If you toss a coin and I roll a dice, the probability that you get heads is $1/2$ regardless of the outcome of my dice. Events that in this way are unrelated to each other are called *independent*.

Definition 6.26. Two events A , B are called (statistically²¹) *independent* if

$$P(A \cap B) = P(A)P(B) \quad (9)$$

- Notation: $A \perp\!\!\!\perp B$
- Read “ A and B are independent” or “ A is independent of B ”
- We call (9) the *multiplication rule* for probabilities.
- If two events are not independent, they are *dependent*. Intuitively, if two events are dependent, the probability of one changes with the knowledge of whether the other has occurred.

6.27. Intuition: Again, here is how you should think about independent events: “If one event has occurred, the probability of the other does not change.”

$$P(A|B) = P(A) \quad \text{and} \quad P(B|A) = P(B). \quad (10)$$

In other words, “the unconditional and the conditional probabilities are the same”. We can almost use (10) as the definitions for independence. However, we use (9) instead because it (1) also works with events whose probabilities are zero and (2) also has clear symmetry in the expression (so that $A \perp\!\!\!\perp B$ and $B \perp\!\!\!\perp A$ can clearly be seen as the same). In fact, in 6.31, we show how (10) can be used to define independence with extra condition that deals with the case when zero probability is involved.

²¹Sometimes our definition for independence above does not agree with the everyday-language use of the word “independence”. Hence, many authors use the term “statistically independence” to distinguish it from other definitions.

Example 6.28 (Slides). [26, Ex. 5.4] Which of the following pairs of events are independent?

(a) The card is a club, and the card is black.

(b) The card is a king, and the card is black.

6.29. An event with probability 0 or 1 is independent of any event (including itself).

- In particular, \emptyset and Ω are independent of any events.

6.30. An event A is independent of itself if and only if $P(A)$ is 0 or 1.

6.31. Two events A, B with positive probabilities are independent if and only if $P(B|A) = P(B)$, which is equivalent to $P(A|B) = P(A)$.

When A and/or B has zero probability, A and B are automatically independent.

6.32. When A and B have nonzero probabilities, the following statements are equivalent:

6.33. The following four statements are equivalent:

$$A \perp\!\!\!\perp B, \quad A \perp\!\!\!\perp B^c, \quad A^c \perp\!\!\!\perp B, \quad A^c \perp\!\!\!\perp B^c.$$

Example 6.34. If $P(A|B) = 0.4$, $P(B) = 0.8$, and $P(A) = 0.5$, are the events A and B independent? [15]

6.35. Keep in mind that **independent and disjoint** are *not synonyms*. In some contexts these words can have similar meanings, but this is not the case in probability.

- If two events cannot occur at the same time (they are disjoint), are they independent? At first you might think so. After all, they have nothing to do with each other, right? Wrong! They have a lot to do with each other. If one has occurred, we know for certain that the other cannot occur. [19, p 12]
- To check whether A and B are disjoint, you only need to look at the sets themselves and see whether they have shared element(s). This can be answered without knowing probabilities.

To check whether A and B are independent, you need to look at the probabilities $P(A)$, $P(B)$, and $P(A \cap B)$.

- Reminder: If events A and B are disjoint, you calculate the probability of the union $A \cup B$ by adding the probabilities of A and B . For independent events A and B you calculate the probability of the intersection $A \cap B$ by multiplying the probabilities of A and B .

- The two statements $A \perp B$ and $A \perp\!\!\!\perp B$ can occur simultaneously only when $P(A) = 0$ and/or $P(B) = 0$.

- Reverse is not true in general.

Example 6.36. Experiment of flipping a fair coin twice. $\Omega = \{HH, HT, TH, TT\}$. Define event A to be the event that the first flip gives a H; that is $A = \{HH, HT\}$. Event B is the event that the second flip gives a H; that is $B = \{HH, TH\}$. Note that even though the events A and B are not disjoint, they are independent.

Example 6.37 (Slides). *Prosecutor’s fallacy*: In 1999, a British jury convicted Sally Clark of murdering two of her children who had died suddenly at the ages of 11 and 8 weeks, respectively. A pediatrician called in as an expert witness claimed that the chance of having two cases of infant sudden death syndrome, or “cot deaths,” in the same family was 1 in 73 million. There was no physical or other evidence of murder, nor was there a motive. Most likely, the jury was so impressed with the seemingly astronomical odds against the incidents that they convicted. But where did the number come from? Data suggested that a baby born into a family similar to the Clarks faced a 1 in 8,500 chance of dying a cot death. Two cot deaths in the same family, it was argued, therefore had a probability of $(1/8,500)^2$ which is roughly equal to 1/73,000,000.

Did you spot the error? I hope you did. The computation assumes that successive cot deaths in the same family are *independent* events. This assumption is clearly questionable, and even a person without any medical expertise might suspect that genetic factors play a role. Indeed, it has been estimated that if there is one cot death, the next child faces a much larger risk, perhaps

around $1/100$. To find the probability of having two cot deaths in the same family, we should thus use conditional probabilities and arrive at the computation $1/8,500 \times 1/100$, which equals $1/850,000$. Now, this is still a small number and might not have made the jurors judge differently. But what does the probability $1/850,000$ have to do with Sallys guilt? Nothing! When her first child died, it was certified to have been from natural causes and there was no suspicion of foul play. The probability that it would happen again without foul play was $1/100$, and if that number had been presented to the jury, Sally would not have had to spend three years in jail before the verdict was finally overturned and the expert witness (certainly no expert in probability) found guilty of “serious professional misconduct.”

You may still ask the question what the probability $1/100$ has to do with Sallys guilt. Is this the probability that she is innocent? Not at all. That would mean that 99% of all mothers who experience two cot deaths are murderers! The number $1/100$ is simply the probability of a second cot death, which only means that among all families who experience one cot death, about 1% will suffer through another. If probability arguments are used in court cases, it is very important that all involved parties understand some basic probability. In Sallys case, nobody did.

References: [14, 118–119] and [19, 22–23].

Definition 6.38. Three events A_1, A_2, A_3 are independent if and only if

$$P(A_1 \cap A_2) = P(A_1)P(A_2)$$

$$P(A_1 \cap A_3) = P(A_1)P(A_3)$$

$$P(A_2 \cap A_3) = P(A_2)P(A_3)$$

$$P(A_1 \cap A_2 \cap A_3) = P(A_1)P(A_2)P(A_3)$$

Remarks:

- (a) When the first three equations hold, we say that the three events are ***pairwise independent***.
- (b) We may use the term “mutually independence” to further emphasize that we have “independence” instead of “pairwise independence”.

Definition 6.39. The events A_1, A_2, \dots, A_n are *independent* if and only if for any subcollection $A_{i_1}, A_{i_2}, \dots, A_{i_k}$,

$$P(A_{i_1} \cap A_{i_2} \cap \dots \cap A_{i_k}) = P(A_{i_1}) \times P(A_{i_2}) \times \dots \times P(A_{i_k}).$$

- Note that part of the requirement is that

$$P(A_1 \cap A_2 \cap \dots \cap A_n) = P(A_1) \times P(A_2) \times \dots \times P(A_n).$$

Therefore, if someone tells us that the events A_1, A_2, \dots, A_n are independent, then one of the properties that we can conclude is that

$$P(A_1 \cap A_2 \cap \dots \cap A_n) = P(A_1) \times P(A_2) \times \dots \times P(A_n).$$

- Equivalently, this is the same as the requirement that

$$P\left(\bigcap_{j \in J} A_j\right) = \prod_{j \in J} P(A_j) \quad \forall J \subset [n] \text{ and } |J| \geq 2$$

- Note that the case when $j = 1$ automatically holds. The case when $j = 0$ can be regarded as the \emptyset event case, which is also trivially true.

6.3 Bernoulli Trials

Example 6.40. Consider the following random experiments

- Flip a coin 10 times. We are interested in the number of heads obtained.
- Of all bits transmitted through a digital transmission channel, 10% are received in error. We are interested in the number of bits in error in the next five bits transmitted.
- A multiple-choice test contains 10 questions, each with four choices, and you guess at each question. We are interested in the number of questions answered correctly.

These examples illustrate that a general probability model that includes these experiments as particular cases would be very useful.

Example 6.41. Each of the random experiments in Example 6.40 can be thought of as consisting of a series of repeated, random trials. In all cases, we are interested in the number of trials that meet a specified criterion. The outcome from each trial either meets the criterion or it does not; consequently, each trial can be summarized as resulting in either a success or a failure.

Definition 6.42. A *Bernoulli trial* involves performing an experiment once and noting whether a particular event A occurs.

The outcome of the Bernoulli trial is said to be

- (a) a “success” if A occurs and
- (b) a “failure” otherwise.

We may view the outcome of a single Bernoulli trial as the outcome of a toss of an unfair coin for which the probability of heads (success) is $p = P(A)$ and the probability of tails (failure) is $1 - p$.

- The labeling (“success” and “failure”) is not meant to be literal and sometimes has nothing to do with the everyday meaning of the words. We can just as well use A and B or 0 and 1 .

Example 6.43. Examples of Bernoulli trials: Flipping a coin, deciding to vote for candidate A or candidate B , giving birth to a boy or girl, buying or not buying a product, being cured or not being cured, even dying or living are examples of Bernoulli trials.

- Actions that have multiple outcomes can also be modeled as Bernoulli trials if the question you are asking can be phrased in a way that has a yes or no answer, such as “Did the dice land on the number 4?” or “Is there any ice left on the North Pole?”

Definition 6.44. (Independent) *Bernoulli Trials* = a Bernoulli trial is repeated many times.

- (a) It is usually assumed that the trials are independent. This implies that the outcome from one trial has no effect on the outcome to be obtained from any other trial.

(b) Furthermore, it is often reasonable to assume that the probability of a success in each trial is constant.

An outcome of the complete experiment is a sequence of successes and failures which can be denoted by a *sequence of ones and zeroes*.

Example 6.45. If we toss *unfair coin* n times, we obtain the space $\Omega = \{H, T\}^n$ consisting of 2^n elements of the form $(\omega_1, \omega_2, \dots, \omega_n)$ where $\omega_i = H$ or T .

Example 6.46. What is the probability of two failures and three successes in five Bernoulli trials with success probability p .

We observe that the outcomes with three successes in five trials are 11100, 11010, 11001, 10110, 10101, 10011, 01110, 01101, 01011, and 00111. We note that the probability of each outcome is a product of five probabilities, each related to one Bernoulli trial. In outcomes with three successes, three of the probabilities are p and the other two are $1 - p$. Therefore, each outcome with three successes has probability $(1 - p)^2 p^3$. There are 10 of them. Hence, the total probability is $10(1 - p)^2 p^3$

6.47. The probability of exactly n_1 success in $n = n_0 + n_1$ bernoulli trials is

$$\binom{n}{n_1} (1 - p)^{n-n_1} p^{n_1} = \binom{n}{n_0} (1 - p)^{n_0} p^{n-n_0}.$$

Example 6.48. At least one occurrence of a 1-in- n -chance event in n repeated trials:

Example 6.49. *Digital communication over unreliable channels*: Consider a communication system below

Here, we consider a simple channel called **binary symmetric channel**:

This channel can be described as a channel that introduces random bit errors with probability p .

A crude digital communication system would put binary information into the channel directly; the receiver then takes whatever value that shows up at the channel output as what the sender transmitted. Such communication system would directly suffer bit error probability of p .

In situation where this error rate is not acceptable, error control techniques are introduced to reduce the error rate in the delivered information.

One method of reducing the error rate is to use error-correcting codes:

A simple error-correcting code is the *repetition code* described below:

- (a) At the transmitter, the “encoder” box performs the following task:

- (i) To send a 1, it will send 11111 through the channel.
 - (ii) To send a 0, it will send 00000 through the channel.
- (b) When the five bits pass through the channel, it may be corrupted. Assume that the channel is binary symmetric and that it acts on each of the bit independently.
- (c) At the receiver, we (or more specifically, the decoder box) get 5 bits, but some of the bits may be changed by the channel. To determine what was sent from the transmitter, the receiver apply the *majority rule*: Among the 5 received bits,
- (i) if $\#1 > \#0$, then it claims that “1” we transmitted,
 - (ii) if $\#0 > \#1$, then it claims that “0” we transmitted.

Exercise 6.50 (F2011). Kakashi and Gai are eternal rivals. Kakashi is a little stronger than Gai and hence for each time that they fight, the probability that Kakashi wins is 0.55. In a competition, they fight n times (where n is odd). We will assume that the results of the fights are independent. The one who wins more will win the competition.

Suppose $n = 3$, what is the probability that Kakashi wins the competition.