

RAFAEL PEÑALOZA

PROBABILITY THEORY AND STATISTICS

FREE UNIVERSITY OF BOZEN-BOLZANO

Disclaimer

THIS MANUSCRIPT CONTAINS the lecture notes for the course on Probability Theory and Statistics at the Free University of Bozen-Bolzano. Most of the content is based on the Sheldon M. Ross's book.¹ There are currently five editions of that book; any of them should suffice for the needs of the course.²

This manuscript **does not** replace the material presented on the blackboard during the course. It may (in fact, it does) contain errors, typos, and miss relevant details discussed at the classroom.³ Use it with care. The best advice is for you to attend the lecture and take notes by yourself, and use this text as supporting material.

Background information and other supplemental material like handouts, exercises, and code are provided at the OLE for the course. If you have not yet registered, **do it now**.

¹

² The third edition is available for free at this location.

³ Needless to say, all errors found here are my fault, and my fault only. If you notice any mistake, **please let me know** so that I can fix it.

Contents

Introduction to Probability Theory 5

Random Variables 17

Special Random Variables 34

Sampling 45

Parameter Estimation 50

Hypothesis Testing 59

Regression 67

Introduction to Probability Theory

PROBABILITY THEORY IS a means for measuring and speaking about the uncertainty about occurrences of events. The event will happen or not, but we do not know which will be the case. Suppose for example that I roll a die. The die will necessarily show one face, but we do not know which one; the event “the die will show a 6” is uncertain.

Probabilities *quantify* this uncertainty. In this case, we would say something like “the probability of the die showing a 6 is $1/6$,” since there are six possible outcomes that we consider equally likely. There are two ways to understand this statement:

- if we repeat the experiment a large enough number of times, then in $1/6$ of them the outcome will be a 6;
- we believe that there is a 1 in 6 chance of this specific roll landing on 6.

The first is the *frequentist* view. This view suggests that a probability is intrinsic to the event, and it can be studied and determined by repeating an experiment. The second is the *subjective* (or *Bayesian*) view, where the probability refers to the belief of the agent stating the probability. Although these two views have some deep philosophical differences, they do not affect the study of the mathematical properties of probabilities. So we often disregard these differences when studying the theory of probabilities formally.⁴

Events

WE WILL DEAL with *experiments* whose outcome is uncertain, but where the set of all possible outcomes is known. This set of all possible outcomes is called the *sample space*, denoted by \mathcal{S} . When rolling a die, the sample space is $\mathcal{S} = \{1, \dots, 6\}$. Sample spaces may be very simple, or very complex. For example, in a race the sample space may describe all the possible orderings in which the participants can finish. Indeed, \mathcal{S} may be infinite.

Any subset of the state space is called an *event*. If the outcome of the experiment is contained in a given event \mathcal{E} , then we say that \mathcal{E} *occurred*. For example, the event of the die rolling an even number is $\mathcal{E} = \{2, 4, 6\}$.⁵

Since events are sets, we can define new events through set operations over other events. Given two events \mathcal{E} and \mathcal{F} , their *union* ($\mathcal{E} \cup \mathcal{F}$), *intersection* ($\mathcal{E} \mathcal{F}$), and *complement* ($\bar{\mathcal{E}}$) are also events.⁶ In our dice example,

⁴ An accessible discussion on the Bayesian view, and its importance in modern statistics is available here.

⁵ Obviously, more complex events are possible for other spaces.

⁶ For a brief introduction to set theory and Venn diagrams, see the relevant handout.

if $\mathcal{E} = \{2, 4, 6\}$ is the event of rolling an even number, and $\mathcal{F} = \{2, 3, 5\}$ refers to the event of rolling a prime number, then

- $\mathcal{E} \cup \mathcal{F} = \{2, 3, 4, 5, 6\}$: rolling an even or a prime number;
- $\mathcal{E} \mathcal{F} = \{2\}$: rolling a prime even number; and
- $\overline{\mathcal{E}} = \{1, 3, 5\}$: rolling an odd number.

It is possible that an event is empty, for example by taking the intersection of two events that do not share any outcomes like $\{2\}$ and $\{3\}$. This event is denoted by \emptyset ; e.g. $\{2\}\{3\} = \emptyset$. If $\mathcal{E}\mathcal{F} = \emptyset$, then \mathcal{E} and \mathcal{F} are *mutually exclusive*. Notice that the whole sample space is also an event, and that $\overline{\mathcal{F}} = \emptyset$.⁷ Just as with sets, we can consider inclusions between events, too. \mathcal{E} is *contained* in \mathcal{F} if all outcomes in \mathcal{E} are also in \mathcal{F} . In this case, we write $\mathcal{E} \subseteq \mathcal{F}$ or $\mathcal{F} \supseteq \mathcal{E}$. If $\mathcal{E} \subseteq \mathcal{F}$ and $\mathcal{F} \subseteq \mathcal{E}$, then \mathcal{E} and \mathcal{F} are *equivalent* ($\mathcal{E} \equiv \mathcal{F}$). For example, rolling a two is contained in rolling an even number, and is equivalent to rolling a prime even number.

⁷ In terms of language, we can see the union as a (non-exclusive) *or*, the intersection as an *and*, and the complement as a *no*.

Axioms of Probability

RECALL THAT THE FREQUENTIST view states that if we repeat an experiment repeatedly under the exact same conditions, then the proportion of times in which the outcome belongs to a given event \mathcal{E} will converge to a constant, which reflects the probability of \mathcal{E} .

From a mathematical point of view, for every event \mathcal{E} over a sample space \mathcal{S} there is a number, called the *probability of \mathcal{E}* and denoted as $P(\mathcal{E})$ that satisfies the following three axioms:

Axiom 1 $0 \leq P(\mathcal{E}) \leq 1$;

Axiom 2 $P(\mathcal{S}) = 1$; and

Axiom 3 For any sequence of mutually exclusive events $\mathcal{E}_1, \mathcal{E}_2, \dots$,⁸ and for any $n \in \mathbb{N}$ it holds that $P(\bigcup_{i=1}^n \mathcal{E}_i) = \sum_{i=1}^n P(\mathcal{E}_i)$.

⁸ That is, events such that $\mathcal{E}_i \mathcal{E}_j = \emptyset$ whenever $i \neq j$.

The first axiom says that the probability of any event is always a number between 0 and 1. Axiom 2 gives a completeness statement on probabilities: the probability of observing any outcome from the sample space is always 1. The last axiom shows how to aggregate the probabilities of mutually exclusive events: the probability of their union is always the sum of their individual probabilities.

The frequentist view of probabilities clearly satisfies these three axioms: for any event \mathcal{E} , the proportion (or frequency) of times in which a repeated experiment falls in an outcome from \mathcal{E} is necessarily between 0 and 1; in every repetition of the experiment the outcome belongs to the sample space \mathcal{S} ; and if two events \mathcal{E} and \mathcal{F} do not have any outcome in common, then the proportion of the time in which the outcome is in \mathcal{E} or in \mathcal{F} is the sum of their respective frequencies. As an example of this last axiom, suppose that \mathcal{E} is the event where a roll of a die falls in an even number, and \mathcal{F} that it falls on the number 1. Then, we expect \mathcal{E} to hold half (50 percent) of the time, and \mathcal{F} $\frac{1}{6}$ -th of the time. Then $\mathcal{E} \cup \mathcal{F}$ should be observed approximately 66% of the time.

These simple axioms allow us derive many properties of probabilities, and understand them formally. We start with two simple propositions that already showcase some interesting characteristics of the theory.

Proposition 1. For every event \mathcal{E} , $P(\overline{\mathcal{E}}) = 1 - P(\mathcal{E})$.

Proof. By definition, \mathcal{E} and $\overline{\mathcal{E}}$ are disjoint, and $\mathcal{S} = \mathcal{E} \cup \overline{\mathcal{E}}$. Using Axiom 3, it follows that $P(\mathcal{S}) = P(\mathcal{E} \cup \overline{\mathcal{E}}) = P(\mathcal{E}) + P(\overline{\mathcal{E}})$. By Axiom 2 it also follows that $P(\mathcal{S}) = 1$, and hence $P(\overline{\mathcal{E}}) = 1 - P(\mathcal{E})$. \square

In other words, the probability of an event *not* occurring is always one minus the probability of it occurring. For example, if the probability of rolling an even number in a die is 0.6, then the probability of rolling an odd number must be 0.4.

Notice that Axiom 3 provides a means for computing the probability of the union of two or more events, but only if they are mutually exclusive. We often need to compute such probabilities for events that are not necessarily disjoint. We can do that, as long as we know the probability of the intersection.

Proposition 2. For every two events \mathcal{E}, \mathcal{F} , $P(\mathcal{E} \cup \mathcal{F}) = P(\mathcal{E}) + P(\mathcal{F}) - P(\mathcal{E} \mathcal{F})$.

Proof. To help in this proof, consider the Venn diagram in Figure 1. The three regions I, II, and III are mutually exclusive. Hence, it follows that

$$P(\mathcal{E} \cup \mathcal{F}) = P(\text{I}) + P(\text{II}) + P(\text{III});$$

$$P(\mathcal{E}) = P(\text{I}) + P(\text{II});$$

$$P(\mathcal{F}) = P(\text{II}) + P(\text{III}).$$

Thus, $P(\mathcal{E} \cup \mathcal{F}) = P(\mathcal{E}) + P(\mathcal{F}) - P(\text{II})$. Since $\text{II} \equiv \mathcal{E} \mathcal{F}$, this finishes the proof. \square

Example 3. 23 percent of adults drink beer; 7 percent drink wine, and 2 percent drink both, beer and wine. What percentage of people drink neither beer nor wine?

Solution. To know how many people drink neither beer nor wine, it suffices to know how many people drink at least one of these beverages (the desired event is the complement of this one). Let \mathcal{E} be the event of a person drinking beer, and \mathcal{F} the event of drinking wine. The probability of drinking one of them is⁹

$$P(\mathcal{E} \cup \mathcal{F}) = P(\mathcal{E}) + P(\mathcal{F}) - P(\mathcal{E} \mathcal{F}) = 0.23 + 0.07 - 0.02 = 0.28.$$

The probability of a person being abstemious is $1 - 0.28 = 0.72$. \triangle

ANOTHER IMPORTANT NOTION when speaking about uncertainty and probability is that of odds.

Definition 4. The *odds* of an event \mathcal{E} is defined by $\frac{P(\mathcal{E})}{P(\overline{\mathcal{E}})} = \frac{P(\mathcal{E})}{1 - P(\mathcal{E})}$.

Intuitively, the odds of an event \mathcal{E} tells us how much more likely it is that \mathcal{E} occurs than that it does not occur. If the odds of \mathcal{E} is greater than 1, then it is more likely to see \mathcal{E} than not; and *vice versa* if it is smaller than 1.




Figure 1: Three relevant regions in a Venn diagram.

⁹ To see this in a different way, think: how many people drink *only* beer? These are those that drink beer, minus those that drink beer and wine: $P(\mathcal{E}) - P(\mathcal{E} \mathcal{F})$. Those that drink *only* wine are $P(\mathcal{F}) - P(\mathcal{E} \mathcal{F})$. So, the people that drink beer or wine are those that drink only beer, plus those that drink only wine, plus those that drink both: $P(\mathcal{E}) - P(\mathcal{E} \mathcal{F}) + P(\mathcal{F}) - P(\mathcal{E} \mathcal{F}) + P(\mathcal{E} \mathcal{F})$.

For example, if rolling a 2 in a die has probability $P(\mathcal{E}) = \frac{1}{6}$, then the odds of this event is $\frac{1}{6} / \frac{5}{6} = \frac{1}{5}$. Consequently, it is five times more likely not to see a 2 than it is to roll the 2. In this case, we say that “the odds are 1 to 5 against the event \mathcal{E} .”¹⁰

¹⁰ Or, dually, *in favour of* $\bar{\mathcal{E}}$.

Uniformity

THERE ARE MANY CASES where it is natural to assume that every point in the sample space \mathcal{S} is equally likely to occur. If the sample space is *finite*, (for example, if $\mathcal{S} = \{1, \dots, n\}$ for some $n \in \mathbb{N}$), then this means that

$$P(\{1\}) = \dots = P(\{n\}) = p.$$

Using Axioms 2 and 3, it follows that

$$1 = P(\mathcal{S}) = P(\{1\}) + \dots + P(\{n\}) = np,$$

and hence for every i , $P(\{i\}) = p = \frac{1}{n}$. That is, every outcome has the same probability of occurring determined by the total size of the sample space.

We can use Axiom 3 again to generalise this result to arbitrary events. For an event \mathcal{E} , let $\#\mathcal{E}$ denote the number of elements (outcomes) in \mathcal{E} . If all outcomes are equally likely to occur in a sample space with n elements,¹¹ then

$$P(\mathcal{E}) = \frac{\#\mathcal{E}}{n}.$$

That is, the probability of \mathcal{E} is the proportion of elements in the sample space that belong to \mathcal{E} .

For that reason, it is often important to be able to *count* the number of different ways in which an event may occur. One very helpful rule for counting is the basic principle of counting.

Principle (Basic Principle of Counting). If two experiments are performed, such that Experiment 1 can result in any of m different outcomes, and for each of them Experiment 2 may result in any of n outcomes, then together there are mn different possible outcomes for the two experiments.

To show that this is the case, it suffices to enumerate all possible outcomes of the two experiments, arranging them in a matrix where every row corresponds to the outcome of Experiment 1, and every column to the outcome of Experiment 2. This matrix has overall mn elements (see Figure 2).

Example 5. A drawer contains 8 black socks and 6 white socks. We “randomly” take two socks from the drawer. What is the probability of them *not* forming a pair?¹²

Solution. We first measure the size of the sample space: the possible outcomes of removing two socks from the drawer. There are 14 socks originally inside. First we can choose any of those 14 socks, and then any of the 13 remaining ones. Overall, there are $14 \cdot 13 = 182$ possible combinations. In order to *not* have a pair, we have to extract first a black one followed by a white one ($8 \cdot 6 = 48$ combinations) or first a white one followed by a black

Axioms of probability:

1. $0 \leq P(\mathcal{E}) \leq 1$
2. $P(\mathcal{S}) = 1$
3. $P(\bigcup_{i=1}^n \mathcal{E}_i) = \sum_{i=1}^n P(\mathcal{E}_i)$ if \mathcal{E}_i s are mutually exclusive.

¹¹ We often refer to this case as having a *uniform* sample space.

(1, 1)	(1, 2)	...	(1, n)
(2, 1)	(2, 2)	...	(2, n)
\vdots	\vdots	\ddots	\vdots
(m, 1)	(m, 2)	...	(m, n)

Figure 2: Possible outcomes for the Basic Principle of Counting

¹² That is, we want the probability of the two selected socks having different color.

one $6 \cdot 8 = 48$ combinations). Hence, assuming that each combination of the sample space is equally likely to occur,¹³ we see that the probability is $\frac{48+48}{182} = \frac{96}{182} \approx 0.53$.

Obviously, the basic principle of counting can be generalised to cases where more than two experiments are performed. If r experiments are performed, such that Experiment 1 can result in any of n_1 different outcomes, and for each of them Experiment 2 may result in any of n_2 outcomes, and so on, then together there are $n_1 \cdot n_2 \cdots n_r$ different possible outcomes for the r experiments.

Consider for example the question of counting the number of ways in which a collection of n different objects can be organised over a line. For a small enough n , we can find this number through enumeration: all the linear ordering of the letters a, b , and c are abc, acb, bac, bca, cab , and cba ; that is, there are 6 such orderings. We often call each of these orderings a *permutation*. There are six permutations on any set of three objects. Rather than enumerating them all, which becomes impossible (or at least tedious) as the number of objects grows, one can also count the number of permutations using the basic principle: the first object in the permutation may be any of 3 choices, the second any of the remaining 2 choices, and the last one is determined by the previous outcomes (that is, only one choice remains). In other words, there are $3 \cdot 2 \cdot 1 = 6$ permutations.

If instead of 3 we have n objects, then the same argument can be used to conclude that there are $n(n-1)(n-2) \cdots 3 \cdot 2 \cdot 1$ different permutations of them. This expression is known as n factorial and denoted by $n!$.¹⁴

Example 6. We want to organize in a bookshelf 10 books, which are divided in the following subjects: 4 are computer science, 3 are mathematics, 2 are statistics, and 1 is history. If we want all the books from the same subject to appear together, how many different arrangements are possible?

Solution. First we consider the number of possible different permutations of the subject matters: since there are 4 different subjects, we can choose any of $4!$ different orderings of them. Once that we have chosen the order of the subjects, we must choose the order of the books within the subject. The CS books allow $4!$ different ordering, the mathematics ones have $3!$ orders, statistics $2!$, and the history book has only $1 = 1!$ ordering. Thus, overall there are $4!4!3!2!1! = 6,912$ possible ways to arrange the bookshelf keeping books from the same subject together.

Example 7. A course has 5 male and 3 female students. After an exam, all students are ranked according to their performance. Assuming that the scores of the exams are all different, (a) how many different rankings are possible? (b) If all rankings are equally likely, what is the probability of women getting the 3 highest scores?

Solution. (a) There are 8 people taking the test, so there are $8! = 40320$ possible rankings. (b) There are $3!$ rankings between the female students, and $5!$ rankings among males; thus through the basic principle of counting, there are $5!3!$ rankings where the first ranked are females. The desired probability is then $\frac{5!3!}{8!} = \frac{3 \cdot 2 \cdot 1}{8 \cdot 7 \cdot 6} = 1/56$.

¹³ That is what we often mean when saying that an experiment is done “randomly.”

¹⁴ That is, $n! = n(n-1)(n-2) \cdots 3 \cdot 2 \cdot 1$. For the sake of completeness, we define $0! := 1$.

RATHER THAN ORDERING the objects in a collection with n elements, we sometimes want to count the number of groups with r objects that can be formed. For instance: how many groups of 3 elements can be formed with the five items A, B, C, D , and E ? One way to solve this problem is the following. We have 5 ways to select the first element, 4 to select the second, and 3 to select the third.[◦] Thus, there are $5 \cdot 4 \cdot 3 = 60$ ways of selecting the group of 3 objects when the order in which they are selected is relevant. Notice, however, that if we are only interested in the groups of objects, then this method counts each group 6 times.¹⁵ Thus, the total number of groups that can be formed (ignoring the ordering) is $\frac{5 \cdot 4 \cdot 3}{3 \cdot 2 \cdot 1} = 10$.

◦Do you understand why?

¹⁵ For example, if the group selected is A, B, C , then it will count as different the cases where we extract ABC , and where we extract CBA .

In general, there are $n(n-1) \cdots (n-r+1)$ different ways to obtain groups of r elements from a collection of n objects, if the order in which they are extracted is relevant. Since each group is counted $r!$ times through this method, when the order is not relevant there are

$$\frac{n(n-1) \cdots (n-r+1)}{r!} = \frac{n!}{(n-r)!r!}$$

groups of r elements from a collection of n objects.

Definition 8. Let $r \leq n$. The *combinations* of r objects in n is

$$\binom{n}{r} := \frac{n!}{(n-r)!r!}.$$

Again, $\binom{n}{r}$ is the number of groups of size r than can be extracted from a collection of n elements, when their ordering is not important; e.g., there are $\binom{7}{2} = \frac{7 \cdot 6}{2 \cdot 1} = 21$ different pairs that can be selected from a group of 7 people. Recall that $0! = 1$. Then $\binom{n}{0} = \binom{n}{n} = 1$.

Example 9. Five people are selected randomly from a group containing 5 men and 8 women. What is the probability that 3 women and 2 men are chosen?

Solution. By *random selection* we mean that each of the $\binom{13}{5}$ possible combinations is equally likely. There are $\binom{5}{2}$ possible combinations of 2 men, and $\binom{8}{3}$ of three women. From the basic principle of counting, the probability of this selection is then

$$\frac{\binom{5}{2}\binom{8}{3}}{\binom{13}{5}} = \frac{5 \cdot 4 \cdot 8 \cdot 7 \cdot 6 \cdot 5!}{13 \cdot 12 \cdot 11 \cdot 10 \cdot 9 \cdot 2 \cdot 3 \cdot 2} = \frac{8 \cdot 7 \cdot 5 \cdot 2}{13 \cdot 11 \cdot 9} = \frac{560}{1287}. \quad \triangle$$

Example 10. If we randomly select k elements from a collection of n objects, what is the probability that a given object is among the k selected?

Solution. There is one way of choosing the selected item, and $\binom{n-1}{k-1}$ ways of selecting $k-1$ elements from the remaining objects in the collection. From the basic principle of counting, there are $1 \cdot \binom{n-1}{k-1}$ different subsets of k of the n elements that include the selected one. Since the total possible choices are $\binom{n}{k}$, the probability that a particular object is among the k selected is[◦]

$$\frac{\binom{n-1}{k-1}}{\binom{n}{k}} = \frac{(n-1)!(n-k)!k!}{(n-k)!(k-1)!n!} = \frac{k}{n}. \quad \triangle$$

◦What is the probability that, in a group of 10 people, they all have a different birthday?

Conditional Probability

ONE OF THE FUNDAMENTAL notions in probability theory is that of *conditional probabilities*. Conditional probabilities are a way of measuring uncertainty when some additional information (or evidence) is available: it allows us to update our beliefs based on some previous observations from the world. However, they are also useful by themselves, as one can often simplify a problem of computing a probability by considering the condition of a secondary event occurring or not.

To illustrate conditional probabilities, suppose that we roll a pair of dice. The sample space of this experiment is the set containing 36 outcomes $\mathcal{S} = \{(i, j) \mid 1 \leq i, j \leq 6\}$, where (i, j) refers to the outcome where the first die lands on i and the second on j . If each outcome is equally likely to occur, each of them has probability $1/36$.¹⁶ Suppose that we observe that the first die landed on 3. What is the probability that the sum of the two dice is exactly 8?

¹⁶ We say that the dice are *fair*.

Since we already know that the first die landed on 3, there are 6 possible outcomes remaining for our experiment: $(3, 1), (3, 2), (3, 3), (3, 4), (3, 5)$, and $(3, 6)$. In addition, as originally each of these outcomes was equally likely, the same should hold for this restricted setting; that is, given that the first die landed on 3, each of these outcomes now has probability $\frac{1}{6}$.¹⁷ To conclude, the probability of the sum being 8 is then the probability of $(3, 5)$, which is $\frac{1}{6}$.

¹⁷ In fact, the (conditional) probability of all the other 30 events is now 0, since they cannot happen, given our evidence.

Let \mathcal{E} and \mathcal{F} denote the event that the sum of the dice is 8, and the event that the first die lands on 3, respectively. Then, what we have just computed is called the *conditional probability* of \mathcal{E} given \mathcal{F} , which is denoted by $P(\mathcal{E} \mid \mathcal{F})$.

To define the general notion of conditional probability, we follow the same intuition showcased by this example. Given that the event \mathcal{F} holds, in order to observe \mathcal{E} we in fact need to observe both events simultaneously; that is $\mathcal{E}\mathcal{F}$. However, since \mathcal{F} is already known, we can consider it as the new sample space (all outcomes out of it are now impossible) and hence the probability of observing \mathcal{E} becomes the probability of $\mathcal{E}\mathcal{F}$ relative to that of \mathcal{F} .

Definition 11. Let \mathcal{E} and \mathcal{F} be two events over the same sample space, such that $P(\mathcal{F}) > 0$. The *conditional probability* of \mathcal{E} given \mathcal{F} is

$$P(\mathcal{E} \mid \mathcal{F}) = \frac{P(\mathcal{E}\mathcal{F})}{P(\mathcal{F})}$$

Notice that this definition only makes sense if $P(\mathcal{F}) > 0$. In fact, given our intuition that \mathcal{F} is an observed evidence, it is reasonable to make such an assumption, as we do not expect to observe an event of probability 0.

Although we originally motivated conditional probabilities as a way to update beliefs in the presence of evidence,¹⁸ Definition 11 is also consistent with the frequentist view. Suppose that an experiment is repeated a large number n of times. Since $P(\mathcal{F})$ expresses the proportion of times where \mathcal{F} is observed, this event will occur approximately $nP(\mathcal{F})$ times. Likewise, $nP(\mathcal{E}\mathcal{F})$ times \mathcal{E} and \mathcal{F} will both happen. Thus, out of the ap-

¹⁸ Following the subjective view

proximately $nP(\mathcal{F})$ experiments whose outcome is in \mathcal{F} , $nP(\mathcal{E}\mathcal{F})$ also belong to \mathcal{E} . In other words, for the experiments in \mathcal{F} , the proportion whose outcome is also in \mathcal{E} approximates $\frac{nP(\mathcal{E}\mathcal{F})}{nP(\mathcal{F})} = \frac{P(\mathcal{E}\mathcal{F})}{P(\mathcal{F})}$. This approximation gets tighter as the number of repetitions of the experiment increases.

Example 12. Suppose that a box contains 4 defective (do not work), 8 partially defective (work only briefly), and 20 working transistors. We choose one of them randomly, and use it. If it does not fail immediately, what is the probability that it is a working transistor?

Solution. Since the transistor did not fail immediately, we know that it is not one of the 4 defective ones. So we need to compute

$$P(\text{correct} \mid \overline{\text{defective}}) = \frac{P(\text{correct}, \overline{\text{defective}})}{P(\overline{\text{defective}})} = \frac{P(\text{correct})}{P(\overline{\text{defective}})},$$

where the last equation holds because every correct transistor is necessarily non-defective. Since the transistor was chosen randomly,¹⁹ we have

$$P(\text{correct} \mid \overline{\text{defective}}) = \frac{\frac{20}{32}}{\frac{28}{32}} = \frac{5}{7}. \quad \triangle$$

Example 13. You toss a coin twice, and on each throw you bet 5€ that it will fall on ‘heads.’ Knowing that you won at least one of the bets, what is the probability that you won 10€? (Assume that the coin is fair).

Solution. The sample space for the two coin tosses can be represented as $\mathcal{S} = \{(h, h), (h, t), (t, h), (t, t)\}$.²¹ You win 10€ if both tosses land on heads; call this event \mathcal{E} . If \mathcal{F} denotes the event that at least one toss is heads, then the desired probability is

$$P(\mathcal{E} \mid \mathcal{F}) = \frac{P(\mathcal{E}\mathcal{F})}{P(\mathcal{F})} = \frac{P(\{(h, h)\})}{P(\{(h, h), (h, t), (t, h)\})} = \frac{1/4}{3/4} = \frac{1}{3}. \quad \triangle$$

The equation for conditional probabilities from Definition 11 can be rewritten as $P(\mathcal{E}\mathcal{F}) = P(\mathcal{E} \mid \mathcal{F})P(\mathcal{F})$.²² This variation is helpful for computing probabilities of intersections of events, if the conditional probabilities are known.

Example 14. Your favourite football team has a 20% chance of reaching the final game, and if it does, then it has a 50% chance of winning. What is the probability of the team becoming the champion?

Solution. Let \mathcal{F} be the event that the team reaches the final game, and \mathcal{E} the event that it wins it. Then the probability of it being the champion is $P(\mathcal{E}\mathcal{F}) = P(\mathcal{E} \mid \mathcal{F})P(\mathcal{F}) = 0.5 \cdot 0.2 = 0.1$. \triangle

Bayes' Formula

FOR ANY TWO EVENTS \mathcal{E} and \mathcal{F} , \mathcal{E} can be equivalently expressed as a union $\mathcal{E}\mathcal{F} \cup \mathcal{E}\overline{\mathcal{F}}$ (see Figure 3). Indeed, any outcome in \mathcal{E} must either be in \mathcal{F} or in $\overline{\mathcal{F}}$, and hence also in the intersection of those events with \mathcal{E} . Since $\mathcal{E}\mathcal{F}$ and $\mathcal{E}\overline{\mathcal{F}}$ are mutually exclusive, Axiom 3 entails that

$$P(\mathcal{E}) = P(\mathcal{E}\mathcal{F}) + P(\mathcal{E}\overline{\mathcal{F}}) = P(\mathcal{E} \mid \mathcal{F})P(\mathcal{F}) + P(\mathcal{E} \mid \overline{\mathcal{F}})P(\overline{\mathcal{F}}). \quad (1)$$

¹⁹ Each transistor is equally likely to be chosen.

²⁰ Notice that the same probability can be computed through counting: since the transistor is not defective, the problem reduces to computing the probability that a transistor, chosen at random from a box with 20 correct and 8 partially defective transistors, is correct. This is, of course, $\frac{20}{28}$.

²¹ Where (h, t) represents that the first toss was heads, and the second tails.

²² In fact, to avoid the need to require $P(\mathcal{F}) > 0$, conditional probabilities are often defined just using this variant.




Figure 3: $\mathcal{E} = \mathcal{E}\mathcal{F} \cup \mathcal{E}\overline{\mathcal{F}}$.

This means that the probability of \mathcal{E} is a weighted sum of the conditional probabilities over any other event \mathcal{F} and its complement $\overline{\mathcal{F}}$, where the weights are given by the relative probabilities of the conditioning events to happen. This is very useful, because it allows to find probabilities by first conditioning relative to another event happening or not. There are many instances where this approach allows us to compute probabilities that would, otherwise, be impossible.

Example 15. There exist two kinds of people: cautious, and risk-takers. Within a one-year period, cautious people have an accident with probability of 0.2, while for risk-takers this probability is 0.4. If 30% of the population is risk-taker, what is the probability that a randomly chosen individual has an accident in this one-year period?

Solution. We condition over the chosen individual being a risk-taker or not. Let \mathcal{E} be the event that the individual has an accident in the one-year period, and \mathcal{F} the event that they are risk-takers.²³ Then,

²³ $\overline{\mathcal{F}}$ means that they are cautious.

$$P(\mathcal{E}) = P(\mathcal{E} | \mathcal{F})P(\mathcal{F}) + P(\mathcal{E} | \overline{\mathcal{F}})P(\overline{\mathcal{F}}) = 0.4 \cdot 0.3 + 0.2 \cdot 0.7 = 0.26 \quad \triangle$$

THE FORMULA IN EQUATION (1) is useful not only for computing probabilities of events by themselves, but also to update initial probability assessments in the presence of additional or new information. We show how this works with an example.

Example 16. Consider again Example 15, and suppose that the randomly chosen individual has an accident. What is the probability that they are a risk-taker?

Solution. Remember that, originally (lacking any further evidence) the individual has a 30% probability of being a risk-taker ($P(\mathcal{F}) = 0.3$). However, given the new information that they had an accident, we can re-evaluate this probability as follows:

$$P(\mathcal{F} | \mathcal{E}) = \frac{P(\mathcal{E} \mathcal{F})}{P(\mathcal{E})} = \frac{P(\mathcal{E} | \mathcal{F})P(\mathcal{F})}{P(\mathcal{E})} = \frac{0.3 \cdot 0.4}{0.26} = \frac{6}{13} \approx 0.4615. \quad \triangle$$

THIS SAME FORMULA²⁴ can be generalised by allowing for a more fine-grained partition of the sample space as follows. Let $\mathcal{F}_1, \dots, \mathcal{F}_n$ be a *partition* of the sample space \mathcal{S} .²⁵ In terms of events, this means that exactly one of these \mathcal{F}_i s must occur and, as before, $\mathcal{E} \equiv \bigcup_{i=1}^n \mathcal{E} \mathcal{F}_i$. Since these events are mutually exclusive, we get

²⁴ In equation (1)

²⁵ $\mathcal{F}_1, \dots, \mathcal{F}_n$ is a partition of the space \mathcal{S} if they are mutually exclusive events and $\bigcup_{i=1}^n \mathcal{F}_i = \mathcal{S}$.

$$P(\mathcal{E}) = \sum_{i=1}^n P(\mathcal{E} \mathcal{F}_i) = \sum_{i=1}^n P(\mathcal{E} | \mathcal{F}_i)P(\mathcal{F}_i). \quad (2)$$

That is, $P(\mathcal{E})$ can be computed by first conditioning over a partition of the sample space, and computing a weighted average of the results, where the weight corresponds to the probability of each of the events \mathcal{F}_i . Notice that the original formula is just a special case of this: \mathcal{G} and $\overline{\mathcal{G}}$ form a partition.

Suppose that we know that \mathcal{E} happened, and we are interested in knowing which of the events in the partition was observed. Using equation (2)

we get

$$P(\mathcal{F}_j | \mathcal{E}) = \frac{P(\mathcal{E} \mathcal{F}_j)}{P(\mathcal{E})} = \frac{P(\mathcal{E} | \mathcal{F}_j)P(\mathcal{F}_j)}{\sum_{i=1}^n P(\mathcal{E} | \mathcal{F}_i)P(\mathcal{F}_i)}. \quad (3)$$

This equation is known as *Bayes' formula*.²⁶ If we think of the events \mathcal{F}_j as possible *hypotheses* about some subject of study, Bayes' formula can be interpreted as showing how opinions held about these hypotheses (i.e., $P(\mathcal{F}_j)$) should change on basis of the evidence of the experiment.²⁷

Example 17. Consider a bug in a program, which may be in any of three different pieces of code with equal probability. For each of the three pieces i , let $1 - \alpha_i$ be the probability of finding the bug when searching in that piece given that it is, in fact, there.²⁸ If we search in the first piece, and we do not find any bug, what is the probability that the bug is in each of the three pieces $i \in \{1, 2, 3\}$?

Solution. Let $\mathcal{F}_i, 1 \leq i \leq 3$ be the event that the bug is in the i -th piece of code, and \mathcal{E} the event that the search in the first piece was unsuccessful. Then we have:

$$P(\mathcal{F}_1 | \mathcal{E}) = \frac{P(\mathcal{E} | \mathcal{F}_1)P(\mathcal{F}_1)}{\sum_{i=1}^3 P(\mathcal{E} | \mathcal{F}_i)P(\mathcal{F}_i)} = \frac{\alpha_1 \cdot (1/3)}{\alpha_1 \cdot (1/3) + 1 \cdot (1/3) + 1 \cdot (1/3)} = \frac{\alpha_1}{\alpha_1 + 2},$$

and for $j, 2 \leq i \leq 3$,

$$P(\mathcal{F}_j | \mathcal{E}) = \frac{P(\mathcal{E} | \mathcal{F}_j)P(\mathcal{F}_j)}{\sum_{i=1}^3 P(\mathcal{E} | \mathcal{F}_i)P(\mathcal{F}_i)} = \frac{1 \cdot (1/3)}{\alpha_1 \cdot (1/3) + 1 \cdot (1/3) + 1 \cdot (1/3)} = \frac{1}{\alpha_1 + 2}.$$

For instance, if $\alpha_1 = 0.4$, then the conditional probability that the bug is in the first piece, given that we did not find it there is $\frac{1}{6}$.[◦] \triangle

²⁶ After Thomas Bayes.

²⁷ The former is known as the *prior* probability, and the latter is the *posterior*.

²⁸ α_i is the *overlook probability*: how likely it is to oversee the bug when looking in the right place? Some bugs are difficult to catch.

[◦]What about the other pieces?

Independent Events

AS IT CAN BE seen from the previous examples, in general the conditional probability of an event \mathcal{E} given \mathcal{F} is not equal to the (unconditional) probability $P(\mathcal{E})$. That is, having evidence about \mathcal{F} generally changes the likelihood of observing \mathcal{E} . In the special case where $P(\mathcal{E} | \mathcal{F}) = P(\mathcal{E})$, we say that \mathcal{E} and \mathcal{F} are *independent*. In other words, \mathcal{E} and \mathcal{F} are independent if knowledge about \mathcal{F} does not change the probability of \mathcal{E} occurring.

Recall that $P(\mathcal{E}\mathcal{F}) = P(\mathcal{E} | \mathcal{F})P(\mathcal{F})$. Thus, if \mathcal{E} and \mathcal{F} are independent we have $P(\mathcal{E}\mathcal{F}) = P(\mathcal{E})P(\mathcal{F})$. The converse implication holds too.

Definition 18. Two events \mathcal{E}, \mathcal{F} are *independent* if $P(\mathcal{E}\mathcal{F}) = P(\mathcal{E})P(\mathcal{F})$. If \mathcal{E} and \mathcal{F} are not independent, they are *dependent*.

Example 19. A card is selected at random from a standard deck of 52 cards.²⁹ Let \mathcal{E} be the event of selecting an ace, and \mathcal{F} the event of selecting a red card. Then, \mathcal{E} and \mathcal{F} are independent because $P(\mathcal{E}\mathcal{F}) = \frac{1}{26}$, $P(\mathcal{E}) = \frac{4}{52}$, and $P(\mathcal{F}) = \frac{26}{52}$. \triangle

²⁹ French playing cards.

Proposition 20. If \mathcal{E} and \mathcal{F} are independent, then so are \mathcal{E} and $\overline{\mathcal{F}}$.³⁰

³⁰ That is, if \mathcal{E} and \mathcal{F} are independent, the likelihood of \mathcal{E} is unchanged by information about \mathcal{F} , whether it holds or not.

Proof. Assume that \mathcal{E} and \mathcal{F} are independent. Then we have

$$P(\mathcal{E}) = P(\mathcal{E}\mathcal{F}) + P(\mathcal{E}\overline{\mathcal{F}}) = P(\mathcal{E})P(\mathcal{F}) + P(\mathcal{E}\overline{\mathcal{F}}),$$

where the last equality follows from the independence of \mathcal{E} and \mathcal{F} . Then

$$P(\mathcal{E}\overline{\mathcal{F}}) = P(\mathcal{E})(1 - P(\mathcal{F})) = P(\mathcal{E})P(\overline{\mathcal{F}}). \quad \square$$

We know then that independence is symmetric and closed under complementation. However, surprisingly, it is *not* closed under intersections. That is, if \mathcal{E} is independent of \mathcal{F} and of \mathcal{G} , it is not necessarily the case that \mathcal{E} and $\mathcal{F}\mathcal{G}$ are independent.

Example 21. You throw two fair dice. Let \mathcal{E} be the event that the sum of the two dice is 7, \mathcal{F} the event that the first die is 1, and \mathcal{G} the event that the second die is 6. It can be shown that \mathcal{E} is independent of both \mathcal{F} and \mathcal{G} .[◦] But, obviously, \mathcal{E} is not independent of $\mathcal{F}\mathcal{G}$; in fact, $P(\mathcal{E} | \mathcal{F}\mathcal{G}) = 1$.

[◦]You will show this at the lab.

This example shows that defining independence between more than two events is more complex than merely checking that all pairs of events are independent between them. We get the following definition.

Definition 22. Three events \mathcal{E} , \mathcal{F} , and \mathcal{G} are *independent* if:³¹

$$\begin{aligned} P(\mathcal{E}\mathcal{F}\mathcal{G}) &= P(\mathcal{E})P(\mathcal{F})P(\mathcal{G}), \\ P(\mathcal{E}\mathcal{F}) &= P(\mathcal{E})P(\mathcal{F}), \\ P(\mathcal{E}\mathcal{G}) &= P(\mathcal{E})P(\mathcal{G}), \quad \text{and} \\ P(\mathcal{F}\mathcal{G}) &= P(\mathcal{F})P(\mathcal{G}). \end{aligned}$$

³¹ Informally, this can be understood as every subset of events being mutually independent.

If \mathcal{E} , \mathcal{F} , and \mathcal{G} are independent, then \mathcal{E} will necessarily be independent of any event formed by combinations of \mathcal{F} and \mathcal{G} . For example, \mathcal{E} is independent of $\mathcal{F} \cup \mathcal{G}$ because

$$\begin{aligned} P(\mathcal{E}(\mathcal{F} \cup \mathcal{G})) &= P(\mathcal{E}\mathcal{F} \cup \mathcal{E}\mathcal{G}) = P(\mathcal{E}\mathcal{F}) + P(\mathcal{E}\mathcal{G}) - P(\mathcal{E}\mathcal{F}\mathcal{G}) \\ &= P(\mathcal{E})P(\mathcal{F}) + P(\mathcal{E})P(\mathcal{G}) - P(\mathcal{E})P(\mathcal{F})P(\mathcal{G}) \\ &= P(\mathcal{E})(P(\mathcal{F}) + P(\mathcal{G}) - P(\mathcal{F}\mathcal{G})) = P(\mathcal{E})P(\mathcal{F} \cup \mathcal{G}). \end{aligned}$$

Definition 22 is generalized to more than three events in the obvious way. The events $\mathcal{E}_1, \dots, \mathcal{E}_n$ are *independent* iff for every subset $\mathcal{F}_1, \dots, \mathcal{F}_m$, $m \leq n$ of these events, it holds that $P(\mathcal{F}_1 \cdots \mathcal{F}_m) = P(\mathcal{F}_1) \cdots P(\mathcal{F}_m)$.

Considering many different independent events becomes important when we deal with a sequence of repetitions of an experiment. For example, if an experiment consists in rolling a die several times, we may see each roll of the die as a subexperiment, and it makes sense to assume that the outcomes of previous (or future) rolls have no effect on the outcome of the current roll.

Example 23. Consider the parallel system in Figure 4, which works if at least one of the n components work. Each component i , independently of all others, works with probability p_i , $1 \leq i \leq n$, what is the probability that the system functions?




Figure 4: A parallel system with current flowing from s to t .

Solution. Let \mathcal{E} be the event that the system functions, and $\mathcal{F}_i, 1 \leq i \leq n$ be the event that the component i functions. Then $P(\overline{\mathcal{E}})$ is exactly the probability that all components fail; that is

$$P(\mathcal{E}) = 1 - P(\overline{\mathcal{E}}) = 1 - P(\overline{\mathcal{F}_1} \cdots \overline{\mathcal{F}_n}) = 1 - \prod_{i=1}^n P(\overline{\mathcal{F}_i}),$$

where the last equality arises from the independence of the events.³²

³² The last expression can be rewritten to use the probabilities of the events, and not their complements, if needed.

Random Variables

USUALLY, WHEN RUNNING experiments, we are only interested in some numerical values determined by the results. For example, after rolling two dice, we may be interested in their sum, but not on the specific individual die-values that led to that sum.³³ These quantities of interest are known as *random variables*.³⁴ Since the values of random variables are determined by the outcome of an experiment, they are associated with a probability degree.

Consider for example the random variable \mathcal{X} given by the sum of two fair dice. \mathcal{X} can take only values between 2 and 12, and the probability of each of these values is

$$P([\mathcal{X} = 2]) = P(\{(1, 1)\}) = \frac{1}{36}$$

$$P([\mathcal{X} = 3]) = P(\{(1, 2), (2, 1)\}) = \frac{2}{36}$$

$$P([\mathcal{X} = 4]) = P(\{(1, 3), (2, 2), (3, 1)\}) = \frac{3}{36}$$

$$P([\mathcal{X} = 5]) = P(\{(1, 4), (2, 3), (3, 2), (4, 1)\}) = \frac{4}{36}$$

$$P([\mathcal{X} = 6]) = P(\{(1, 5), (2, 4), (3, 3), (4, 2), (5, 1)\}) = \frac{5}{36}$$

$$P([\mathcal{X} = 7]) = P(\{(1, 6), (2, 5), (3, 4), (4, 3), (5, 2), (6, 1)\}) = \frac{6}{36}$$

$$P([\mathcal{X} = 8]) = P(\{(2, 6), (3, 5), (4, 4), (5, 3), (6, 2)\}) = \frac{5}{36}$$

$$P([\mathcal{X} = 9]) = P(\{(3, 6), (4, 5), (5, 4), (6, 3)\}) = \frac{4}{36}$$

$$P([\mathcal{X} = 10]) = P(\{(4, 6), (5, 5), (6, 4)\}) = \frac{3}{36}$$

$$P([\mathcal{X} = 11]) = P(\{(5, 6), (6, 5)\}) = \frac{2}{36}$$

$$P([\mathcal{X} = 12]) = P(\{(6, 6)\}) = \frac{1}{36}.$$

Since \mathcal{X} must take some value, the sum of all these probabilities must be 1. This fact can be easily verified from this distribution.

Another random variable \mathcal{Y} of interest can be the value of the first die. In this case, \mathcal{Y} is equally likely to take any of the values from 1 to 6.

These examples presented random variables taking finitely many different values. Random variables whose values can be described as a (finite or infinite) sequence x_1, x_2, \dots are called *discrete*. There are also random variables that can take a continuum of possible values; these are called *continuous*. For example, the weight of a person can take any value in

³³ If the sum is 7, we do not care which of the possible 6 ways to get this sum was actually observed.

³⁴ This is a bad choice of name, as they are not *random*, nor *variables*.

some interval (a, b) .³⁵

Definition 24. The *cumulative distribution function* (often called simply *distribution function*) of a random variable \mathcal{X} is the function F defined for every real number x by:³⁶

$$F(x) := P[\mathcal{X} \leq x].$$

The notation $\mathcal{X} \sim F$ expresses that F is the distribution function of \mathcal{X} .

All probability questions about \mathcal{X} can be answered in terms of its distribution function F . For example, to compute $P[a < \mathcal{X} \leq b]$, we use the fact that $[\mathcal{X} \leq b]$ can be decomposed into the two mutually exclusive events $[\mathcal{X} \leq a]$ and $[a < \mathcal{X} \leq b]$. Hence, $P[\mathcal{X} \leq b] = P[\mathcal{X} \leq a] + P[a < \mathcal{X} \leq b]$. From this, we can deduce

$$P[a < \mathcal{X} \leq b] = F(b) - F(a).$$

Example 25. Consider the continuous random variable \mathcal{X} with the distribution function

$$F(x) = \begin{cases} 0 & x \leq 0 \\ 1 - \exp(-x^2) & x > 0 \end{cases}$$

The probability that \mathcal{X} is greater than 1 is

$$P[\mathcal{X} > 1] = 1 - P[\mathcal{X} \leq 1] = 1 - F(1) = e^{-1} = 0.368. \quad \triangle$$

Types of Random Variables

IF A RANDOM VARIABLE (RV) \mathcal{X} is discrete,^o then we can define its *probability mass function* as $p(x) := P[\mathcal{X} = x]$. Since the variable is discrete, and \mathcal{X} must take one value, we know that $p(x)$ is positive for at most countably many elements,³⁷ and that $\sum_{i=1}^{\infty} p(x_i) = 1$.

Example 26. Consider a RV \mathcal{X} that can take values from $\{1, 2, 3\}$. If we know that $p(2) = \frac{1}{3}$ and $p(3) = \frac{1}{6}$, then we can deduce that

$$p(1) = 1 - p(2) - p(3) = 1 - \frac{1}{3} - \frac{1}{6} = \frac{1}{2}.$$

This function is shown in Figure 5. \triangle

The cumulative distribution function F can be computed in these cases from p by adding all the relevant values: $F(x) = \sum_{y \leq x} p(y)$. Whenever \mathcal{X} is a discrete RV whose possible values are $x_1 < x_2 < \dots$, the distribution function F is a *step function*: it remains constant for the interval $[x_{i-1}, x_i)$, and then makes a jump of size $p(x_i)$ at x_i . For instance, if \mathcal{X} has the probability mass function given in Example 26, then the cumulative distribution function of \mathcal{X} is

$$F(x) = \begin{cases} 0 & x < 1 \\ \frac{1}{2} & 1 \leq x < 2 \\ \frac{5}{6} & 2 \leq x < 3 \\ 1 & 3 \leq x, \end{cases}$$

as depicted in Figure 6.

³⁵ There are also *mixed* random variables, but we will not consider them much during this lecture.

³⁶ For this course, we consider only random variables taking real numbers as values.

^oThat is, if it can only take values from a sequence x_1, x_2, \dots

³⁷ At most those that can be taken by \mathcal{X} .




Figure 5: Probability mass function from Example 26.




Figure 6: Cumulative distribution function from the RV of Example 26.

WE OFTEN NEED TO consider random variables that can take all the possible values from a (real) interval.

Definition 27. The RV \mathcal{X} is *continuous* if there exists a non-negative function f defined over all real numbers $x \in \mathbb{R}$ such that for any set $B \subseteq \mathbb{R}$

$$P[\mathcal{X} \in B] = \int_B f(x) dx.$$

This function f is called the *probability density function* of \mathcal{X} .


Essentially, to compute the probability of an event defined by a set of values B , we find the integral of the probability density function, which is the intuitive generalization of the mass function for continuous variables. Since \mathcal{X} must take some value in $(-\infty, \infty)$, $f(x)$ must satisfy

$$1 = P[\mathcal{X} \in (-\infty, \infty)] = \int_{-\infty}^{\infty} f(x) dx.$$

All probability statements can be answered integrating over $f(x)$; for instance, if $B = [a, b]$, then $P[a \leq \mathcal{X} \leq b] = \int_a^b f(x) dx$. In particular, if $a = b$ then we have $P[\mathcal{X} = a] = \int_a^a f(x) dx = 0$.³⁸ As an example, Figure 7 shows the probability density function

$$f(x) = \begin{cases} e^{-x} & x \geq 0 \\ 0 & x < 0, \end{cases}$$

and the area that represents $P[a < \mathcal{X} < b]$.



³⁸ If \mathcal{X} is continuous, then the probability of observing any specific value is always 0.

Figure 7: A probability density function, and the area representing $P[a < \mathcal{X} < b]$.

The probability density function f and the cumulative distribution F are related by the expression

$$F(a) = P[\mathcal{X} \leq a] = \int_{-\infty}^a f(x) dx,$$

or equivalently, $\frac{d}{dx}F(x) = f(x)$. In other words, the density is the derivate of the cumulative distribution.

If the probability density function is continuous, then taking a small ε yields, from Definition 27,

$$P[a - \frac{\varepsilon}{2} \leq \mathcal{X} \leq a + \frac{\varepsilon}{2}] = \int_{a - \frac{\varepsilon}{2}}^{a + \frac{\varepsilon}{2}} f(x) dx \approx \varepsilon f(a).$$

In other words, f tells us how likely it is to see a value that is close enough to a .

Example 28. Consider the continuous random variable \mathcal{X} with the probability density function

$$f(x) = \begin{cases} c(4x - 2x^2) & 0 < x < 2 \\ 0 & \text{otherwise,} \end{cases}$$

for some constant c .

1. What is the value of c ?
2. Compute $P[\mathcal{X} > 1]$.

Solution. Since f is a density function, it must hold that

$$1 = \int_{-\infty}^{\infty} f(x) dx = \int_0^2 c(4x - 2x^2) dx = c \left(2x^2 - \frac{2x^3}{3} \right) \Big|_0^2,$$

and hence $c = \frac{3}{8}$.

From this, we can conclude that

$$P[\mathcal{X} > 1] = \int_1^{\infty} f(x) dx = \frac{3}{8} \int_1^2 (4x - 2x^2) dx = \frac{1}{2}. \quad \triangle$$

Joint Distributions

WHEN DEALING WITH EXPERIMENTS, we are often interested in the relationship between two or more random variables, rather than merely observing one. For example, for health policies, we may be interested in the relationship between hours spent sitting down and the incidence of back pain, or want to understand the relationship between working hours and productivity of employees.

To consider two random variables \mathcal{X} and \mathcal{Y} simultaneously, we need a *joint cumulative probability distribution function* F that specifies the probability of \mathcal{X} and \mathcal{Y} to be below a given value; that is,

$$F(x, y) = P[\mathcal{X} \leq x, \mathcal{Y} \leq y].$$

Knowing the joint probability distribution allows us to find the probabilities of different statements concerning with the variables \mathcal{X} and \mathcal{Y} . For example, the distribution function $F_{\mathcal{X}}$ of \mathcal{X} is obtained by not imposing any limit on the values of \mathcal{Y} ; that is,

$$F_{\mathcal{X}}(x) = P[\mathcal{X} \leq x] = P[\mathcal{X} \leq x, \mathcal{Y} < \infty] = F(x, \infty).$$

Similarly, the cumulative distribution function of \mathcal{Y} is $F_{\mathcal{Y}}(y) = F(\infty, y)$.³⁹

For discrete random variables \mathcal{X} and \mathcal{Y} taking values x_1, x_2, \dots , and y_1, y_2, \dots , respectively, the *joint probability mass function* p of \mathcal{X} and \mathcal{Y} is defined in the obvious way: $p(x, y) = P[\mathcal{X} = x, \mathcal{Y} = y]$.

The individual probability mass functions of each of the variables can be easily obtained by eliminating the variable that is not of interest. E.g.,

³⁹ Formally, these would in fact be limits (i.e., $F_{\mathcal{X}}(x) = \lim_{y \rightarrow \infty} F(x, y)$), but this intuition serves to our purposes.

since \mathcal{Y} must take some value y_j , the event $[\mathcal{X} = x]$ is equivalent to the union of the mutually exclusive events $[\mathcal{X} = x, \mathcal{Y} = y_j]$. Using Axiom 3 of probabilities, we have

$$P[\mathcal{X} = x] = P\left(\bigcup_j [\mathcal{X} = x, \mathcal{Y} = y_j]\right) = \sum_j P[\mathcal{X} = x, \mathcal{Y} = y_j] = \sum_j p(x, y_j).$$

Analogously, $P[\mathcal{Y} = y] = \sum_i p(x_i, y)$.

This says that the joint probability mass function fully determines the individual probability mass functions for each of the random variables involved. However, the converse does not hold: knowing $P[\mathcal{X} = x]$ and $P[\mathcal{Y} = y]$ is not sufficient for deriving $P[\mathcal{X} = x, \mathcal{Y} = y]$.

Example 29. Consider a box of batteries, where 2 are new, 3 are partially charged, and 4 are completely empty. We randomly select 3 batteries from this box. Let \mathcal{X} and \mathcal{Y} denote the number of new, and the number of partially charged batteries selected, respectively. The joint probability mass function $p(x, y) = P[\mathcal{X} = x, \mathcal{Y} = y]$ of \mathcal{X} and \mathcal{Y} is:[◦]

$$\begin{aligned} p(0, 0) &= \frac{\binom{4}{3}}{\binom{9}{3}} = \frac{4}{84} & p(0, 1) &= \frac{\binom{3}{1}\binom{4}{2}}{\binom{9}{3}} = \frac{18}{84} \\ p(0, 2) &= \frac{\binom{3}{2}\binom{4}{1}}{\binom{9}{3}} = \frac{12}{84} & p(0, 3) &= \frac{\binom{3}{3}}{\binom{9}{3}} = \frac{1}{84} \\ p(1, 0) &= \frac{\binom{2}{1}\binom{4}{2}}{\binom{9}{3}} = \frac{12}{84} & p(1, 1) &= \frac{\binom{2}{1}\binom{3}{1}\binom{4}{1}}{\binom{9}{3}} = \frac{24}{84} \\ p(1, 2) &= \frac{\binom{2}{1}\binom{3}{2}}{\binom{9}{3}} = \frac{6}{84} & p(2, 0) &= \frac{\binom{2}{2}\binom{4}{1}}{\binom{9}{3}} = \frac{4}{84} \\ p(2, 1) &= \frac{\binom{2}{2}\binom{3}{1}}{\binom{9}{3}} = \frac{3}{84} \end{aligned}$$

A simpler way to express all these probabilities is shown in Table 1. △

Notice that the probability mass function of \mathcal{X} is obtained by summing the elements in a row, while the mass function of \mathcal{Y} appears from summing through the columns. These probabilities are often known as the *marginal* probability mass functions of \mathcal{X} and \mathcal{Y} , respectively.⁴⁰ To check the correctness of such a probability table, one should check that the sum of the marginal row and of the marginal column is 1.[◦]

Example 30. Consider a community where 15 percent of the families have no children, 20 percent have 1, 35 percent have 2, and 30 percent have 3. Suppose that each child is equally likely to be a boy or a girl. From a randomly chosen family, let \mathcal{X} be the random variable measuring the number of boys, and \mathcal{Y} the number of girls. The joint probability mass function is shown in Table 2. The cells in the first row are computed next.[◦]

$$P[\mathcal{X} = 0, \mathcal{Y} = 0] = P[\text{no children}] = 0.15$$

$$P[\mathcal{X} = 0, \mathcal{Y} = 1] = P[1 \text{ child that is girl}]$$

$$= P[1 \text{ child}]P[1 \text{ girl} | 1 \text{ child}] = 0.2 \cdot 0.5 = 0.1$$

$$P[\mathcal{X} = 0, \mathcal{Y} = 2] = P[2 \text{ children that are girls}]$$

Axioms of probability:

1. $0 \leq P(\mathcal{E}) \leq 1$
2. $P(\mathcal{S}) = 1$
3. $P(\bigcup_{i=1}^n \mathcal{E}_i) = \sum_{i=1}^n P(\mathcal{E}_i)$ if \mathcal{E}_i s are mutually exclusive.

[◦]Do you understand why?

Table 1: The joint mass function $p(x, y)$ from Example 29. The denominator 84 in the values is not included.

$i \backslash j$	0	1	2	3	Sum
0	4	18	12	1	35
1	12	24	6	0	42
2	4	3	0	0	7
Sum	20	45	18	1	

⁴⁰ To remember this name, think that they appear in the *margin* of the joint probability table.

[◦]Do you understand why? In the case of Table 1, the sum should be 84, because the denominator 84 is not written.

[◦]Verify the others!

$i \backslash j$	0	1	2	3	Sum
0	0.1500	0.1000	0.0875	0.0375	0.3750
1	0.1000	0.1750	0.1125	0	0.3875
2	0.0875	0.1125	0	0	0.2000
3	0.0375	0	0	0	0.0375
Sum	0.3750	0.3875	0.2000	0.0375	

Table 2: The joint mass $p(x, y)$ from Example 30.

$$= P[2 \text{ children}]P[2 \text{ girls} \mid 2 \text{ children}] = 0.35 \cdot (0.5)^2 = 0.0875$$

$$P[\mathcal{X} = 0, \mathcal{Y} = 3] = P[3 \text{ children}]P[3 \text{ girls} \mid 3 \text{ children}] = 0.3 \cdot (0.5)^3 = 0.0375.$$

From Table 2 it can be seen that, for example, the probability of having at least one girl is 0.625.⁴¹ \triangle

Definition 31. The RVs \mathcal{X} and \mathcal{Y} are *jointly continuous* if there exists a function $f(x, y)$ defined for all $x, y \in \mathbb{R}$ such that for every set $C \subseteq \mathbb{R} \times \mathbb{R}$ ⁴²

$$P[(\mathcal{X}, \mathcal{Y}) \in C] = \iint_{(x, y) \in C} f(x, y) dx dy.$$

The function $f(x, y)$ is called the *joint probability density function* of \mathcal{X} and \mathcal{Y} .

If $A, B \subseteq \mathbb{R}$ are two sets of real numbers, then it follows from this definition that

$$P[\mathcal{X} \in A, \mathcal{Y} \in B] = \int_B \int_A f(x, y) dx dy.$$

From $F(a, b) = P[\mathcal{X} \leq a, \mathcal{Y} \leq b] = \int_{-\infty}^b \int_{-\infty}^a f(x, y) dx dy$, it follows (via differentiation) that whenever the partial derivatives are defined,

$$f(a, b) = \frac{\partial^2}{\partial a \partial b} F(a, b).$$

As in the case of single random variables, $f(a, b)$ is a measure of how likely it is for the random vector $(\mathcal{X}, \mathcal{Y})$ to appear near (a, b) ; but the probability of being *exactly* (a, b) remains 0.

Notice that if \mathcal{X} and \mathcal{Y} are jointly continuous, then each of them is continuous individually. Moreover, the probability density function of \mathcal{X} is $f_{\mathcal{X}}(x) := \int_{-\infty}^{\infty} f(x, y) dy$. Hence,

$$P[\mathcal{X} \in A] = P[\mathcal{X} \in A, \mathcal{Y} \in \mathbb{R}] = \int_A \int_{-\infty}^{\infty} f(x, y) dy dx = \int_A f_{\mathcal{X}}(x) dx.$$

Example 32. Consider the joint density function for \mathcal{X} and \mathcal{Y}

$$f(x, y) = \begin{cases} 2e^{-x}e^{-2y} & x, y > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Compute $P[\mathcal{X} > 1, \mathcal{Y} < 1]$, $P[\mathcal{X} < \mathcal{Y}]$, and $P[\mathcal{X} < a]$.⁴³

⁴³ Recall that $\int e^{kx} dx = (1/k)e^{kx} + c$.

Solution.

$$\begin{aligned} P[\mathcal{X} > 1, \mathcal{Y} < 1] &= \int_0^1 \int_1^{\infty} 2e^{-x}e^{-2y} dx dy \\ &= \int_0^1 2e^{-2y} (-e^{-x}) \Big|_1^{\infty} dy = e^{-1} \int_0^1 2e^{-2y} dy \end{aligned}$$

$$\begin{aligned}
&= e^{-1}(-e^{-2y}|_0^1) = e^{-1}(1 - e^{-2}) \\
P[\mathcal{X} < \mathcal{Y}] &= \iint_{(x,y)|x < y} 2e^{-x}e^{-2y} dx dy = \int_0^\infty \int_0^y 2e^{-x}e^{-2y} dx dy \\
&= \int_0^\infty 2e^{-2y}(1 - e^{-y}) dy = \int_0^\infty 2e^{-2y} dy - \int_0^\infty 2e^{-3y} dy \\
&= 1 - \frac{2}{3} = \frac{1}{3} \\
P[\mathcal{X} < a] &= \int_0^a \int_0^\infty 2e^{-x}e^{-2y} dy dx = \int_0^a e^{-x} dx = 1 - e^{-a} \quad \triangle
\end{aligned}$$

Independent Random Variables

THE RANDOM VARIABLES \mathcal{X} and \mathcal{Y} are *independent* if for any two sets $A, B \subseteq \mathbb{R}$ it holds that $P[\mathcal{X} \in A, \mathcal{Y} \in B] = P[\mathcal{X} \in A]P[\mathcal{Y} \in B]$.[◦] That is, they are independent if for all possible sets A, B , the events $\mathcal{E}_A = [\mathcal{X} \in A]$ and $\mathcal{F}_B = [\mathcal{Y} \in B]$ are independent. This definition of independence is equivalent to requiring that for all $a, b \in \mathbb{R}$

$$P[\mathcal{X} \leq a, \mathcal{Y} \leq b] = P[\mathcal{X} \leq a]P[\mathcal{Y} \leq b],$$

or, alternatively, that $F(a, b) = F_{\mathcal{X}}(a)F_{\mathcal{Y}}(b)$.^{44◦}

If \mathcal{X} and \mathcal{Y} are discrete, independence is equivalent to requiring, for all $x, y \in \mathbb{R}$ that $p(x, y) = p_{\mathcal{X}}(x)p_{\mathcal{Y}}(y)$, where $p_{\mathcal{X}}$ and $p_{\mathcal{Y}}$ are the probability mass functions. In the continuous case, it refers to $f(x, y) = f_{\mathcal{X}}(x)f_{\mathcal{Y}}(y)$.⁴⁵

Example 33. Let \mathcal{X} and \mathcal{Y} be two independent random variables, each having the density function

$$f(x) = \begin{cases} e^{-x} & x > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Find the density function of \mathcal{X}/\mathcal{Y} .

Solution. We first compute the distribution function $F_{\mathcal{X}/\mathcal{Y}}$ of \mathcal{X}/\mathcal{Y} . Given $a > 0$,

$$\begin{aligned}
F_{\mathcal{X}/\mathcal{Y}}(a) &= P[\mathcal{X}/\mathcal{Y} \leq a] = \iint_{x/y \leq a} f(x, y) dx dy \\
&= \iint_{x/y \leq a} e^{-x}e^{-y} dx dy = \int_0^\infty \int_0^{ay} e^{-x}e^{-y} dx dy \\
&= \int_0^\infty (1 - e^{-ay})e^{-y} dy = \left(-e^{-y} + \frac{e^{-(a+1)y}}{a+1} \right) \Big|_0^\infty = 1 - \frac{1}{a+1}.
\end{aligned}$$

Differentiating this function over the variable a , we obtain the density function $f_{\mathcal{X}/\mathcal{Y}}(a) = 1/(a+1)^2$, for $0 < a < \infty$. \triangle

WE CAN GENERALIZE the notion of joint probability distributions from 2 to n random variables, defining the analogous notion of the joint cumulative probability distribution F , the joint probability mass function p if they are discrete, the joint probability density function f if they are continuous, and independence if every finite subcollection of random variables is independent.

[◦] Recall Definition 18.

⁴⁴ To prove this, notice that one direction corresponds directly to the definition. The other direction follows from the axioms of probability, with caveats of which sets are allowed.

[◦] F is the joint distribution function.

⁴⁵ As for events, random variables that are not independent are called *dependent*.

Example 34. Consider a stock, whose price changes every day independently with the probability mass function $p(x)$ of the stock changing by x defined as

$$p(x) = \begin{cases} 0.05 & x \in \{-3, 3\} \\ 0.10 & x \in \{-2, 2\} \\ 0.15 & x \in \{-1, 1\} \\ 0.40 & x = 0. \end{cases}$$

If \mathcal{X}_i denotes the change in day i , then the probability that it increases in the next days by 1, 2, and 0 is

$$P[\mathcal{X}_1 = 1, \mathcal{X}_2 = 2, \mathcal{X}_3 = 0] = p(1)p(2)p(0) = 0.15 \cdot 0.1 \cdot 0.4 = 0.006. \quad \triangle$$

Expectation

A FUNDAMENTAL NOTION IN probability theory is the expectation of a random variable. If the random variable \mathcal{X} is discrete and takes values x_1, x_2, \dots , then the *expectation* or *expected value* of \mathcal{X} is

$$E[\mathcal{X}] := \sum_i x_i P[\mathcal{X} = x_i] = \sum_i x_i p(x_i).$$

That is, $E[\mathcal{X}]$ is the weighted average of the values of \mathcal{X} , where the weights are given by their probability of occurrence. For example, if the probability mass function of \mathcal{X} is such that $p(0) = \frac{1}{2} = p(1)$, then $E[\mathcal{X}] = 0\frac{1}{2} + 1\frac{1}{2} = \frac{1}{2}$,⁴⁶ but if $p(0) = \frac{1}{3}$ and $p(1) = \frac{2}{3}$, then $E[\mathcal{X}] = 0\frac{1}{3} + 1\frac{2}{3} = \frac{2}{3}$. This follows from the fact that the value 1 is twice as likely to appear as the value 0.

⁴⁶ The usual average of the values of \mathcal{X} .

An intuitive motivation for the notion of expectation arises from the frequentist view. If an experiment is repeated independently, then the proportion of times we observe a given event \mathcal{E} is $P(\mathcal{E})$. Consider a random variable \mathcal{X} taking values x_1, x_2, \dots , which represents the winnings of one execution of a bet, and let p be its mass function.⁴⁷ Suppose that we repeat the bet n times, for n sufficiently large. Then, in approximately $np(x_i)$ of those times, we win x_i units. Since this is true for all i , overall we win $\sum_i x_i \cdot np(x_i)$. Thus, in average, on each bet, we will win

⁴⁷ That is, with probability $p(x_i)$, we win x_i units.

$$\sum_i \frac{x_i \cdot np(x_i)}{n} = \sum_i x_i p(x_i) = E[\mathcal{X}].$$

Example 35. If \mathcal{X} is the outcome of rolling a fair die, then $p(i) = 1/6$ for all $i, 1 \leq i \leq 6$. Hence, we get

$$E[\mathcal{X}] = 1\left(\frac{1}{6}\right) + 2\left(\frac{1}{6}\right) + 3\left(\frac{1}{6}\right) + 4\left(\frac{1}{6}\right) + 5\left(\frac{1}{6}\right) + 6\left(\frac{1}{6}\right) = \frac{21}{6} = \frac{7}{2}. \quad \triangle$$

Notice from this example that the expectation of \mathcal{X} might be a value that can never occur in \mathcal{X} . The expected value is thus not what we will likely observe, but rather the average of the results over a long run of repetitions of the experiment.⁴⁸

⁴⁸ If we roll a die many times, the average of the results will converge to 3.5.

The *indicator variable* for an event \mathcal{E} is a random variable that takes value 1 if \mathcal{E} occurs, and 0 otherwise. If \mathcal{X} is the indicator variable of \mathcal{E} , then $E[\mathcal{X}] = 1P(\mathcal{E}) + 0P(\overline{\mathcal{E}}) = P(\mathcal{E})$. That is, the expected value of an indicator variable is exactly the probability of the event it indicates.

ONE CAN ALSO DEFINE the expectation for continuous random variables. Similarly to the discrete case, for a continuous random variable \mathcal{X} with probability density function f , we have

$$E[\mathcal{X}] := \int_{-\infty}^{\infty} xf(x)dx.$$

Example 36. Your IKEA order will arrive at some point after 17:00. You know by experience that the number of hours \mathcal{X} that you have to wait for the arrival is a random variable with the probability density function

$$f(x) = \begin{cases} \frac{1}{2} & 0 < x < 2 \\ 0 & \text{otherwise.} \end{cases}$$

The expected amount of time that you will wait is then

$$E[\mathcal{X}] = \int_0^2 \frac{x}{2} dx = 1.$$

That is, on average, you have to wait one hour. \triangle

The notion of expectation can be intuitively understood as that of the center of gravity. Suppose, for a discrete random variable \mathcal{X} , that the probability mass distribution is interpreted literally, with a physical object with mass proportional to $p(x_i)$ located at each point x_i . Then, $E[\mathcal{X}]$ is the point where such a structure would balance. See Figure 8. It is also important to notice that for every random variable \mathcal{X} , $E[\mathcal{X}]$ has the same units of measurement as \mathcal{X} .




Figure 8: Expected value as the center of gravity.

Properties of Expectation

CONSIDER A GIVEN RANDOM variable \mathcal{X} with its probability distribution. Often, we are not interested in $E[\mathcal{X}]$, but rather on a *function* of \mathcal{X} ; say $g(\mathcal{X})$. Notice that $g(\mathcal{X})$ is itself a random variable, and hence has a probability distribution that must depend on the distribution of \mathcal{X} . If we can compute this distribution, then $E[g(\mathcal{X})]$ can be readily obtained from it.

Example 37. Let \mathcal{X} be the random variable with probability mass function $p(0) = 0.3$, $p(1) = 0.5$, and $p(2) = 0.2$. Compute $E[\mathcal{X}^2]$.

Solution. If $\mathcal{Y} = \mathcal{X}^2$, then \mathcal{Y} can take the values $0 = 0^2$, $1 = 1^2$ and $4 = 2^2$ with probabilities 0.3, 0.5, and 0.2, respectively. Then

$$E[\mathcal{X}^2] = E[\mathcal{Y}] = 0 \cdot 0.3 + 1 \cdot 0.5 + 4 \cdot 0.2 = 1.3. \quad \triangle$$

Example 38. The time that it takes to correct a bug in a piece of software is a random variable \mathcal{X} with density function

$$f(x) = \begin{cases} 1 & 0 < x < 1 \\ 0 & \text{otherwise.} \end{cases}$$

If the cost of keeping a bug for time x is x^3 , what is the expected cost of such a bug?

Solution. Let $\mathcal{Y} = \mathcal{X}^3$ be the random variable for the cost. The distribution $F_{\mathcal{Y}}$ of \mathcal{Y} is given, for every $0 \leq a \leq 1$ by

$$F_{\mathcal{Y}}(a) = P[\mathcal{Y} \leq a] = P[\mathcal{X}^3 \leq a] = P[\mathcal{X} \leq a^{1/3}] = \int_0^{a^{1/3}} 1 dx = a^{1/3}.$$

Through differentiation, we get the density $f_{\mathcal{Y}}(a) = \frac{1}{3}a^{-2/3}$ for $0 \leq a < 1$. Hence,

$$\begin{aligned} E[\mathcal{X}^3] &= E[\mathcal{Y}] = \int_{-\infty}^{\infty} a f_{\mathcal{Y}}(a) da \\ &= \int_0^1 a \frac{1}{3} a^{-2/3} da = \frac{1}{3} \int_0^1 a^{1/3} da \\ &= \frac{1}{3} \frac{3}{4} a^{4/3} \Big|_0^1 = 1/4 \end{aligned} \quad \triangle$$

As we can see, this approach always allows us to compute the expected value of any function of \mathcal{X} , if we know the distribution of \mathcal{X} . However, the process is far from obvious. An easy way to do compute this expectation is based on the following intuition: since $g(\mathcal{X})$ takes the value $g(x)$ whenever $\mathcal{X} = x$, $E[g(\mathcal{X})]$ should be the weighted average of the values of $g(x)$ with the weights given by $p(x)$ for each possible value x .⁴⁹

Proposition 39 (Expectation of Functions). *Let \mathcal{X} be a random variable, and g any real-valued function. Then:*

1. *if \mathcal{X} is discrete with probability mass function $p(x)$, then*

$$E[g(\mathcal{X})] = \sum_x g(x)p(x);$$

2. *if \mathcal{X} is continuous with probability density function $f(x)$, then*

$$E[g(\mathcal{X})] = \int_{-\infty}^{\infty} g(x)p(x)dx.$$

Corollary 40. *For any two constants $a, b \in \mathbb{R}$, $E[a\mathcal{X} + b] = aE[\mathcal{X}] + b$.*

Proof. We show here the discrete case.⁵⁰

$$E[a\mathcal{X} + b] = \sum_x (ax + b)p(x) = a \sum_x xp(x) + b \sum_x p(x) = aE[\mathcal{X}] + b \quad \square$$

Notice in particular that $E[b] = b$ and $E[a\mathcal{X}] = aE[\mathcal{X}]$.

The expected value of \mathcal{X} is also known as the *mean* or the *first moment* of \mathcal{X} . In general, the *n-th moment* of \mathcal{X} , for any $n \in \mathbb{N}$, is defined as $E[\mathcal{X}^n]$, which we know how to compute by Proposition 39.

PROPOSITION 39 CAN BE EXTENDED to deal with several random variables. In the case of two random variables \mathcal{X} and \mathcal{Y} , if g is a function over pairs of real numbers, then

$$E[g(\mathcal{X}, \mathcal{Y})] = \begin{cases} \sum_y \sum_x g(x, y)p(x, y) & \mathcal{X}, \mathcal{Y} \text{ discrete} \\ \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x, y)f(x, y)dx dy & \mathcal{X}, \mathcal{Y} \text{ jointly continuous.} \end{cases}$$

In particular, if $g(\mathcal{X}, \mathcal{Y}) = \mathcal{X} + \mathcal{Y}$, where \mathcal{X}, \mathcal{Y} are jointly continuous, we get

$$E[\mathcal{X} + \mathcal{Y}] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} (x + y)f(x, y)dx dy$$

⁴⁹ This intuition holds in general. We can verify that the following proposition holds for the previous two examples.

⁵⁰ The continuous case is left as an exercise to the interested student.

$$\begin{aligned}
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x f(x, y) dx dy + \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} y f(x, y) dx dy \\
&= E[\mathcal{X}] + E[\mathcal{Y}],
\end{aligned} \tag{†}$$

where (†) is a consequence of using $g_{\mathcal{X}}(\mathcal{X}, \mathcal{Y}) = \mathcal{X}$; that is

$$E[\mathcal{X}] = E[g_{\mathcal{X}}(\mathcal{X}, \mathcal{Y})] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x f(x, y) dx dy,$$

and similarly for $g_{\mathcal{Y}}(\mathcal{X}, \mathcal{Y}) = \mathcal{Y}$.[◦]

Since sums are associative, we can simply repeat this argument to obtain the general form for sums of random variables: for any $n \in \mathbb{N}$,

$$E\left[\sum_{i=1}^n \mathcal{X}_i\right] = \sum_{i=1}^n E[\mathcal{X}_i].$$

Example 41. Two fair dice are rolled. Find the expected value of their sum.

Solution. If \mathcal{X} is the random variable representing the sum, we know that we can compute $E[\mathcal{X}] = \sum_{i=1}^{12} iP[\mathcal{X} = i]$. But a simpler approach is to consider two variables \mathcal{Y}_1 and \mathcal{Y}_2 , which represent the result of each individual die. Then $\mathcal{X} = \mathcal{Y}_1 + \mathcal{Y}_2$, and $E[\mathcal{X}] = E[\mathcal{Y}_1] + E[\mathcal{Y}_2] = 7$.⁵¹

[◦]Try to prove that the same holds for discrete variables.

⁵¹ See Example 35.

Example 42. You have n different pairs of shoes, which are all mixed up in your closet. If you pair them completely at random (selecting one left and one right shoe), what is the expected number of correct pairs you get?

Solution. If \mathcal{X} is the number of correct pairs, we can compute $E[\mathcal{X}]$ via the random variables $\mathcal{X}_1, \dots, \mathcal{X}_n$, where for each $i, 1 \leq i \leq n$,

Indicator variables.

$$\mathcal{X}_i = \begin{cases} 1 & i\text{-th pair is correct} \\ 0 & \text{otherwise.} \end{cases}$$

Then, $\mathcal{X} = \sum_{i=1}^n \mathcal{X}_i$. Since the i -th left shoe is equally likely to be matched with any of the n right shoes, it follows that $P[\mathcal{X}_i = 1] = 1/n$,⁵² and hence $E[\mathcal{X}_i] = 1/n$. From this, we obtain $E[\mathcal{X}] = \sum_{i=1}^n E[\mathcal{X}_i] = n(1/n) = 1$.⁵³ \triangle

⁵² Probability of the i -th left shoe matching its right one.

⁵³ This means that regardless of how many pairs we start with, in average only one pair will be correct.

AN IMPORTANT PROPERTY of the mean arises when trying to predict the value of a random variable \mathcal{X} . If we predict that \mathcal{X} will take the value c (for example, when placing a bet), then the *square of the prediction error* is $(\mathcal{X} - c)^2$. Let $\mu = E[\mathcal{X}]$; then,

$$\begin{aligned}
E[(\mathcal{X} - c)^2] &= E[(\mathcal{X} - \mu + \mu - c)^2] = E[(\mathcal{X} - \mu)^2 + 2(\mathcal{X} - \mu)(\mu - c) + (\mu - c)^2] \\
&= E[(\mathcal{X} - \mu)^2] + 2(\mu - c)E[\mathcal{X} - \mu] + (\mu - c)^2 \\
&= E[(\mathcal{X} - \mu)^2] + (\mu - c)^2 \\
&\geq E[(\mathcal{X} - \mu)^2],
\end{aligned} \tag{†}$$

where (†) follows from $E[\mathcal{X} - \mu] = E[\mathcal{X}] - \mu = 0$. Thus, the average squared error is minimised when we predict that \mathcal{X} is equal to its mean μ .⁵⁴

⁵⁴ The best predictor of a random variable, in terms of minimizing the expected square of its error, is just the mean.

Variance

THE EXPECTED VALUE OF a random variable \mathcal{X} provides the weighted average of the values it may take, but this information does not tell us much about the *distribution* of the values; e.g., how far apart are the values from each other.⁵⁵ For example, the random variables \mathcal{X} and \mathcal{Y} with $p_{\mathcal{X}}(0) = 1$ and $p_{\mathcal{Y}}(-100) = p_{\mathcal{Y}}(100) = 0.5$, respectively, have both the same expectation; namely 0. However, the values of \mathcal{Y} are much more separated between themselves than those of \mathcal{X} , which is a constant.

⁵⁵ The *spread* or *variation* of the values.

In order to measure the variation of the values of \mathcal{X} , we can try to see how far is \mathcal{X} from its mean μ in average; that is, $E[|\mathcal{X} - \mu|]$. However, dealing with absolute values is often problematic mathematically.⁵⁶ For that reason we consider the expected squared difference between them.

⁵⁶ It requires a case analysis for when \mathcal{X} is larger or smaller than μ .

Definition 43. If \mathcal{X} is a random variable with $E[\mathcal{X}] = \mu$, then the *variance* of \mathcal{X} is $Var(\mathcal{X}) := E[(\mathcal{X} - \mu)^2]$.

Notice that

$$\begin{aligned} Var(\mathcal{X}) &= E[(\mathcal{X} - \mu)^2] = E[\mathcal{X}^2 - 2\mu\mathcal{X} + \mu^2] = E[\mathcal{X}^2] - 2\mu E[\mathcal{X}] + \mu^2 \\ &= E[\mathcal{X}^2] - \mu^2 = E[\mathcal{X}^2] - (E[\mathcal{X}])^2. \end{aligned}$$

This is often a simpler way to compute the variance.

Example 44. Let \mathcal{X} represent the outcome of rolling a fair die. We know that $P[\mathcal{X} = i] = 1/6$ for $1 \leq i \leq 6$. Then,

$$E[\mathcal{X}^2] = \frac{1}{6} \sum_{i=1}^6 i^2 = \frac{1}{6} (1^2 + 2^2 + 3^2 + 4^2 + 5^2 + 6^2) = \frac{91}{6}.$$

We have also seen that $E[\mathcal{X}] = 7/2$. Thus,

$$Var(\mathcal{X}) = E[\mathcal{X}^2] - (E[\mathcal{X}])^2 = \frac{91}{6} - \frac{49}{4} = \frac{35}{12}. \quad \triangle$$

The variance of a linear transformation of \mathcal{X} is also easy to compute. Let $\mu = E[\mathcal{X}]$.⁵⁷ Then,

⁵⁷ Recall that $E[a\mathcal{X} + b] = aE[\mathcal{X}] + b$.

$$\begin{aligned} Var(a\mathcal{X} + b) &= E[(a\mathcal{X} + b - E[a\mathcal{X} + b])^2] = E[(a\mathcal{X} + b - a\mu - b)^2] \\ &= E[(a\mathcal{X} - a\mu)^2] = E[a^2(\mathcal{X} - \mu)^2] = a^2 E[(\mathcal{X} - \mu)^2] \\ &= a^2 Var(\mathcal{X}). \end{aligned}$$

In particular, this means that $Var(b) = 0$ and $Var(\mathcal{X} + b) = Var(\mathcal{X})$ for any constant b .⁵⁸ That is, constants have variance 0, and shifting the values of \mathcal{X} by a constant does not affect its variance.[◦] However, scaling \mathcal{X} by a constant scales quadratically the variance; i.e. $Var(a\mathcal{X}) = a^2 Var(\mathcal{X})$.

⁵⁸ By setting $a = 0$ and $a = 1$, respectively.

[◦] Think why this must be the case.

The value $\sqrt{Var(\mathcal{X})}$, called the *standard deviation* of \mathcal{X} , has the same units as the mean.

Covariance

WE HAVE SEEN THAT the mean of a sum of random variables is equal to the sum of their means. This result does not carry out to variances in general. In fact, we know already that

$$Var(\mathcal{X} + \mathcal{X}) = Var(2\mathcal{X}) = 4Var(\mathcal{X}) \neq Var(\mathcal{X}) + Var(\mathcal{X}).$$

However, if two random variables are independent, then the variance of their sum corresponds to the sum of their variances. We will show this using the *covariance*.

Definition 45. Let \mathcal{X} and \mathcal{Y} be two random variables, and $\mu_{\mathcal{X}} = E[\mathcal{X}]$, $\mu_{\mathcal{Y}} = E[\mathcal{Y}]$. The *covariance* of \mathcal{X} and \mathcal{Y} is

$$\text{Cov}(\mathcal{X}, \mathcal{Y}) = E[(\mathcal{X} - \mu_{\mathcal{X}})(\mathcal{Y} - \mu_{\mathcal{Y}})]$$

Expanding this definition, we obtain

$$\begin{aligned} \text{Cov}(\mathcal{X}, \mathcal{Y}) &= E[\mathcal{X}\mathcal{Y} - \mu_{\mathcal{X}}\mathcal{Y} - \mu_{\mathcal{Y}}\mathcal{X} + \mu_{\mathcal{X}}\mu_{\mathcal{Y}}] \\ &= E[\mathcal{X}\mathcal{Y}] - \mu_{\mathcal{X}}E[\mathcal{Y}] - \mu_{\mathcal{Y}}E[\mathcal{X}] + \mu_{\mathcal{X}}\mu_{\mathcal{Y}} \\ &= E[\mathcal{X}\mathcal{Y}] - \mu_{\mathcal{X}}\mu_{\mathcal{Y}} - \mu_{\mathcal{Y}}\mu_{\mathcal{X}} + \mu_{\mathcal{X}}\mu_{\mathcal{Y}} = E[\mathcal{X}\mathcal{Y}] - E[\mathcal{X}]E[\mathcal{Y}]. \end{aligned}$$

Notice that the covariance is symmetric, and $\text{Cov}(\mathcal{X}, \mathcal{X}) = \text{Var}(\mathcal{X})$. In addition, for any constant a , $\text{Cov}(a\mathcal{X}, \mathcal{Y}) = a\text{Cov}(\mathcal{X}, \mathcal{Y})$. The covariance is also additive.

Proposition 46. $\text{Cov}(\mathcal{X}_1 + \mathcal{X}_2, \mathcal{Y}) = \text{Cov}(\mathcal{X}_1, \mathcal{Y}) + \text{Cov}(\mathcal{X}_2, \mathcal{Y})$.

Proof.

$$\begin{aligned} \text{Cov}(\mathcal{X}_1 + \mathcal{X}_2, \mathcal{Y}) &= E[(\mathcal{X}_1 + \mathcal{X}_2)\mathcal{Y}] - E[\mathcal{X}_1 + \mathcal{X}_2]E[\mathcal{Y}] \\ &= E[\mathcal{X}_1\mathcal{Y}] + E[\mathcal{X}_2\mathcal{Y}] - (E[\mathcal{X}_1] + E[\mathcal{X}_2])E[\mathcal{Y}] \\ &= E[\mathcal{X}_1\mathcal{Y}] - E[\mathcal{X}_1]E[\mathcal{Y}] + E[\mathcal{X}_2\mathcal{Y}] - E[\mathcal{X}_2]E[\mathcal{Y}] \\ &= \text{Cov}(\mathcal{X}_1, \mathcal{Y}) + \text{Cov}(\mathcal{X}_2, \mathcal{Y}). \quad \square \end{aligned}$$

This can be easily generalized to arbitrary sums and, using the symmetry of the covariance, we obtain the following theorem.

Theorem 47. $\text{Cov}(\sum_{i=1}^n \mathcal{X}_i, \sum_{j=1}^m \mathcal{Y}_j) = \sum_{i=1}^n \sum_{j=1}^m \text{Cov}(\mathcal{X}_i, \mathcal{Y}_j)$.

We can now compute the variance of the sum of random variables as:⁵⁹

⁵⁹ Recall that $\text{Cov}(\mathcal{X}, \mathcal{X}) = \text{Var}(\mathcal{X})$.

$$\begin{aligned} \text{Var}\left(\sum_{i=1}^n \mathcal{X}_i\right) &= \text{Cov}\left(\sum_{i=1}^n \mathcal{X}_i, \sum_{i=1}^n \mathcal{X}_i\right) \\ &= \sum_{i=1}^n \sum_{j=1}^n \text{Cov}(\mathcal{X}_i, \mathcal{X}_j) \\ &= \sum_{i=1}^n \left[\sum_{j \neq i} \text{Cov}(\mathcal{X}_i, \mathcal{X}_j) + \text{Cov}(\mathcal{X}_i, \mathcal{X}_i) \right] \\ &= \sum_{i=1}^n \sum_{j \neq i} \text{Cov}(\mathcal{X}_i, \mathcal{X}_j) + \sum_{i=1}^n \text{Cov}(\mathcal{X}_i, \mathcal{X}_i) \\ &= \sum_{i=1}^n \sum_{j \neq i} \text{Cov}(\mathcal{X}_i, \mathcal{X}_j) + \sum_{i=1}^n \text{Var}(\mathcal{X}_i). \end{aligned}$$

If $n = 2$, then this means that

$$\begin{aligned} \text{Var}(\mathcal{X} + \mathcal{Y}) &= \text{Var}(\mathcal{X}) + \text{Var}(\mathcal{Y}) + \text{Cov}(\mathcal{X}, \mathcal{Y}) + \text{Cov}(\mathcal{Y}, \mathcal{X}) \\ &= \text{Var}(\mathcal{X}) + \text{Var}(\mathcal{Y}) + 2\text{Cov}(\mathcal{X}, \mathcal{Y}). \end{aligned}$$

Then we get to the result that we have hinted before.

Theorem 48. If \mathcal{X} and \mathcal{Y} are independent, then $\text{Cov}(\mathcal{X}, \mathcal{Y}) = 0$.⁶⁰

Proof. It suffices to show that $E[\mathcal{X}\mathcal{Y}] = E[\mathcal{X}]E[\mathcal{Y}]$. We show it for the discrete case only.⁶¹

$$\begin{aligned} E[\mathcal{X}\mathcal{Y}] &= \sum_j \sum_i x_i y_j P[\mathcal{X} = x_i, \mathcal{Y} = y_j] \\ &= \sum_j \sum_i x_i y_j P[\mathcal{X} = x_i] P[\mathcal{Y} = y_j] \quad (\dagger) \\ &= \sum_j y_j P[\mathcal{Y} = y_j] \sum_i x_i P[\mathcal{X} = x_i] = E[\mathcal{X}]E[\mathcal{Y}], \end{aligned}$$

where (\dagger) follows by independence. \square

In particular, if $\mathcal{X}_1, \dots, \mathcal{X}_n$ are independent, then

$$\text{Var}\left(\sum_{i=1}^n \mathcal{X}_i\right) = \sum_{i=1}^n \text{Var}(\mathcal{X}_i).$$

Example 49. Compute the variance of the sum of 10 independent rolls of a fair die.

Solution. If \mathcal{X}_i denotes the i -th roll, we get⁶²

$$\text{Var}\left(\sum_{i=1}^{10} \mathcal{X}_i\right) = \sum_{i=1}^{10} \text{Var}(\mathcal{X}_i) = 10\left(\frac{35}{12}\right) = \frac{175}{6}. \quad \triangle$$

⁶⁰ Notice that this is not an equivalence: two variables may have covariance 0, and still not be independent.

⁶¹ All other cases are analogous.

⁶² Recall Example 44.

Intuitively, the covariance describes the relationship between two variables. Consider for example the indicator variables \mathcal{X}, \mathcal{Y} for the events \mathcal{E} and \mathcal{F} , respectively. It follows that

$$\text{Cov}(\mathcal{X}, \mathcal{Y}) = E[\mathcal{X}\mathcal{Y}] - E[\mathcal{X}]E[\mathcal{Y}] = P[\mathcal{X} = 1, \mathcal{Y} = 1] - P[\mathcal{X} = 1]P[\mathcal{Y} = 1].$$

Then, $\text{Cov}(\mathcal{X}, \mathcal{Y}) > 0$ if and only if $P[\mathcal{X} = 1, \mathcal{Y} = 1] > P[\mathcal{X} = 1]P[\mathcal{Y} = 1]$ or, equivalently, $P(\mathcal{Y} = 1 \mid \mathcal{X} = 1) > P[\mathcal{Y} = 1]$. In words, the covariance tells us whether it is more or less likely to observe $[\mathcal{Y} = 1]$ when we know that $\mathcal{X} = 1$.⁶²

⁶² And by symmetry, also the dual.

In general, a positive covariance between two RVs expresses that both variables grow together,⁶³ while a negative covariance says that one decreases while the other increases. The strength of the relationship between the two variables is indicated by their *correlation*, defined as

⁶³ That is, \mathcal{Y} grows as \mathcal{X} grows.

$$\text{Corr}(\mathcal{X}, \mathcal{Y}) = \frac{\text{Cov}(\mathcal{X}, \mathcal{Y})}{\sqrt{\text{Var}(\mathcal{X})\text{Var}(\mathcal{Y})}}.$$

This dimensionless value is always between -1 and 1.




Figure 9: Random variables \mathcal{X} and \mathcal{Y} with correlations (a) 0.75; (b) 0.2; (c) 0; and (d) -0.75 .

Moment Generating Functions

EVERY RANDOM VARIABLE \mathcal{X} has a *moment generating function* (mgf) defined for every $t \in \mathbb{N}$ as

$$\phi(t) := E[e^{t\mathcal{X}}] = \begin{cases} \sum_x e^{tx} p(x) & \text{if } \mathcal{X} \text{ is discrete} \\ \int_x e^{tx} f(x) dx & \text{if } \mathcal{X} \text{ is continuous} \end{cases}$$

This function is called moment generating because all the moments of \mathcal{X} can be obtained by differentiating it and evaluating at zero.⁶⁴ For example,

$$\phi'(t) = \frac{d}{dt} E[e^{t\mathcal{X}}] = E\left[\frac{d}{dt} e^{t\mathcal{X}}\right] = E[\mathcal{X} e^{t\mathcal{X}}].$$

Hence, $\phi'(0) = E[\mathcal{X}]$. Similarly, $\phi''(t) = E[\mathcal{X}^2 e^{t\mathcal{X}}]$ and so $\phi''(0) = E[\mathcal{X}^2]$. In general, the n -th moment of \mathcal{X} is the n -th derivative of ϕ evaluated at $t = 0$.

An important property is that the mgf of the sum of *independent* random variables is the product of their individual moment generating functions. For two random variables, the mgf of $\mathcal{X} + \mathcal{Y}$ is

$$\begin{aligned} \phi_{\mathcal{X}+\mathcal{Y}}(t) &= E[e^{t(\mathcal{X}+\mathcal{Y})}] = E[e^{t\mathcal{X}} e^{t\mathcal{Y}}] \\ &= E[e^{t\mathcal{X}}] E[e^{t\mathcal{Y}}] = \phi_{\mathcal{X}}(t) \phi_{\mathcal{Y}}(t), \end{aligned} \quad (\dagger)$$

where (\dagger) follows because \mathcal{X} and \mathcal{Y} , and hence also $e^{t\mathcal{X}}$ and $e^{t\mathcal{Y}}$ are independent.⁶⁵

Interestingly, mgfs *uniquely* determine the distribution of random variables; that is, there is a one-to-one correspondence between distributions and their moment generating function.

⁶⁴ The n -th moment of \mathcal{X} is $E[\mathcal{X}^n]$.

⁶⁵ In the proof of Theorem 48, we have shown that if two random variables are independent, then $E[\mathcal{X}\mathcal{Y}] = E[\mathcal{X}]E[\mathcal{Y}]$.

The Weak Law of Large Numbers

WE NOW PROVE two important results.

Theorem 50 (Markov's Inequality). *Let \mathcal{X} be a random variable that takes only non-negative values. Then, for every $a > 0$*

$$P[\mathcal{X} \geq a] \leq \frac{E[\mathcal{X}]}{a}.$$

Proof. We prove it for a continuous variable with density f .

$$\begin{aligned} E[\mathcal{X}] &= \int_0^\infty x f(x) dx = \int_0^a x f(x) dx + \int_a^\infty x f(x) dx \\ &\geq \int_a^\infty x f(x) dx \geq \int_a^\infty a f(x) dx = a \int_a^\infty f(x) dx = a P[\mathcal{X} \geq a]. \quad \square \end{aligned}$$

Corollary 51 (Chebyshev's Inequality). *If \mathcal{X} is a random variable with mean μ and variance σ^2 , then for every $k > 0$ we have*

$$P[|\mathcal{X} - \mu| \geq k] \leq \frac{\sigma^2}{k^2}.$$

Proof. Notice that $(\mathcal{X} - \mu)^2$ is a non-negative random variable. So we can apply Markov's inequality, with $a = k^2$ to get

$$P[(\mathcal{X} - \mu)^2 \geq k^2] \leq \frac{E[(\mathcal{X} - \mu)^2]}{k^2}.$$

Since $(\mathcal{X} - \mu)^2 \geq k^2$ holds if and only if $|\mathcal{X} - \mu| \geq k$, this implies

$$P[|\mathcal{X} - \mu| \geq k] \leq \frac{E[(\mathcal{X} - \mu)^2]}{k^2} \leq \frac{\sigma^2}{k^2}. \quad \square$$

The Markov and Chebishev inequalities allow bounding probabilities even if the general distribution of \mathcal{X} is unknown, as long as the mean, and potentially the variance, is known.⁶⁶

⁶⁶ If the distribution is known, the probabilities can usually be computed precisely, without the need to approximate.

Example 52. Suppose that the number of hours that a person works per week is a random variable with mean 40.

1. What can be said about the probability that this week they will work more than 60 hours?
2. If the variance of the working hours per week is 16, what is the probability that this week's work will be between 32 and 48 hours?

Solution. Let \mathcal{X} be the number of working hours in a week.

1. Using Markov's inequality, $P[\mathcal{X} \geq 60] \leq \frac{E[\mathcal{X}]}{60} = \frac{40}{60} = 2/3$.
2. Using Chebyshev's inequality, $P[|\mathcal{X} - 40| \geq 8] \leq \frac{16}{64} = 1/4$. So, we know that $P[|\mathcal{X} - 40| \leq 8] = 1 - P[|\mathcal{X} - 40| \geq 8] \geq 0.75$. \triangle

If we use Chebyshev's inequality with distance $k\sigma$, we get that

$$P[|\mathcal{X} - \mu| \geq k\sigma] \leq \frac{\sigma^2}{k^2\sigma^2} = \frac{1}{k^2}.$$

That is, the probability that \mathcal{X} differs from its mean by at least k standard deviations is bounded by $\frac{1}{k^2}$.⁶⁷

⁶⁷ Decreases quadratically.

A consequence of this inequality is that if one takes several independent and identically distributed random variables, then their average will tend to their mean (with probability 1).⁶⁸

⁶⁸ This is the basis of many statistical analyses, for repeated experiments

Theorem 53 (Weak Law of Large Numbers). *Let $\mathcal{X}_1, \mathcal{X}_2, \dots$ be a sequence of independent, identically distributed RVs with mean $E[\mathcal{X}_i] = \mu$. Then, for any $\epsilon > 0$*

$$\lim_{n \rightarrow \infty} P \left[\left| \frac{\sum_{i=1}^n \mathcal{X}_i}{n} - \mu \right| > \epsilon \right] = 0.$$

Proof. Suppose for simplicity that the random variables have a finite variance σ^2 . Since all the variables are independent, if $\mathcal{Y} = \frac{\sum_{i=1}^n \mathcal{X}_i}{n}$, then

$$E[\mathcal{Y}] = \mu \quad \text{Var}(\mathcal{Y}) = \frac{\sigma^2}{n}.$$

Applying Chebyshev's inequality results in

$$P \left[\left| \frac{\sum_{i=1}^n \mathcal{X}_i}{n} - \mu \right| > \epsilon \right] \leq \frac{\sigma^2}{n\epsilon^2}. \quad \square$$

Suppose for example that we independently repeat a clinical trial. Let \mathcal{E} be a fixed event that has probability $P(\mathcal{E})$ of occurring at each trial. If \mathcal{X}_i is the indicator variable for \mathcal{E} occurring at trial i , then $\sum_{i=1}^n$ is the number of times that \mathcal{E} is observed in the first n trials. Since $E[\mathcal{X}_i] = P(\mathcal{E})$, it follows from the weak law of large numbers that the probability that the proportion of trials in which we observe \mathcal{E} differs from $P(\mathcal{E})$ by more than ε tends to 0 as n grows.[◦]

[◦]Think of the name “large numbers.”

Special Random Variables

WE INTRODUCE SOME OF the most commonly observed and used RVs.

Bernoulli and Binomial

CONSIDER AN EXPERIMENT WHOSE outcome may be a *success* or a *failure*, and \mathcal{X} its indicator variable.⁶⁹ The probability mass function of \mathcal{X} is defined by $P[\mathcal{X} = 1] = p$; its *probability of success*. Such a random variable is called *Bernoulli*⁷⁰ and, as we know already, its expected value is equal to its probability of success p .

If we independently repeat the experiment n times, and \mathcal{X} represents the number of successes observed, then \mathcal{X} is a *binomial* random variable with parameters (n, p) . The probability mass function of this random variable is defined for all $i, 0 \leq i \leq n$ as⁷¹

$$P[\mathcal{X} = i] = \binom{n}{i} p^i (1-p)^{n-i}.$$

The probability mass functions of three binomial random variables with parameters $(10, 0.5)$, $(10, 0.4)$, and $(10, 0.75)$, respectively, are depicted in Figure 10. Notice that two of them lean (or *skew*) away from the center.

Example 54. A system with n components works if at least half of its components function. If each component works, independently, with probability p ,

1. for which values of p will a 5-component system be more reliable than a 3-component one?
2. In general, when is a $(2k+1)$ -component system more reliable than a $(2k-1)$ -component one?

Solution. For the first question, we know that a 5-component system and a 3-component system will work, respectively, with probability

$$\begin{aligned} \binom{5}{3} p^3 (1-p)^2 + \binom{5}{4} p^4 (1-p) + p^5, \\ \binom{3}{2} p^2 (1-p) + p^3. \end{aligned}$$

The 5-component system is more reliable if

$$\binom{5}{3} p^3 (1-p)^2 + \binom{5}{4} p^4 (1-p) + p^5 \geq \binom{3}{2} p^2 (1-p) + p^3$$

⁶⁹ That is, $\mathcal{X} = 1$ if the experiment succeeds, and $\mathcal{X} = 0$ otherwise.

⁷⁰ After James Bernoulli.

⁷¹ Recall that $\binom{n}{r} = \frac{n!}{r!(n-r)!}$.




Figure 10: Three binomial probability mass functions.

or, equivalently, if $3(p-1)^2(2p-1) \geq 0$; that is if $p \geq 1/2$.

In general, consider a system with $2k+1$ components and let \mathcal{X} denote the number of the first $2k-1$ that function. The system will work if any of the following three conditions hold:

1. $\mathcal{X} \geq k+1$;
2. $\mathcal{X} = k$ and at least one of the last two components functions; or
3. $\mathcal{X} = k-1$ and the last two components function.

The probability of this system working is

$$P[\mathcal{X} \geq k+1] + P[\mathcal{X} = k](1 - (1-p)^2) + P[\mathcal{X} = k-1]p^2.$$

Similarly, the probability that the smaller system works is[◦]

$$P[\mathcal{X} = k] + P[\mathcal{X} \geq k+1].$$

This means that the larger system is more reliable if⁷²

$$\begin{aligned} 0 &\leq P_{2k+1}(\text{works}) - P_{2k-1}(\text{works}) \\ &= P[\mathcal{X} = k-1]p^2 - P[\mathcal{X} = k](1-p)^2 \\ &= \binom{2k-1}{k-1} p^{k-1}(1-p)^k p^2 - \binom{2k-1}{k} p^k(1-p)^{k-1}(1-p)^2 \\ &= \binom{2k-1}{k} p^k(1-p)^k(p - (1-p)) = \binom{2k-1}{k} p^k(1-p)^k(2p-1), \quad (\dagger) \end{aligned}$$

where (\dagger) follows from $\binom{2k-1}{k} = \binom{2k-1}{k-1}$. Thus, the larger system is more reliable if and only if $p \geq 1/2$.⁷³ \triangle

Recall that a binomial random variable \mathcal{X} represents the number of successes of n independent trials, where each trial has probability p of occurring. Then $\mathcal{X} = \sum_{i=1}^n \mathcal{X}_i$, where the \mathcal{X}_i s are independent Bernoulli random variables. In particular,⁷⁴

$$\begin{aligned} E[\mathcal{X}_i] &= P[\mathcal{X}_i = 1] = p \\ \text{Var}(\mathcal{X}_i) &= E[\mathcal{X}_i^2] - (E[\mathcal{X}_i])^2 = p - p^2 = p(1-p). \end{aligned}$$

From this, it immediately follows that

$$\begin{aligned} E[\mathcal{X}] &= \sum_{i=1}^n E[\mathcal{X}_i] = np \\ \text{Var}(\mathcal{X}) &= \sum_{i=1}^n \text{Var}(\mathcal{X}_i) = np(1-p). \end{aligned}$$

In addition, if \mathcal{X}_1 and \mathcal{X}_2 are two independent binomial random variables with parameters (n_i, p) for $i = 1, 2$, then their sum is also a binomial with parameters $(n_1 + n_2, p)$.⁷⁵

TO COMPUTE THE DISTRIBUTION function of a binomial \mathcal{X} with parameters (n, p) it is helpful to use the equation

$$P[\mathcal{X} = k+1] = \frac{p}{1-p} \frac{n-k}{k+1} P[\mathcal{X} = k].$$

[◦]By independence.

⁷² The probability of the larger system to work is larger than the one of the smaller system.

⁷³ For any number N , $N(2p-1) \geq 0$ iff $2Np \geq N$.

⁷⁴ Notice that $\mathcal{X}^2 = \mathcal{X}$ because \mathcal{X} can only take values 0 or 1.

⁷⁵ It is important that the probability in both binomials is the same. \mathcal{X}_i denotes the number of successes in n_i independent trials. So their sum corresponds to the successes in the sum of their (independently made) trials.

Example 55. Let \mathcal{X} be a binomial with $n = 5$ and $p = 0.3$. Then, starting with $P[\mathcal{X} = 0] = (0.7)^5 = 0.168$ we obtain

$$\begin{aligned} P[\mathcal{X} = 1] &= \frac{3}{7} \frac{5}{1} P[\mathcal{X} = 0] \\ P[\mathcal{X} = 2] &= \frac{3}{7} \frac{4}{2} P[\mathcal{X} = 1] \\ P[\mathcal{X} = 3] &= \frac{3}{7} \frac{3}{3} P[\mathcal{X} = 2] \\ P[\mathcal{X} = 4] &= \frac{3}{7} \frac{2}{4} P[\mathcal{X} = 3] \\ P[\mathcal{X} = 5] &= \frac{3}{7} \frac{1}{5} P[\mathcal{X} = 4]. \end{aligned}$$

△

It is obviously possible to generalize the Bernoulli and binomial random variables to allow more than two values to be taken. In this case, we would speak of a *multinomial* RV. Understanding the properties of such RVs is left as an exercise to the interested student.

Poisson

A POISSON RANDOM VARIABLE with parameter $\lambda > 0$ takes values in the natural numbers, and has the probability mass function, for every $i \in \mathbb{N}$,⁷⁶

$$P[\mathcal{X} = i] = e^{-\lambda} \frac{\lambda^i}{i!};$$

see Figure 11.⁷⁷

To determine the mean and variance of a Poisson random variable, we compute its moment generating function, and its two first derivatives.

$$\begin{aligned} \phi(t) &= E[e^{t\mathcal{X}}] = \sum_{i=0}^{\infty} e^{ti} e^{-\lambda} (\lambda^i / i!) \\ &= e^{-\lambda} \sum_{i=0}^{\infty} (\lambda e^t)^i / i! \\ &= e^{-\lambda} e^{\lambda e^t} = \exp[\lambda(e^t - 1)] \\ \phi'(t) &= \lambda e^t \exp[\lambda(e^t - 1)] \\ \phi''(t) &= (\lambda e^t)^2 \exp[\lambda(e^t - 1)] + \lambda e^t \exp[\lambda(e^t - 1)] \end{aligned}$$

From these equations, we deduce

$$\begin{aligned} E[\mathcal{X}] &= \phi'(0) = \lambda \\ \text{Var}(\mathcal{X}) &= \phi''(0) - (E[\mathcal{X}])^2 = \lambda^2 + \lambda - \lambda^2 = \lambda. \end{aligned}$$

That is, the mean and the variance of a Poisson random variable are both equal to the parameter λ .

Poisson random variables provide good approximations for binomial variables with parameters (n, p) if n is large and p is small. Suppose that \mathcal{X} is such a binomial random variable and let $\lambda = np$. Then

$$\begin{aligned} P[\mathcal{X} = i] &= \frac{n!}{(n-i)!i!} p^i (1-p)^{n-i} = \frac{n!}{(n-i)!i!} \left(\frac{\lambda}{n}\right)^i \left(1 - \frac{\lambda}{n}\right)^{n-i} \\ &= \frac{n(n-1)\cdots(n-i+1)}{n^i} \frac{\lambda^i}{i!} \left(1 - \frac{\lambda}{n}\right)^n. \end{aligned}$$

Assuming that n is large and p is small, we can deduce that⁷⁸

$$\left(1 - \frac{\lambda}{n}\right)^n \approx e^{-\lambda}; \quad \left(1 - \frac{\lambda}{n}\right)^i \approx 1; \quad \frac{n(n-1)\cdots(n-i+1)}{n^i} \approx 1.$$

And hence it follows that $P[\mathcal{X} = i] \approx e^{-\lambda} \frac{\lambda^i}{i!}$.⁷⁹

⁷⁶ First defined by S. D. Poisson.

⁷⁷ To see that this is a mass function, recall that for all $x \in \mathbb{R}$, $\sum_{i=1}^{\infty} \frac{x^i}{i!} = e^x$.




Figure 11: Poisson mass function for $\lambda = 4$.

⁷⁸ Recall that $\lim_{n \rightarrow \infty} (1 + x/n)^n = e^x$.

⁷⁹ If a very large number n of independent trials with a very low probability p of success are performed, then the number of successes is approximately a Poisson with $\lambda = np$.

Example 56. Suppose that, in average, there are three accidents in the highway between Trento and Bolzano every week. Compute the probability that there is at least one accident this week.

Solution. Let \mathcal{X} denote the number of accidents this week. It is reasonable to assume that there is a large number of cars passing, each with a very low probability of having an accident. So \mathcal{X} should be approximately Poisson distributed. Thus,

$$P[\mathcal{X} \geq 1] = 1 - P[\mathcal{X} = 0] \approx 1 - e^{-3} \frac{3^0}{0!} = 1 - e^{-3} \approx 0.95. \quad \triangle$$

THE POISSON APPROXIMATION REMAINS valid in more general circumstances. If n independent trials are performed, each with a probability of success p_i , $1 \leq i \leq n$, then if n is large and *each* p_i is small, the number of successful trials is approximately a Poisson with mean $\sum_{i=1}^n p_i$. This remains true without the independence assumption, as long as the dependence is not very strong.

Example 57. Consider n people that leave their umbrella at the entrance of a bar, and pick up one of the umbrellas randomly when leaving. If \mathcal{X} denotes the number of people that take their own umbrella, then for large n \mathcal{X} approximates a Poisson distribution with mean 1. Intuitively, this is the case because we can express $\mathcal{X} = \sum_{i=1}^n \mathcal{X}_i$, where each \mathcal{X}_i is the indicator variable for person i picking their own umbrella. Since each person is equally likely to pick any umbrella, we have that $P[\mathcal{X} = i] = 1/n$.⁸⁰ Moreover $E[\mathcal{X}] = \sum_{i=1}^n E[\mathcal{X}_i] = n(1/n) = 1$.

However, if the j -th person picks up their own umbrella, then the i -th person ($i \neq j$) is equally likely to pick any of the remaining $n-1$ umbrellas; more precisely, $P[\mathcal{X}_i = 1 \mid \mathcal{X}_j = 1] = \frac{1}{n-1}$. This means that \mathcal{X}_i and \mathcal{X}_j are not independent, but their dependence is extremely weak.⁸¹ \triangle

Poisson random variables are also *reproductive*; that is, the sum of two independent Poisson random variables \mathcal{X} and \mathcal{Y} is also Poisson. To prove this, consider the moment generating function of $\mathcal{X} + \mathcal{Y}$ ⁸²

$$\begin{aligned} E[e^{t(\mathcal{X}+\mathcal{Y})}] &= E[e^{t\mathcal{X}} e^{t\mathcal{Y}}] = E[e^{t\mathcal{X}}] E[e^{t\mathcal{Y}}] \\ &= \exp(\lambda_{\mathcal{X}}(e^t - 1)) \exp(\lambda_{\mathcal{Y}}(e^t - 1)) = \exp((\lambda_{\mathcal{X}} + \lambda_{\mathcal{Y}})(e^t - 1)). \end{aligned}$$

Since this is the mgf of a Poisson with mean $\lambda_{\mathcal{X}} + \lambda_{\mathcal{Y}}$, it follows that $\mathcal{X} + \mathcal{Y}$ is indeed a Poisson.⁸³

Example 58. The number of hourly customers at a bar is a Poisson with mean 4. What is the probability that over a 2-hour period there are no more than 3 customers?

Solution. Let \mathcal{X}_i , $i = 1, 2$ be the number of customers during the i -th hour. Assuming that \mathcal{X}_1 and \mathcal{X}_2 are independent, then $\mathcal{X} = \mathcal{X}_1 + \mathcal{X}_2$ is Poisson with mean 8. Thus,

$$P[\mathcal{X}_1 + \mathcal{X}_2 \leq 3] = \sum_{n=0}^3 e^{-8} \frac{8^n}{n!} = 0.423. \quad \triangle$$

⁸⁰ That is, we can approximately simulate the experiment by a binomial with probability $1/n$.

⁸¹ Specially when n tends to be large.

⁸² Assume that the means of \mathcal{X} and \mathcal{Y} are $\lambda_{\mathcal{X}}$ and $\lambda_{\mathcal{Y}}$, respectively.

⁸³ Recall that there is a one-to-one correspondence between moment generating functions and distributions.

CONSIDER NOW A SCENARIO in which a random number N of events will occur, and each of them is independently of type 1 or type 2 with probabilities p and $1 - p$, respectively.⁸⁴ Let $N_i, i = 1, 2$ denote the number of events of type i observed. If N is a Poisson with mean λ then the joint mass function of N_1, N_2 is

$$\begin{aligned} P[N_1 = n, N_2 = m] &= P[N_1 = n, N_2 = m, N = n + m] \\ &= P[N_1 = n, N_2 = m \mid N = n + m] P[N = n + m] \\ &= P[N_1 = n, N_2 = m \mid N = n + m] e^{-\lambda} \frac{\lambda^{n+m}}{(n+m)!}. \end{aligned}$$

Since each of the $n + m$ events are independently of type 1 with probability p , the probability that there are *exactly* n events of type 1 is the probability that the binomial random variable $(n + m, p)$ takes value n . That is,

$$\begin{aligned} P[N_1 = n, N_2 = m] &= \frac{(n+m)!}{n!m!} p^n (1-p)^m e^{-\lambda} \frac{\lambda^{n+m}}{(n+m)!} \\ &= e^{-\lambda p} \frac{(\lambda p)^n}{n!} e^{-\lambda(1-p)} \frac{(\lambda(1-p))^m}{m!}. \end{aligned}$$

The probability mass function of N_1 is[◦]

$$\begin{aligned} P[N_1 = n] &= \sum_{m=0}^{\infty} P[N_1 = n, N_2 = m] \\ &= e^{-\lambda p} \frac{(\lambda p)^n}{n!} \sum_{m=0}^{\infty} e^{-\lambda(1-p)} \frac{(\lambda(1-p))^m}{m!} = e^{-\lambda p} \frac{(\lambda p)^n}{n!}. \end{aligned}$$

Thus, N_1 is a Poisson with mean λp , and similarly N_2 is Poisson with mean $\lambda(1-p)$. Moreover, these two variables are independent.

This result can be generalised to the case where each event can take any of r different categories with probabilities p_1, \dots, p_r . That is, the numbers of type i events ($1 \leq i \leq r$) are independent Poisson random variables with mean λp_i .

To compute the distribution of a Poisson \mathcal{X} with mean λ we notice that

$$\frac{P[\mathcal{X} = i+1]}{P[\mathcal{X} = i]} = \frac{e^{-\lambda} \lambda^{i+1} / (i+1)!}{e^{-\lambda} \lambda^i / i!} = \frac{\lambda}{i+1}.$$

Starting from $P[\mathcal{X} = 0] = e^{-\lambda}$, we can use this equation to successively compute the probability of each successive value.⁸⁵

⁸⁴ For example, the number of customers N in a bar is a random variable; each customer is male with probability p and female with probability $1 - p$.

[◦] Recall that $\sum_{i=1}^{\infty} \frac{x^i}{i!} = e^x$.

⁸⁵ $P[\mathcal{X} = i+1] = (\lambda/(i+1))P[\mathcal{X} = i]$.

Hypergeometric Random Variables

CONSIDER A BIN WITH N working batteries, and M defective ones. We randomly pick up a sample of size n .⁸⁶ If \mathcal{X} is the number of working batteries in the sample, then for all $i, 0 \leq i \leq \min(N, m)$ ⁸⁷

$$P[\mathcal{X} = i] = \frac{\binom{N}{i} \binom{M}{n-i}}{\binom{N+M}{n}}.$$

A random variable with such a mass function is called *hypergeometric* with parameters N, M, n .

⁸⁶ Any of the $\binom{N+M}{n}$ such samples is equally likely.

⁸⁷ To simplify, we assume $\binom{m}{r} = 0$ whenever $r > m$ or $r < 0$.

An intuitive way to see a hypergeometric random variable \mathcal{X} is to think of the sample to be drawn sequentially and defining the random variables

$$\mathcal{X}_i = \begin{cases} 1 & i\text{-th selection is working} \\ 0 & \text{otherwise.} \end{cases}$$

Each selection is equally likely to be any of the $N + M$ batteries; hence $P[\mathcal{X}_i = 1] = \frac{N}{N+M}$. In addition, for any $i \neq j$,⁸⁸

$$P[\mathcal{X}_i = 1, \mathcal{X}_j = 1] = P[\mathcal{X}_i = 1 | \mathcal{X}_j = 1]P[\mathcal{X}_j = 1] = \frac{N-1}{N+M-1} \frac{N}{N+M}.$$

Since $\mathcal{X} = \sum_{i=1}^n \mathcal{X}_i$, it follows that

$$\begin{aligned} E[\mathcal{X}] &= \sum_{i=1}^n E[\mathcal{X}_i] = \sum_{i=1}^n P[\mathcal{X}_i = 1] = \frac{nN}{N+M}, \\ \text{Var}(\mathcal{X}) &= \sum_{i=1}^n \text{Var}(\mathcal{X}_i) + 2 \sum_{1 \leq i < j \leq n} \text{Cov}(\mathcal{X}_i, \mathcal{X}_j). \end{aligned}$$

Since each \mathcal{X}_i is a Bernoulli, $\text{Var}(\mathcal{X}_i) = P[\mathcal{X}_i = 1](1 - P[\mathcal{X}_i = 1]) = \frac{NM}{(N+M)^2}$, and for $i < j$ ⁸⁹ $E[\mathcal{X}_i \mathcal{X}_j] = P[\mathcal{X}_i = 1, \mathcal{X}_j = 1] = \frac{N(N-1)}{(N+M)(N+M-1)}$. Thus,

⁸⁸ Once that we have observed a working battery, another selection can be any of the remaining $N - 1$ working batteries.

⁸⁹ $\mathcal{X}_i \mathcal{X}_j = 1$ iff $\mathcal{X}_i = 1 = \mathcal{X}_j$.

$$\begin{aligned} \text{Cov}(\mathcal{X}_i, \mathcal{X}_j) &= E[\mathcal{X}_i \mathcal{X}_j] - E[\mathcal{X}_i]E[\mathcal{X}_j] \\ &= \frac{N(N-1)}{(N+M)(N+M-1)} - \left(\frac{N}{N+M} \right)^2 \\ &= \frac{-NM}{(N+M)^2(N+M-1)}. \end{aligned}$$

Since there are $\binom{n}{2}$ terms in the sum of covariances, it follows that

$$\begin{aligned} \text{Var}(\mathcal{X}) &= \frac{nNM}{(N+M)^2} - \frac{n(n-1)NM}{(N+M)^2(N+M-1)} \\ &= \frac{nNM}{(N+M)^2} \left[1 - \frac{n-1}{N+M-1} \right]. \end{aligned}$$

Let now $p = N/(N+M)$ be the proportion of batteries that are functioning. It then follows that $E[\mathcal{X}] = np$ and $\text{Var}(\mathcal{X}) = np(1-p) \left[1 - \frac{n-1}{N+M-1} \right]$.

If $N + M$ tends to infinity while preserving the same proportion p , then $\text{Var}(\mathcal{X})$ converges to $np(1-p)$, which is the variance of a binomial with parameters (n, p) .[◦]

[◦] Is this reasonable to you?

Let \mathcal{X} and \mathcal{Y} be two independent binomial random variables with parameters (n, p) and (m, p) , respectively. The conditional mass function of \mathcal{X} given that $\mathcal{X} + \mathcal{Y} = k$ is⁹⁰

$$\begin{aligned} P[\mathcal{X} = i | \mathcal{X} + \mathcal{Y} = k] &= \frac{P[\mathcal{X} = i, \mathcal{Y} = k-i]}{P[\mathcal{X} + \mathcal{Y} = k]} = \frac{P[\mathcal{X} = i]P[\mathcal{Y} = k-i]}{P[\mathcal{X} + \mathcal{Y} = k]} \\ &= \frac{\binom{n}{i} p^i (1-p)^{n-i} \binom{m}{k-i} p^{k-i} (1-p)^{m-(k-i)}}{\binom{n+m}{k} p^k (1-p)^{n+m-k}} \\ &= \frac{\binom{n}{i} \binom{m}{k-i}}{\binom{n+m}{k}}; \end{aligned}$$

⁹⁰ Recall that the sum of two binomials (n, p) and (m, p) is a binomial $(n+m, p)$.

that is, this conditional distribution is a hypergeometric.

Uniform Random Variables

A RANDOM VARIABLE IS *uniformly distributed* over the interval $[\alpha, \beta]$ if its density function is (see Figure 13)

$$f(x) = \begin{cases} \frac{1}{\beta - \alpha} & \alpha \leq x \leq \beta \\ 0 & \text{otherwise.} \end{cases}$$

The uniform distribution expresses that a random variable is equally likely to be near any value within $[\alpha, \beta]$. For a sub-interval of $I \subseteq [\alpha, \beta]$, the probability of \mathcal{X} taking a value in I is the proportional size of I w.r.t. $\beta - \alpha$:

$$P[a < \mathcal{X} < b] = \frac{1}{\beta - \alpha} \int_a^b dx = \frac{b - a}{\beta - \alpha}.$$

For example, if \mathcal{X} is uniformly distributed over the interval $[0, 10]$, then $P[\mathcal{X} > 6] = 0.4$ and $P[2 < \mathcal{X} < 5] = 0.3$.

Intuitively, the mean of a uniform random variable should appear at the middle of its interval.^o We confirm that this is the case:

$$E[\mathcal{X}] = \int_{\alpha}^{\beta} \frac{x}{\beta - \alpha} dx = \frac{\beta^2 - \alpha^2}{2(\beta - \alpha)} = \frac{(\beta + \alpha)(\beta - \alpha)}{2(\beta - \alpha)} = \frac{\alpha + \beta}{2}.$$

To compute the variance, we use⁹¹

$$E[\mathcal{X}^2] = \frac{1}{\beta - \alpha} \int_{\alpha}^{\beta} x^2 dx = \frac{\beta^3 - \alpha^3}{3(\beta - \alpha)} = \frac{\beta^2 + \alpha\beta + \alpha^2}{3},$$

which yields

$$\text{Var}(\mathcal{X}) = \frac{\beta^2 + \alpha\beta + \alpha^2}{3} - \left(\frac{\alpha + \beta}{2}\right)^2 = \frac{\alpha^2 + \beta^2 - 2\alpha\beta}{12} = \frac{(\beta - \alpha)^2}{12}.$$

The value of a uniform $(0, 1)$ random variable is called a *random number*. Modern computers use mathematical methods to generate sequences of independent (pseudo-)random numbers.⁹² Random numbers are used often in clinical trials, for example in what are called *double-blind* tests.⁹³

THE NOTION OF UNIFORM random variables can be extended to joint distributions. The joint probability distribution of \mathcal{X}, \mathcal{Y} is uniform over a region R with area a if $f(x, y) = 1/a$ whenever $(x, y) \in R$, and 0 otherwise. For example, if R is the rectangular region between (α_1, α_2) and (β_1, β_2) , then it is possible to show that \mathcal{X} and \mathcal{Y} are independent uniform distributions over $[\alpha_1, \beta_1]$ and $[\alpha_2, \beta_2]$, respectively.⁹⁴

Exponential Random Variables

A CONTINUOUS RANDOM VARIABLE with density function

$$f(x) = \begin{cases} \lambda e^{-\lambda x} & x \geq 0 \\ 0 & x < 0, \end{cases}$$




Figure 12: Uniform density function of a uniform distribution over $[\alpha, \beta]$.




Figure 13: Probability of an interval in a uniform distribution.

^oDo you understand why?

⁹¹ $x^3 - y^3 = (x^2 + xy + y^2)(x - y)$

⁹² For a *true* random number generator, visit <http://www.random.org>.

⁹³ In a double blind test, a group of volunteers for a trial is divided in two subgroups (of the same size); one groups is given the treatment, and the other one is given a placebo. In this way, it is possible to determine if (and to what extent) the treatment is effective. Random numbers are used to guarantee that the division of the groups is really random, and not biased by some potentially hidden factor.

⁹⁴ To show this, simply compute the individual distributions of \mathcal{X} and \mathcal{Y} .

for some constant λ is called *exponential* (or *exponentially distributed*), and λ is often called the *rate* of the distribution. The cumulative distribution function of such a variable is, for every $x \geq 0$,

$$F(x) = P[\mathcal{X} \leq x] = \int_0^x \lambda e^{-\lambda y} dy = 1 - e^{-\lambda x}.$$

This function often arises as the distribution of the amount of time until some event occurs.⁹⁵

The moment generating function of an exponential is

$$\phi(t) = E[e^{t\mathcal{X}}] = \int_0^\infty e^{tx} \lambda e^{-\lambda x} dx = \lambda \int_0^\infty e^{-(\lambda-t)x} dx = \frac{\lambda}{\lambda-t}, \quad t < \lambda$$

Differentiating, we get

$$\begin{aligned} \phi'(t) &= \frac{\lambda}{(\lambda-t)^2} \\ \phi''(t) &= \frac{2\lambda}{(\lambda-t)^3}, \end{aligned}$$

which implies $E[\mathcal{X}] = \phi'(0) = 1/\lambda$, and $Var(\mathcal{X}) = \phi''(0) - 1/\lambda^2 = 1/\lambda^2$.

The key property of the exponential distribution is that it is *memoryless*; that is, for all $s, t > 0$ it holds that $P[\mathcal{X} > s+t \mid \mathcal{X} > t] = P[\mathcal{X} > s]$.⁹⁶ To show this, notice that the notion of memory less is equivalent to stating that $\frac{P[\mathcal{X} > s+t, \mathcal{X} > t]}{P[\mathcal{X} > t]} = P[\mathcal{X} > s]$ or, equivalently, that

$$P[\mathcal{X} > s+t] = P[\mathcal{X} > t]P[\mathcal{X} > s].$$

If \mathcal{X} is an exponential random variable, this equation is satisfied because $e^{-\lambda(s+t)} = e^{-\lambda s}e^{-\lambda t}$. Interestingly, the exponential is the only type of random variable that is memoryless.

Example 59. Consider 3 identical machines that work for an exponentially distributed amount of time with parameter λ . We use two of them until one breaks, and replace that one with the unused one (call it U). What is the probability that the next machine to break is U?

Solution. Due to the memoryless property, when U starts working the remaining lifetime of the originally working machine still available is equivalent to that of U. Hence, the probability of the next one to break to be U is exactly 0.5.

Proposition 60. If $\mathcal{X}_1, \dots, \mathcal{X}_n$ are independent exponential random variables with parameters $\lambda_1, \dots, \lambda_n$, then $\min\{\mathcal{X}_1, \dots, \mathcal{X}_n\}$ is an exponential with parameter $\sum_{i=1}^n \lambda_i$.

Proof. The proposition follows because⁹⁷

$$\begin{aligned} P[\min\{\mathcal{X}_1, \dots, \mathcal{X}_n\} > x] &= P[\mathcal{X}_1 > x, \dots, \mathcal{X}_n > x] = \prod_{i=1}^n P[\mathcal{X}_i > x] \\ &= \prod_{i=1}^n e^{-\lambda_i x} = e^{-\sum_{i=1}^n \lambda_i x}. \quad \square \end{aligned}$$

This property is useful, for example, if we are interested in understanding when will the first component of a system will fail. Another useful property is that if \mathcal{X} is exponential with parameter λ , then $c\mathcal{X}$ is exponential with parameter λ/c . Indeed,

$$P[c\mathcal{X} \leq x] = P[\mathcal{X} \leq x/c] = 1 - e^{-\lambda x/c}.$$

⁹⁵ For example, the time until it rains, or until you receive a new email tend to be exponentially distributed.

⁹⁶ Intuitively, under the condition that the event has not been observed at time t , the probability of having to wait at least s more is exactly the same as having to wait s at the beginning. Conditioning restarts the clock.

⁹⁷ If \mathcal{X} is exponential with parameter λ , then $F(x) = P[\mathcal{X} \leq x] = 1 - e^{-\lambda x}$.

Normal Random Variables

A RANDOM VARIABLE IS *normally distributed* with parameters μ and σ^2 (denoted as $\mathcal{X} \sim \mathcal{N}(\mu, \sigma^2)$) if it has density

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/2\sigma^2}.$$

This is a bell-shaped symmetric curve that has its highest value ($1/\sqrt{2\pi}\sigma$) at μ (see Figure 14).

The importance of the normal distribution and its use in statistics arises from the *central limit theorem* that, in a nutshell, states that when the number of observations is large enough, random phenomena tend to approximate a normal distribution.⁹⁸ Examples where this can be observed empirically are the heights of people, or the measuring error in physical quantities.

Intuitively, the mean and variance of a normal RV $\mathcal{X} \sim \mathcal{N}(\mu, \sigma^2)$ should be μ and σ^2 , respectively. First, we see that $E[\mathcal{X} - \mu] = 0$,⁹⁹ by defining a new variable $y = (x - \mu)/\sigma$.¹⁰⁰

$$\begin{aligned} E[\mathcal{X} - \mu] &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \frac{(x - \mu)}{\sigma} e^{-(x-\mu)^2/2\sigma^2} dx \\ &= \frac{\sigma}{\sqrt{2\pi}} \int_{-\infty}^{\infty} y e^{-y^2/2} dy \\ &= \frac{\sigma}{\sqrt{2\pi}} \left(-e^{-y^2/2} \right) \Big|_{-\infty}^{\infty} = 0. \end{aligned} \quad (*)$$

To compute the variance, we use $u = y$ and $\frac{dv}{dy} = y e^{-y^2/2}$ to see that¹⁰¹

$$\int_{-\infty}^{\infty} y^2 e^{-y^2/2} dy = -y e^{-y^2/2} \Big|_{-\infty}^{\infty} + \int_{-\infty}^{\infty} e^{-y^2/2} dy = \int_{-\infty}^{\infty} e^{-y^2/2} dy.$$

Then, it follows that

$$\begin{aligned} \text{Var}(\mathcal{X}) &= E[(\mathcal{X} - \mu)^2] = \int_{-\infty}^{\infty} \frac{(x - \mu)^2}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/2\sigma^2} dx \\ &= \sigma^2 \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} y^2 e^{-y^2/2} dy = \sigma^2 \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} dy = \sigma^2, \end{aligned}$$

where the last equation follows from $\frac{1}{\sqrt{2\pi}} e^{-y^2/2}$ being the density function of a normal (with parameters $\mu = 0$ and $\sigma = 1$); so the integral over all reals must be 1.

An important property is that any linear transformation $a\mathcal{X} + b$ of a normal random variable \mathcal{X} is distributed as a normal too.¹⁰² This means that the variable $Z = \frac{\mathcal{X} - \mu}{\sigma}$ is in fact a normal random variable with mean 0 and variance 1. This is called the *standard* or *unit* normal RV and its distribution is denoted by Φ .¹⁰³ We can simplify statements about \mathcal{X} in terms of Z . For example, knowing that $\mathcal{X} < b$ iff $\frac{\mathcal{X} - \mu}{\sigma} < \frac{b - \mu}{\sigma}$ we get that

$$P[\mathcal{X} < b] = P\left[\frac{\mathcal{X} - \mu}{\sigma} < \frac{b - \mu}{\sigma}\right] = \Phi\left(\frac{b - \mu}{\sigma}\right).$$

In other words, for dealing with normal distributions, it suffices to know the values of Φ . There are many software systems and tables providing this




Figure 14: The normal density function.

⁹⁸ This will be covered more extensively in the next chapter.

⁹⁹ Remember that $E[\mathcal{X} + b] = E[\mathcal{X}] + b$.

¹⁰⁰ Using u -substitutions, $\sigma dy = dx$.

¹⁰¹ Recall the *integration by parts* formula: $\int u \frac{dv}{dx} dx = uv - \int v \frac{du}{dx} dx$.

¹⁰² As known already, with mean $a\mu + b$ and variance $a^2\sigma^2$.

¹⁰³ That is, $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-y^2/2} dy$.

information.¹⁰⁴ These tables often provide values only for positive values of x , but since Z is symmetric, negative values can be handled easily. For example, $P[Z < -x] = P[X > x] = 1 - \Phi(x)$. Thus, if $x = 1$, we get that $P[Z < -1] = 1 - \Phi(1) = 1 - 0.8413 = 0.1587$.

Example 61. Data transmission is subject to channel noise disturbances. To reduce the the possibility of error, we encode binary messages using values $-2, 2$, which stand for the original 0 and 1, respectively. Due to noise, when we submit a message $x \in \{-2, 2\}$, it reaches its destination as $R = x + N$. The recipient then decodes the message by concluding, if $R \geq 0.5$ that the message sent was 1, and 0 otherwise. We consider that N has a standard normal distribution.

The probabilities of receiving a wrong signal, due to the noise are

$$\begin{aligned} P[R < 0.5 \mid \text{signal} = 1] &= P[N < -1.5] = 1 - \Phi(1.5) = 0.0668 \\ P[R \geq 0.5 \mid \text{signal} = 0] &= P[N \geq 2.5] = 1 - \Phi(2.5) = 0.0062. \quad \triangle \end{aligned}$$

THE MOMENT GENERATING FUNCTION of a standard normal RV Z is

$$\begin{aligned} E[e^{tZ}] &= \int_{-\infty}^{\infty} e^{tx} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-(x^2-2tx)/2} dx \\ &= e^{-t^2/2} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-(x-t)^2/2} dx = e^{-t^2/2} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} dy = e^{-t^2/2}. \end{aligned}$$

A normal random variable \mathcal{X} with mean μ and variance σ^2 can be expressed as $\sigma Z + \mu$. Thus, its moment generating function is

$$E[e^{t\mathcal{X}}] = E[e^{t\mu + t\sigma Z}] = e^{t\mu} E[e^{t\sigma Z}] = e^{t\mu} e^{-(t\sigma)^2/2} = e^{t\mu - \sigma^2 t^2/2}.$$

We can use this moment generating function to prove that the sum of independent normal random variables is a normal random variable.^o

Example 62. The height of European males is a normal RV with mean 177.6cm and standard deviation 4cm. If a person has two (adult) sons, find the probability that the older one is taller than the younger one by at least 2cm, assuming that the heights for each child are independent.

Solution. Let \mathcal{X}_1 and \mathcal{X}_2 be the heights of the first and second child, respectively. Since $-\mathcal{X}_2$ is a normal with mean -177.6 and variance $4^2 = 16$, $\mathcal{X}_1 - \mathcal{X}_2$ is a normal with mean 0 and variance 32. Then

$$\begin{aligned} P[\mathcal{X}_1 > \mathcal{X}_2 + 2] &= P[\mathcal{X}_1 - \mathcal{X}_2 > 2] = P\left[\frac{\mathcal{X}_1 - \mathcal{X}_2}{\sqrt{32}} > \frac{2}{\sqrt{32}}\right] \\ &= P[Z > 0.3536] \approx 1 - 0.6368 = 0.3632. \quad \triangle \end{aligned}$$

Given any number $\alpha \in (0, 1)$, let z_α be the value such that $P[Z \geq z_\alpha] = \alpha$ (see Figure 15). For example, we have that $z_{0.05} = 1.645$ and $z_{0.01} = 2.33$. The value z_α is the $100(1 - \alpha)$ percentile of Z .¹⁰⁵

Chi-Square

THE SUM OF n independent standard RVs $\mathcal{X} = \sum_{i=1}^n Z_i^2$ is a *chi-square RV* with n degrees of freedom. This is denoted by $\mathcal{X} \sim \chi_n^2$.

¹⁰⁴ See the OLE for one such table. We will see also how to obtain it from R.

^oTry it yourself. What will be the mean and the variance of the sum of two independent standard normal RVs?




Figure 15: $P[Z > z_\alpha] = \alpha$.

¹⁰⁵ The name comes because $100(1 - \alpha)$ percent of the time Z will fall below z_α .

The chi-square distribution is *additive* in the sense that if \mathcal{X} and \mathcal{Y} are independent chi-square with n and m degrees of freedom, respectively, then $\mathcal{X} + \mathcal{Y}$ is chi-square with $n + m$ degrees of freedom. This holds because $\mathcal{X} + \mathcal{Y}$ is the sum of $n + m$ squared independent standard RVs.

If \mathcal{X} is chi-square with n degrees of freedom, then for any $\alpha \in (0, 1)$ the value $\chi_{\alpha,n}^2$ is the value such that $P[\mathcal{X} \geq \chi_{\alpha,n}^2] = \alpha$.

One application where the chi-square distribution may be of interest is when dealing with measurement errors in multiple dimensions. If we have sensors measuring the position of an object in each dimension, then the square of the distance from the measured to the real value behaves as a chi-square, assuming that the measuring error is normal.¹⁰⁶

The t -Distribution

IF A STANDARD NORMAL random variable Z and a chi-square random variable with n degrees of freedom χ_n^2 are independent, then the RV

$$T_n := \frac{Z}{\sqrt{\chi_n^2/n}}$$

has a t -distribution with n degrees of freedom. Figure 16 shows the density function of T_n for different degrees of freedom.

The t -density is symmetric on 0, and approximates the standard normal density as n grows. Recall that χ_n^2 is the sum of the squares of n independent standard normal random variables. From the weak law of large numbers, it follows that as n grows, χ_n^2/n will tend with probability 1 to $E[Z_i^2] = 1$. Thus, for large enough n , $T_n = Z/\sqrt{\chi_n^2/n}$ will approximately have the same distribution as Z (see Figure 17).

The mean and variance of T_n are, for $n > 1$ and $n > 2$,

$$E[T_n] = 0$$

$$\text{Var}(T_n) = \frac{n}{n-2}.$$

Notice that the variance of T_n tends to 1 from above, as n goes to infinity. Given $\alpha \in (0, 1)$, let $t_{\alpha,n}$ be such that $P[T_n \geq t_{\alpha,n}] = \alpha$. Since t is symmetric at 0, it follows that

$$\alpha = P[-T_n \geq t_{\alpha,n}] = P[T_n \leq -t_{\alpha,n}] = 1 - P[T_n > t_{\alpha,n}].$$

This means that $P[T_n \geq t_{\alpha,n}] = 1 - \alpha$ or, equivalently, $-t_{\alpha,n} = t_{1-\alpha,n}$. See Figure 18.

¹⁰⁶ We will use it later to estimate parameters of a distribution.




Figure 16: The density function of T_n for $n = 1, 2, 10$.




Figure 17: The density function of T_5 (solid) and Z (dotted).




Figure 18: The areas before $-t_{\alpha,n}$ and after $t_{\alpha,n}$ have size α .

Sampling

THE GOAL OF STATISTICS is to draw conclusions about a large *population* by observing a suitable part (or a sample) of it. For the sample data to yield meaningful information about the whole population, one must make some assumptions about the relationship between the two of them. First, we assume that the measures on the individuals of the population are random variables having a shared distribution among the whole population. If the sample data is chosen randomly, it is reasonable to believe that they are independent random values from this distribution. Formally, a *sample* is a finite set of independent, identically distributed, random variables.¹⁰⁷

¹⁰⁷ Also called a *random sample*.

Usually, the population distribution F is not known, and we try to use the data to draw conclusions about it. Sometimes, we may know the general shape of F , but not its precise parameters; for example, that F is normal or Poisson, but with unknown mean and variance. In this case, we speak about *parametric* inferences. If, on the contrary, nothing is assumed about F (except, perhaps, that it is continuous, or discrete), we deal with *non-parametric* inference problems.

From now on, a *statistic* is a random variable whose value is determined by the sample data. Two fundamental such statistics are the *sample mean* and the *sample variance*.

Sample Mean and the Central Limit Theorem

CONSIDER A POPULATION OF elements, each of whom is associated with a value derived from a random variable with mean μ and variance σ^2 . These values are known as the *population mean* and *population variance*, respectively.¹⁰⁸ Suppose that we extract a sample from this population, and observe the values $\mathcal{X}_1, \dots, \mathcal{X}_n$ associated to the elements of the sample. The *sample mean* is

$$\overline{\mathcal{X}} := \frac{\sum_{i=1}^n \mathcal{X}_i}{n}.$$

Notice that $\overline{\mathcal{X}}$ is also a random variable; so it has an associated expected value and variance. Since the RVs in a sample are always independent, these are given by

$$\begin{aligned} E[\overline{\mathcal{X}}] &= E\left[\frac{\sum_{i=1}^n \mathcal{X}_i}{n}\right] = \frac{1}{n} \sum_{i=1}^n E[\mathcal{X}_i] = \mu, \\ \text{Var}(\overline{\mathcal{X}}) &= \text{Var}\left(\frac{\sum_{i=1}^n \mathcal{X}_i}{n}\right) = \frac{1}{n^2} \sum_{i=1}^n \text{Var}(\mathcal{X}_i) = \frac{\sigma^2}{n}. \end{aligned}$$

¹⁰⁸ Recall that all the elements in the sample are equally distributed.

In other words, the expected value of the sample mean is the population mean, while its variance is the population variance divided by the size of the sample. Thus, $\bar{\mathcal{X}}$ is centered around μ but its variance tends to 0 as the sample size grows. To understand how the RV $\bar{\mathcal{X}}$ is distributed, we consider the *central limit theorem* which in essence states that the sum of a large number of independent random variables has a distribution that approximates a normal.¹⁰⁹


Theorem 63 (Central Limit Theorem). *If $\mathcal{X}_1, \dots, \mathcal{X}_n$ are independent identically distributed random variables with mean μ and variance σ^2 , then for a large n the distribution of $\sum_{i=1}^n \mathcal{X}_i$ is approximately normal.*¹¹⁰

Example 64. An insurance company with 25,000 policy holders observes that the yearly claim of a holder is a RV with mean €320 and standard deviation €540. What is the probability that the total yearly claim is above €8.3 million?

Solution. Let \mathcal{X} be the total yearly claim, and \mathcal{X}_i the yearly claim of client i ; then $\mathcal{X} = \sum_{i=1}^{25000} \mathcal{X}_i$. From the central limit theorem, \mathcal{X} is approximately a normal with mean $320 \cdot 25000 = 8 \times 10^6$ and standard deviation $540 \cdot \sqrt{25000} = 8.53 \times 10^4$. Thus,

$$P[\mathcal{X} > 8.3 \times 10^6] = P\left[\frac{\mathcal{X} - 8 \times 10^6}{8.53 \times 10^4} > \frac{8.3 \times 10^6 - 8 \times 10^6}{8.53 \times 10^4}\right] \\ \approx P[Z > 3.51] \approx 0.00023. \quad \triangle$$

The central limit theorem has an important use with binomial random variables. Recall that a binomial \mathcal{X} with parameters (n, p) is a sum of n independent Bernoulli random variables $\mathcal{X} = \sum_{i=1}^n \mathcal{X}_i$ with $E[\mathcal{X}_i] = p$ and $Var(\mathcal{X}_i) = p(1-p)$. Then, if n is large enough, $\frac{\mathcal{X} - np}{\sqrt{np(1-p)}}$ approximates a standard normal random variable (see Figure 19).¹¹¹



¹⁰⁹ The proof of this theorem is well beyond the scope of the lecture, but this is a fundamental result in statistics.

¹¹⁰ A consequence of this theorem is that $(\sum_{i=1}^n \mathcal{X}_i - n\mu)/(\sigma\sqrt{n})$ approximates a standard normal random variable.

¹¹¹ Since the binomial is discrete, and the normal is continuous, sometimes it is necessary to make adjustments (called *continuity corrections*) when using this approximation. See footnote 112.

Example 65. An airplane fits 150 passengers. On a busy route, only 30% of the people that buy the ticket take the plane. If the airline starts selling 450 tickets per flight, what is the probability that the flight is overbooked?

Solution. Let \mathcal{X} be the number of passengers in the flight. Assuming that each customer will independently decide to take the flight or not, \mathcal{X} is a binomial $(450, 0.3)$. Since the binomial is discrete and the normal continuous, we compute $P[\mathcal{X} = i]$ as $P[i - 0.5 < \mathcal{X} < i + 0.5]$ when using the normal approximation.¹¹² Then,

$$P[\mathcal{X} > 150.5] = P\left[\frac{\mathcal{X} - 450 \cdot 0.3}{\sqrt{450 \cdot 0.3 \cdot 0.7}} > \frac{150.5 - 450 \cdot 0.3}{\sqrt{450 \cdot 0.3 \cdot 0.7}}\right] \approx P[Z > 1.59] = 0.06.$$

Figure 19: Binomial probability mass functions for $p = 0.6$ with n growing from 10 to 90. The mass converges to a normal density.

¹¹² This is the *continuity correction*.

That is, the probability of overbooking is 6%. \triangle

We have seen two ways to approximate a binomial: the (discrete) Poisson that works well with n large and p small, and the (continuous) normal, which works when $np(1-p)$ is large.¹¹³

¹¹³ In practice, the normal is a good approximation whenever $np(1-p) \geq 10$.

THE CENTRAL LIMIT THEOREM helps to approximate also the distribution of the sample mean. Indeed, $\bar{\mathcal{X}} = \sum_{i=1}^n \mathcal{X}_i / n$ will be approximately normal when n is large. In particular, $\frac{\bar{\mathcal{X}} - \mu}{\sigma/\sqrt{n}}$ approximates a standard normal RV.¹¹⁴ Note that this approximation works, regardless of the distribution of the original RV \mathcal{X} .

¹¹⁴ Since $E[\bar{\mathcal{X}}] = \mu$ and $\text{Var}(\bar{\mathcal{X}}) = \sigma^2/n$.

Example 66. The measurements for the distance d to a star are affected by atmospheric disturbances. To get a precise reading, a measurement is repeated several times, and their average is used as an estimate for the actual distance. Assuming that each measurement is an independent random variable with mean d and standard deviation 2, how many measurements we need to have a 95% certainty of the estimate being within a value of 0.5 from the actual distance?

Solution. After n measurements, the sample mean $\bar{\mathcal{X}}$ of these measurements will be approximately a normal with mean d and standard deviation $2/\sqrt{n}$. Thus,¹¹⁵

$$\begin{aligned} P[-0.5 < \bar{\mathcal{X}} - d < 0.5] &= P\left[\frac{-0.5}{2/\sqrt{n}} < \frac{\bar{\mathcal{X}} - d}{2/\sqrt{n}} < \frac{0.5}{2/\sqrt{n}}\right] \\ &\approx P[-\sqrt{n}/4 < Z < \sqrt{n}/4] = 2P[Z < \sqrt{n}/4] - 1. \end{aligned}$$

¹¹⁵ Recall that Z is symmetric around 0.

To get a 95% certainty, we need an n such that $2P[Z < \sqrt{n}/4] - 1 \geq 0.95$; that is, $P[Z < \sqrt{n}/4] \geq 0.975$. Since $P[Z < 1.96] = 0.975$, we need n such that $\sqrt{n}/4 \geq 1.96$. Thus, at least 62 measurements are needed. \triangle

While the central limit theorem guarantees that one will approximate a normal distribution as the sample size n grows, it is not explicit on how large should n be for the approximation to be good enough. In practical terms, a sample size of at least 30 will usually suffice, but depending on the population distribution, a much smaller sample might still work; in some cases, even 5 measurements would be good enough.¹¹⁶

¹¹⁶ Obviously, the larger the sample, the better the result will be.

Sample Variance

CONSIDER AGAIN A RANDOM sample from a distribution, and let $\bar{\mathcal{X}}$ be the sample mean. The *sample variance* is defined by

$$S^2 := \frac{\sum_{i=1}^n (\mathcal{X}_i - \bar{\mathcal{X}})^2}{n-1}.$$

The *sample standard deviation* is $S := \sqrt{S^2}$.

The sample standard deviation is also a random variable, and hence it is interesting to know its expected value. To find out $E[S^2]$ recall that for any sequence of numbers x_1, \dots, x_n , and for $\bar{x} := \sum_{i=1}^n x_i / n$,

$$\sum_{i=1}^n (x_i - \bar{x})^2 = \sum_{i=1}^n x_i^2 - n\bar{x}^2.$$

It then follows that $(n-1)S^2 = \sum_{i=1}^n \mathcal{X}_i^2 - n\bar{\mathcal{X}}^2$. Taking expectations on both sides, and using the fact that all \mathcal{X}_i s are equally distributed,¹¹⁷

$$\begin{aligned} (n-1)E[S^2] &= E\left[\sum_{i=1}^n \mathcal{X}_i^2\right] - nE[\bar{\mathcal{X}}^2] \\ &= nE[\mathcal{X}_1^2] - nE[\bar{\mathcal{X}}^2] \\ &= n\text{Var}(\mathcal{X}_1) + n(E[\mathcal{X}_1])^2 - n\text{Var}(\bar{\mathcal{X}}) - n(E[\bar{\mathcal{X}}])^2 \\ &= n\sigma^2 + n\mu^2 - n(\sigma^2/n) - n\mu^2 = (n-1)\sigma^2. \end{aligned}$$

That is, the expected value of S^2 is the population variance σ^2 .

Sampling from a Normal Population

SUPPOSE THAT WE SAMPLE from a normal distribution $\mathcal{X} \sim \mathcal{N}(\mu, \sigma^2)$. Since the sum of independent normal random variables is also normal, it follows that $\bar{\mathcal{X}}$ is also a normal, and in particular $\frac{\bar{\mathcal{X}} - \mu}{\sigma/\sqrt{n}}$ is a standard normal random variable.

Notice that $\sum_{i=1}^n (x_i - \bar{x})^2 = \sum_{i=1}^n (x_i - \mu)^2 - n(\bar{x} - \mu)^2$.¹¹⁸ Then it follows that

$$\frac{\sum_{i=1}^n (\mathcal{X}_i - \mu)^2}{\sigma^2} = \frac{\sum_{i=1}^n (\mathcal{X}_i - \bar{\mathcal{X}})^2}{\sigma^2} + \frac{n(\bar{\mathcal{X}} - \mu)^2}{\sigma^2}.$$

That is,

$$\sum_{i=1}^n \left(\frac{\mathcal{X}_i - \mu}{\sigma} \right)^2 = \frac{\sum_{i=1}^n (\mathcal{X}_i - \bar{\mathcal{X}})^2}{\sigma^2} + \left(\frac{\sqrt{n}(\bar{\mathcal{X}} - \mu)}{\sigma} \right)^2.$$

Each $(\mathcal{X}_i - \mu)/\sigma$ is a standard normal distribution, and they are all independent. So, the left-hand side of this equation is a chi-square distribution with n degrees of freedom.¹¹⁹ The second term on the right is also a chi-square with 1 degree of freedom. This suggests, given the additive property of chi-square distributions, that the missing term is also a chi-square with $n-1$ degrees of freedom. Indeed, we get the following important result.

Theorem 67. *If $\mathcal{X}_1, \dots, \mathcal{X}_n$ is a sample from a normal population with mean μ and variance σ^2 , then $\bar{\mathcal{X}}$ and S^2 are independent random variables with $\bar{\mathcal{X}}$ a normal and $(n-1)S^2/\sigma^2$ a chi-square with $n-1$ degrees of freedom.*

Importantly, this theorem not only establishes the distributions of the sample mean and sample variance, but also their independence. The latter only holds for normal distributions.

Corollary 68. $\sqrt{n} \frac{(\bar{\mathcal{X}} - \mu)}{S} \sim t_{n-1}$.¹²⁰

Proof. Theorem 67 states that $\sqrt{n} \frac{(\bar{\mathcal{X}} - \mu)}{\sigma}$ is a standard normal independent from $(n-1)S^2/\sigma^2$, which is a chi-square with $n-1$ degrees of freedom. Hence,

$$\frac{\sqrt{n}(\bar{\mathcal{X}} - \mu)/\sigma}{\sqrt{(n-1)S^2/\sigma^2}} = \sqrt{n} \frac{(\bar{\mathcal{X}} - \mu)}{S}$$

is a t with $n-1$ degrees of freedom. \square

¹¹⁷ Recall that $E[\mathcal{Y}^2] = \text{Var}(\mathcal{Y}) + (E[\mathcal{Y}])^2$ for all random variables \mathcal{Y} .

¹¹⁸ Let $y_i = x_i - \mu$. Then $\bar{y} = \bar{x} - \mu$ and

$$\begin{aligned} \sum ((x_i - \mu) - (\bar{x} - \mu))^2 &= \sum (y_i - \bar{y})^2 \\ &= \sum y_i^2 - n\bar{y}^2 \\ &= \sum (x_i - \mu)^2 - n(\bar{x} - \mu)^2. \end{aligned}$$

¹¹⁹ A chi-square with n degrees of freedom is the sum of the squares of n independent standard normal distributions.

¹²⁰ That is, $\sqrt{n} \frac{(\bar{\mathcal{X}} - \mu)}{S}$ has a t -distribution with $n-1$ degrees of freedom.

Sampling from a Finite Population

GIVEN A FINITE POPULATION of size N , a *random sample* of size n is such that any of the $\binom{N}{n}$ subsets of the population of size n is equally likely to be chosen.

Suppose that the proportion of the population with a characteristic of interest is p , and that we have a random sample of size n . For $i, 1 \leq i \leq n$ let \mathcal{X}_i be the indicator variable for the i -th element in the sample having the characteristic. Then $\mathcal{X} = \sum_{i=1}^n \mathcal{X}_i$ is the count of elements in the sample having the characteristic, and the sample mean $\overline{\mathcal{X}} = \mathcal{X}/n$ is the proportion of the sample with the characteristic. We study \mathcal{X} and $\overline{\mathcal{X}}$.

Notice that each \mathcal{X}_i is a Bernoulli with parameter p .¹²¹ However, they are not independent: we know already that $P[\mathcal{X}_2 = 1] = p$, but conditioning over the outcome of the first sample, we get $P[\mathcal{X}_2 = 1 | \mathcal{X}_1 = 1] = \frac{Np-1}{N-1}$. Knowing that we have already seen an element with the characteristic, the second-sampled element is equally likely to be any of the remaining $N-1$ elements, but only $Np-1$ of them have the property.¹²²

When the size of the population N is very large in comparison to the sample size n , the difference between the unconditional and the conditional probabilities can be dismissed.¹²³ For example, if $N = 2000$ and $p = 0.5$, then $P[\mathcal{X}_2 = 1 | \mathcal{X}_1 = 1] = \frac{999}{1999} = 0.4997$; similarly $P[\mathcal{X}_2 = 1 | \mathcal{X}_1 = 0] = \frac{1000}{1999} = 0.5003$. Thus, assuming that N is large, we can think of each \mathcal{X}_i as an independent Bernoulli, meaning that \mathcal{X} is approximately a binomial with parameters (n, p) .¹²⁴ Under this view it follows that

$$\begin{aligned} E[\mathcal{X}] &= np & \text{Var}(\mathcal{X}) &= np(1-p), \\ E[\overline{\mathcal{X}}] &= E[\mathcal{X}]/n = p & \text{Var}(\overline{\mathcal{X}}) &= \text{Var}(\mathcal{X})/n^2 = p(1-p)/n. \end{aligned}$$

Example 69. Suppose that 40% of the population supports a given political candidate. Given a random sample of 150 individuals, find

1. the expected value and variance of the number of sampled individuals that favour the candidate;
2. the probability that more than half of the sample favours the candidate.¹²⁵

Solution. The number of people \mathcal{X} favouring the candidate in the sample is a binomial with parameters $(150, 0.4)$. So, $E[\mathcal{X}] = 150 \cdot 0.4 = 60$ and $\text{Var}(\mathcal{X}) = 150 \cdot 0.4 \cdot 0.6 = 36$.

To compute $P[\mathcal{X} \geq 76]$, we can use the normal approximation:^o

$$\begin{aligned} P[\mathcal{X} \geq 76] &= P[\mathcal{X} \geq 75.5] = P\left[\frac{\mathcal{X} - 60}{6} \geq \frac{75.5 - 60}{6}\right] \\ &\approx P[Z \geq 2.5833] \approx 0.0049. \end{aligned} \quad \triangle$$

It is important to notice that, although we considered it here only for the Bernoulli case, in general even if the random variables can take more than two values, when the population is large with respect to the sample size, we can always assume that the sample data are independent random variables with the population distribution.¹²⁶

¹²¹ Every \mathcal{X}_i is an indicator variable for the i -th element from the sample to have the characteristic. But every element from the population is equally likely to be the i -th sampled.

¹²² See Example 57.

¹²³ If we are sampling from the whole Italian population, this is a good assumption. If instead, we sample from the students in the classroom, it might not make much sense.

¹²⁴ Notice that \mathcal{X} is in fact a hypergeometric random variable. Thus, we have shown that binomials approximate hypergeometrics when the number of elements is large w.r.t. to the number of selections.

¹²⁵ When polling we are often trying to answer the dual questions: given the values of the sample, what is the probability that the full population supports the candidate. This is a topic for a different chapter.

^o Do not forget the continuity correction!

¹²⁶ And hence, apply the central limit theorem.

Parameter Estimation

ONE OF THE MAIN problems in statistics is to understand the distribution of a population given the data from a sample. Often, we can assume that the distribution is known up to a vector of unknown parameters. For example, we may say that it is a normal with unknown mean and variance,¹²⁷ or an exponential with unknown rate. In that case, we use the data to estimate the value of the missing parameters. These estimates may be precise (*point estimates*), or give a larger range (*interval estimates*).

¹²⁷ Recall the law of large numbers.

Any statistic used to estimate the value of an unknown parameter θ is called an *estimator* of θ . The observed value of this estimator is an *estimate*.

Maximum Likelihood Estimators

SUPPOSE THAT WE OBSERVE a sequence $\mathcal{X}_1, \dots, \mathcal{X}_n$ of random variables whose joint distribution is known except for an unknown parameter θ . We want to use the observed values to understand θ . If, for example, the observations are all from independent exponential random variables with the same unknown rate θ , their joint density function is

$$\begin{aligned} f(x_1, \dots, x_n) &= f_{\mathcal{X}_1}(x_1) \cdots f_{\mathcal{X}_n}(x_n) \\ &= \prod_{i=1}^n \frac{1}{\theta} e^{-x_i/\theta} & (0 < x_i < \infty) \\ &= \frac{1}{\theta^n} \exp\left(-\sum_{i=1}^n x_i/\theta\right). & (0 < x_i < \infty) \end{aligned}$$

One of the most used type of estimator is the *maximum likelihood* estimator. Let $f(x_1, \dots, x_n | \theta)$ be the joint probability mass or density function of the variables $\mathcal{X}_1, \dots, \mathcal{X}_n$,¹²⁸ this is a function of θ since this parameter is unknown. This function f represents the likelihood of observing x_1, \dots, x_n when θ is the *true* value of the parameter. The maximum likelihood estimate $\hat{\theta}$ of θ is the value that maximises $f(x_1, \dots, x_n | \theta)$ given the observation x_1, \dots, x_n . One useful property for computing this estimator is that $f(x_1, \dots, x_n | \theta)$ and $\log(f(x_1, \dots, x_n | \theta))$ are maximised at the same value of θ so it suffices to maximise the latter function.

¹²⁸ Depending on whether they are discrete or jointly continuous.

Maximum Likelihood Estimator of a Bernoulli Parameter

CONSIDER A SEQUENCE OF n independent trials that have each a probability of success p , which we want to estimate.¹²⁹ The trial consists of n Bernoulli random variables \mathcal{X}_i with an unknown parameter p . Then, for $x = 0, 1$ we have $P[\mathcal{X}_i = x] = p^x(1-p)^{1-x}$, and the joint probability mass function of the data is

$$\begin{aligned} f(x_1, \dots, x_n | p) &= P[\mathcal{X}_1 = x_1, \dots, \mathcal{X}_n = x_n | p] \\ &= \prod_{i=1}^n p^{x_i} (1-p)^{1-x_i} = p^{\sum_{i=1}^n x_i} (1-p)^{n-\sum_{i=1}^n x_i}. \end{aligned}$$

We want to determine the value of p that maximises this function. This is easier to do by taking logarithms:

$$\log(f(x_1, \dots, x_n | p)) = \sum_{i=1}^n x_i \log p + \left(n - \sum_{i=1}^n x_i \right) \log(1-p).$$

To optimise, we can find the point where the derivative (w.r.t. p) evaluates to 0; that is,[◦]

$$\frac{d}{dp} \log(f(x_1, \dots, x_n | p)) = \frac{\sum_{i=1}^n x_i}{p} - \frac{n - \sum_{i=1}^n x_i}{1-p} = 0.$$

Solving this equation yields the estimator $\frac{\hat{p}}{\sum_{i=1}^n x_i} = \frac{1-\hat{p}}{n-\sum_{i=1}^n x_i}$, or equivalently

$$\hat{p} = \frac{\sum_{i=1}^n x_i}{n}.$$

That is, the maximum likelihood estimator of p is the proportion of successes observed in the trials.[◦]

¹²⁹ This is the estimator of the parameter of a Bernoulli distribution.

[◦] Recall that the derivative of $\log p$ is $1/p$.

[◦] Is this intuitive?

Maximum Likelihood Estimator of a Poisson Parameter

CONSIDER NOW INDEPENDENT POISSON random variables $\mathcal{X}_1, \dots, \mathcal{X}_n$ all with the same (unknown) parameter λ . To estimate λ , we first compute the likelihood function

$$f(x_1, \dots, x_n | \lambda) = \prod_{i=1}^n \frac{e^{-\lambda} \lambda^{x_i}}{x_i!} = \frac{e^{-n\lambda} \lambda^{\sum_{i=1}^n x_i}}{\prod_{i=1}^n x_i!}.$$

Taking the logarithm yields

$$\log(f(x_1, \dots, x_n | \lambda)) = -n\lambda + \sum_{i=1}^n x_i \log(\lambda) - \log\left(\prod_{i=1}^n x_i\right).$$

To find the maximum, we evaluate the derivative (w.r.t. λ) at 0

$$\frac{d}{d\lambda} \log(f(x_1, \dots, x_n | \lambda)) = -n + \frac{\sum_{i=1}^n x_i}{\lambda} = 0,$$

which yields $\hat{\lambda} = \frac{\sum_{i=1}^n x_i}{n}$.

For example, suppose that the number of customers at a bar in one hour is a Poisson with an unknown parameter λ that we want to estimate. If in a working day of 12 hours, the bar gets 180 customers, then the maximum likelihood estimate for λ is $156/12 = 15$.¹³⁰

¹³⁰ That is, on average, there will be 15 customers every hour.

Estimators for Normal Parameters

CONSIDER NOW THE MORE complex case of estimating both parameters from a normal distribution.¹³¹ If $\mathcal{X}_1, \dots, \mathcal{X}_n$ are independent identically distributed normal random variables, their joint density, and its logarithm are

$$\begin{aligned} f(x_1, \dots, x_n | \mu, \sigma) &= \prod_{i=1}^n \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x_i - \mu)^2}{2\sigma^2}\right) \\ &= \left(\frac{1}{2\pi}\right)^{n/2} \frac{1}{\sigma^n} \exp\left(-\frac{\sum_{i=1}^n (x_i - \mu)^2}{2\sigma^2}\right) \\ \log(f(x_1, \dots, x_n | \mu, \sigma)) &= -\frac{n}{2} \log(2\pi) - n \log(\sigma) - \frac{\sum_{i=1}^n (x_i - \mu)^2}{2\sigma^2} \end{aligned}$$

As usual, the values of μ and σ that maximise this function can be found by differentiating and evaluating at 0.

$$\begin{aligned} \frac{\partial}{\partial \mu} \log(f(x_1, \dots, x_n | \mu, \sigma)) &= \frac{\sum_{i=1}^n (x_i - \mu)}{\sigma^2} = 0, \\ \frac{\partial}{\partial \sigma} \log(f(x_1, \dots, x_n | \mu, \sigma)) &= -\frac{n}{\sigma} + \frac{-\sum_{i=1}^n (x_i - \mu)^2}{\sigma^3} = 0. \end{aligned}$$

This yields, when solving,

$$\begin{aligned} \hat{\mu} &= \sum_{i=1}^n x_i / n \\ \hat{\sigma} &= \left(\sum_{i=1}^n (x_i - \hat{\mu})^2 / n \right)^{1/2}. \end{aligned}$$

Note that the maximum likelihood estimator for the standard deviation is *not* the same as the sample standard deviation,¹³² but their difference decreases as n grows.

IN ALL THE PREVIOUS cases, the maximum likelihood estimator for the mean was the sample mean. We show that this is not necessarily the case. If $\mathcal{X}_1, \dots, \mathcal{X}_n$ are a sample from a uniform distribution over $(0, \theta)$, where θ is unknown, then we have

$$f(x_1, \dots, x_n | \theta) = \begin{cases} \frac{1}{\theta^n} & 0 \leq x_i \leq \theta, 1 \leq i \leq n \\ 0 & \text{otherwise.} \end{cases}$$

To maximise this function, we need to choose θ to be as small as possible.[°] Since θ cannot be smaller than any of the observed x_i s,[°] the maximum likelihood estimator of θ is $\hat{\theta} = \max\{x_1, \dots, x_n\}$. Thus, the maximum likelihood estimator for the mean of the distribution is $\max\{x_1, \dots, x_n\}/2$.¹³³

Interval Estimates

SUPPOSE THAT WE HAVE a sample from a normal population with unknown mean μ and *known* variance σ^2 . The maximum likelihood estimator for μ is $\bar{\mathcal{X}} = \sum_{i=1}^n \mathcal{X}_i / n$. However, we do not expect the value of $\bar{\mathcal{X}}$

¹³¹ In this case, the mean μ and the standard deviation σ .

¹³² The former divides by n while the latter divides by $(n - 1)$.

[°] As θ grows, $1/\theta^n$ decreases.

[°] All the x_i s are observations that should be in the interval $(0, \theta)$. If $\theta < x_i$, then the observation x_i would be impossible.

¹³³ Remember that the mean of a uniform distribution over $(0, a)$ is $a/2$.

to be *exactly* μ .^o For this reason, we often prefer to compute an interval guaranteed to contain μ (with some degree of confidence). This interval is found through the probability distribution of the point estimator. In this example, $\bar{\mathcal{X}} \sim \mathcal{N}(\mu, \sigma^2/n)$ and hence $\frac{\bar{\mathcal{X}} - \mu}{\sigma/\sqrt{n}} \sim Z$. Thus, it follows that¹³⁴

$$\begin{aligned} 0.95 &= P \left[-1.96 < \frac{\sqrt{n}}{\sigma} (\bar{\mathcal{X}} - \mu) < 1.96 \right] \\ &= P \left[-1.96 \frac{\sigma}{\sqrt{n}} < \bar{\mathcal{X}} - \mu < 1.96 \frac{\sigma}{\sqrt{n}} \right] \\ &= P \left[-1.96 \frac{\sigma}{\sqrt{n}} < \mu - \bar{\mathcal{X}} < 1.96 \frac{\sigma}{\sqrt{n}} \right] \quad (\text{symmetry}) \\ &= P \left[\bar{\mathcal{X}} - 1.96 \frac{\sigma}{\sqrt{n}} < \mu < \bar{\mathcal{X}} + 1.96 \frac{\sigma}{\sqrt{n}} \right]. \end{aligned}$$

In other words, there is a 95% chance that the *true mean* μ lies within a distance $1.96 \frac{\sigma}{\sqrt{n}}$ from $\bar{\mathcal{X}}$. If we observe that $\bar{\mathcal{X}} = \bar{x}$, then we have a 95% confidence that μ is in the interval

$$\left(\bar{x} - 1.96 \frac{\sigma}{\sqrt{n}}, \bar{x} + 1.96 \frac{\sigma}{\sqrt{n}} \right).$$

This is known as the *95 percent confidence interval* of μ .¹³⁵

Example 70. When sending a signal μ over a noisy channel, it is received as $\mu + N$ where N is a normal with mean 0 and variance 4. To reduce the error, each signal is sent 9 times, and the receiver is tasked with estimating the original signal. Suppose that we receive the signals 5, 8.5, 12, 15, 7, 9, 7.5, 6.5, and 10.5. Then $\bar{x} = 81/9 = 9$. Through the maximum likelihood estimator, we conclude that the signal was 9. To have a better understanding of the signal, however, we decide to compute its 95% confidence interval, which is

$$\left(9 - 1.96 \frac{2}{3}, 9 + 1.96 \frac{2}{3} \right) = (7.69, 10.31).$$

We can guarantee with 95% confidence that the *true* original signal lies in this interval. \triangle

An interval like the one from this example is known as a *two-sided confidence interval* since we allow for the real mean μ to be on either side of the point estimator. Alternatively, we may prefer a *one-sided confidence interval*; for example, finding a value x such that we can assert, with 95% confidence, that $\mu > x$. To find this value x , notice that $P[Z < 1.645] = 0.95$. This means that

$$0.95 = P \left[\sqrt{n} \frac{\bar{\mathcal{X}} - \mu}{\sigma} < 1.645 \right] = P \left[\bar{\mathcal{X}} - 1.645 \frac{\sigma}{\sqrt{n}} < \mu \right].$$

The 95% *one-sided upper confidence interval* for μ is $(\bar{x} - 1.645\sigma/\sqrt{n}, \infty)$, where \bar{x} is the observed value of the sample mean. Similarly, the 95% *one-sided lower confidence interval* for μ is $(-\infty, \bar{x} + 1.645\sigma/\sqrt{n})$.

Example 71. The upper 95% confidence interval estimate of μ from Example 70 is $(9 - 1.645 \frac{2}{3}, \infty) = (7.903, \infty)$. \triangle

The use of 95% in the previous discussion is completely arbitrary; it is possible to compute the confidence intervals for any desired level of confidence.¹³⁶ Recall that for any $\alpha \in (0, 1)$, $P[-z_{\alpha/2} < Z < z_{\alpha/2}] = 1 - \alpha$;¹³⁷ see

^oIn fact, $\bar{\mathcal{X}}$ is a random variable, too.

¹³⁴ Recall Example 66.

¹³⁵ One important note on confidence intervals: a “95% confidence interval” for μ does not mean that μ has a probability of 95% of belonging to this interval; in fact, there are no random variables involved to make a probabilistic statement. What it says is that when using this method, in 95% of the cases, the interval will in fact contain the *true mean* μ (but we will miss it 5% of the time).

¹³⁶ In practice, one would like the confidence to be high, but the interval not to be too wide.

¹³⁷ For $\alpha \in (0, 1)$, z_{α} is the value such that $P[Z < z_{\alpha}] = \alpha$.

Figure 20. In the computation of the 95% confidence interval, we used the fact that for $\alpha = 0.05$,¹³⁸ $z_{\alpha/2} = z_{0.025} = 1.96$. Following the same computations, we can generalise the idea to *any* confidence $1 - \alpha$: the $100(1 - \alpha)$ percent two-sided confidence interval for μ is

$$(\bar{x} - z_{\alpha/2} \frac{\sigma}{\sqrt{n}}, \bar{x} + z_{\alpha/2} \frac{\sigma}{\sqrt{n}}).$$

Similarly, the one-sided confidence intervals for the same level of confidence are $(\bar{x} - z_{\alpha} \frac{\sigma}{\sqrt{n}}, \infty)$, and $(-\infty, \bar{x} + z_{\alpha} \frac{\sigma}{\sqrt{n}})$, respectively, where \bar{x} is the observed sample mean.

Example 72. Suppose that in Example 70 we are interested in the 99% two-sided, and one sided upper confidence interval estimates for μ . Then, knowing that $z_{0.005} = 2.58$ and $z_{0.01} = 2.33$, we get that these intervals are $9 \pm 2.58^{2/3} = (7.28, 10.72)$ and $(9 - 2.33^{2/3}, \infty) = (7.447, \infty)$, respectively.

SOMETIMES, WE ARE INTERESTED in finding a small confidence interval, without reducing the level of confidence. For example, we may want to guarantee with 99% confidence, that μ is within an interval of length 1.¹³⁸ Since $z_{0.005} = 2.58$, the 99% confidence interval for μ over a sample of size n is $(\bar{x} - 2.58 \frac{\sigma}{\sqrt{n}}, \bar{x} + 2.58 \frac{\sigma}{\sqrt{n}})$, which has length $5.16 \frac{\sigma}{\sqrt{n}}$. Thus, to obtain an interval of length 1, we need to choose an n such that $5.16 \frac{\sigma}{\sqrt{n}} = 1$; that is, $n = (5.16\sigma)^2$.¹³⁹

Example 73. A packaging machine fills a product to a random weight with unknown mean and standard deviation of 0.5kg. To be 95% certain that our estimate of the mean is correct to $\pm 200g$, then we need a sample size n such that $1.96 \frac{\sigma}{\sqrt{n}} \leq 0.2$.¹⁴⁰ That is, $n \geq (9.8 \cdot 0.5)^2 = 4.9^2 = 24.01$. \triangle

SUPPOSE NOW THAT WE want to construct a confidence interval for μ from a sample of a normal distribution with unknown mean and variance. Since σ is unknown, we cannot simply build a standard normal distribution as before.¹⁴¹ However, for the sample standard deviation S ,¹⁴² we know from Corollary 68 that $\sqrt{n} \frac{(\mathcal{X} - \mu)}{S}$ is a t -random variable with $n - 1$ degrees of freedom. Thus, for every $\alpha \in (0, 1)$,

$$P \left[-t_{\alpha/2, n-1} < \sqrt{n} \frac{(\mathcal{X} - \mu)}{S} < t_{\alpha/2, n-1} \right] = 1 - \alpha.$$

Following an analogous argument as for the normal before, if \bar{x} and s are the observed values for \mathcal{X} and S , respectively, we can say with $100(1 - \alpha)$ confidence that $\mu \in (\bar{x} - t_{\alpha/2, n-1} \frac{S}{\sqrt{n}}, \bar{x} + t_{\alpha/2, n-1} \frac{S}{\sqrt{n}})$.¹⁴³

Example 74. Consider again Example 70, but with an unknown noise variance. To compute a 95% confidence interval for the signal μ , we use again $\bar{x} = 9$, and compute $s^2 = \frac{\sum_{i=1}^n x_i^2 - 9(\bar{x})^2}{8} = 9.5$,¹⁴⁴ that is, $s = 3.802$. Knowing that $t_{0.025, 8} = 2.306$, a 95% confidence interval for μ is

$$\left(9 - 2.306 \frac{3.802}{3}, 9 + 2.306 \frac{3.802}{3} \right) = (6.63, 11.37). \quad \triangle$$

Similarly, we can compute the one-sided upper and lower confidence intervals, respectively as $(\bar{x} - t_{\alpha, n-1} \frac{S}{\sqrt{n}}, \infty)$ and $(-\infty, \bar{x} + t_{\alpha, n-1} \frac{S}{\sqrt{n}})$.

$$^{\circ} 0.95 = 1 - 0.05.$$




Figure 20: $P[-z_{\alpha/2} < Z < z_{\alpha/2}] = 1 - \alpha$.

¹³⁸ In this case, the only parameter that we can manipulate is the sample size.

¹³⁹ This can be generalised to other levels of confidence in the obvious way.

¹⁴⁰ Remember that $z_{0.025} = 1.96$.

¹⁴¹ By saying that $\sqrt{n}(\mathcal{X} - \mu)/\sigma \sim \mathcal{N}(0, 1)$.

¹⁴² $S = \sqrt{\sum_{i=1}^n (\mathcal{X}_i - \bar{\mathcal{X}})^2 / (n - 1)}$.

¹⁴³ Compare this for the case where the variance σ is known: the formula is essentially the same, but now we use an estimate for σ , and need to correct through a t distribution.

¹⁴⁴ Recall that we have previously shown that $(n - 1)S^2 = \sum_{i=1}^n \mathcal{X}_i^2 - n\bar{\mathcal{X}}^2$.

To recall, when the standard deviation is known, the computation of the confidence interval is based on the fact that $\sqrt{n}(\bar{\mathcal{X}} - \mu)/\sigma \sim \mathcal{N}(0, 1)$, while when σ is unknown we estimate S , and use that $\sqrt{n}(\bar{\mathcal{X}} - \mu)/S \sim t_{n-1}$. Notice that the confidence interval does not need to be larger when σ is unknown. In fact, it is completely possible that the sample variance turns out to be much smaller than the population variance. However, we can show that the *mean length* of the interval is larger when σ is unknown. More precisely, we can show that $t_{\alpha, n-1}E[S] \geq z_{\alpha}\sigma$.[◦]

◦Do you understand why this shows that the interval length is larger?

For this section, we have assumed that the sample was taken from a normal distribution. Using the Central Limit Theorem, we can extend these results to any population distribution, provided that the sample is large enough. Indeed, in that case, $\sqrt{n}(\bar{\mathcal{X}} - \mu)/\sigma$ will be approximately normal, and hence $\sqrt{n}(\bar{\mathcal{X}} - \mu)/S$ will approximate a t .

CONSIDER NOW A SAMPLE from a normal distribution with unknown μ and σ , which we want to use to predict the value of a newly sampled element.¹⁴⁵ A natural point predictor for that value is the sample mean $\bar{\mathcal{X}}$. If we want an interval predictor, we can notice that $\bar{\mathcal{X}}$ is a normal with mean μ and variance σ^2/n , and is independent of $\mathcal{X}_{n+1} \sim \mathcal{N}(\mu, \sigma^2)$. Thus, $\mathcal{X}_{n+1} - \bar{\mathcal{X}} \sim \mathcal{N}(0, \sigma^2/n + \sigma^2)$ or, in other words,

¹⁴⁵ That is, we observe $\mathcal{X}_1, \dots, \mathcal{X}_n$, and we want to predict the value of \mathcal{X}_{n+1} .

$$\frac{\mathcal{X}_{n+1} - \bar{\mathcal{X}}}{\sigma\sqrt{1 + 1/n}} \sim \mathcal{N}(0, 1).$$

Since σ remains unknown, we need to estimate it too. As before, using the sample variance $S^2 = \sum_{i=1}^n (\mathcal{X}_i - \bar{\mathcal{X}})^2 / (n-1)$, which is independent of the previous random variable,¹⁴⁶ we can conclude that

¹⁴⁶ See Theorem 67.

$$\frac{\mathcal{X}_{n+1} - \bar{\mathcal{X}}}{S\sqrt{1 + 1/n}} \sim t_{n-1}.$$

Hence, we get that for any $\alpha \in (0, 1)$

$$\begin{aligned} 1 - \alpha &= P \left[-t_{\alpha/2, n-1} < \frac{\mathcal{X}_{n+1} - \bar{\mathcal{X}}}{S\sqrt{1 + 1/n}} < t_{\alpha/2, n-1} \right] \\ &= P[\bar{\mathcal{X}} - t_{\alpha/2, n-1}S\sqrt{1 + 1/n} < \mathcal{X}_{n+1} < \bar{\mathcal{X}} + t_{\alpha/2, n-1}S\sqrt{1 + 1/n}]. \end{aligned}$$

In other words, if \bar{x} and s are the observed values of $\bar{\mathcal{X}}$ and S , respectively, we can predict with $100(1 - \alpha)$ percent confidence that

$$\mathcal{X}_{n+1} \in \left(\bar{x} - t_{\alpha/2, n-1}s\sqrt{1 + 1/n}, \bar{x} + t_{\alpha/2, n-1}s\sqrt{1 + 1/n} \right).$$

JUST AS WE DID for the mean, one can also construct confidence intervals for an unknown variance σ^2 . In this case, recall that $(n-1)\frac{S^2}{\sigma^2} \sim \chi_{n-1}^2$.¹⁴⁷ Since chi-square distributions are not symmetric and are positive,¹⁴⁸ we need to be careful when building the confidence intervals. In this case,

¹⁴⁷ Theorem 67.

¹⁴⁸ A chi-square is a sum of squares, so it must take only positive values.

$$\begin{aligned} 1 - \alpha &= P \left[\chi_{1-\alpha/2, n-1}^2 \leq (n-1)\frac{S^2}{\sigma^2} \leq \chi_{\alpha/2, n-1}^2 \right] \\ &= P \left[\frac{(n-1)S^2}{\chi_{\alpha/2, n-1}^2} \leq \sigma^2 \leq \frac{(n-1)S^2}{\chi_{1-\alpha/2, n-1}^2} \right]. \end{aligned}$$

In other words, if s^2 is the observed sample variance, then a $100(1 - \alpha)$ confidence interval for σ^2 is

$$\left(\frac{(n-1)S^2}{\chi_{\alpha/2, n-1}^2}, \frac{(n-1)S^2}{\chi_{1-\alpha/2, n-1}^2} \right).$$

One sided intervals can be computed following the same ideas.

Estimating the Difference of Means

CONSIDER NOW TWO INDEPENDENT samples $\mathcal{X}_1, \dots, \mathcal{X}_n$ from a normal population with mean μ_1 and variance σ_1^2 , and $\mathcal{Y}_1, \dots, \mathcal{Y}_m$ from a different normal population with mean μ_2 and variance σ_2^2 . We want to estimate $\mu_1 - \mu_2$.¹⁴⁹ It can be shown that $\bar{\mathcal{X}} - \bar{\mathcal{Y}}$ is the maximum likelihood estimator for $\mu_1 - \mu_2$, where $\bar{\mathcal{X}}$ and $\bar{\mathcal{Y}}$ are the sample means.

To find confidence intervals, we need first to understand the distribution of $\bar{\mathcal{X}} - \bar{\mathcal{Y}}$. We know that $\bar{\mathcal{X}} \sim \mathcal{N}(\mu_1, \sigma_1^2/n)$ and $\bar{\mathcal{Y}} \sim \mathcal{N}(\mu_2, \sigma_2^2/m)$. This means that¹⁵⁰

$$\bar{\mathcal{X}} - \bar{\mathcal{Y}} \sim \mathcal{N}\left(\mu_1 - \mu_2, \frac{\sigma_1^2}{n} + \frac{\sigma_2^2}{m}\right),$$

and hence

$$\frac{\bar{\mathcal{X}} - \bar{\mathcal{Y}} - (\mu_1 - \mu_2)}{\sqrt{\frac{\sigma_1^2}{n} + \frac{\sigma_2^2}{m}}} \sim \mathcal{N}(0, 1).$$

If the two variances are known, then it easy to verify that the two-sided and one-sided $100(1 - \alpha)$ confidence interval estimates for $\mu_1 - \mu_2$ are, respectively,

$$\begin{aligned} & \left(\bar{x} - \bar{y} - z_{\alpha/2} \sqrt{\frac{\sigma_1^2}{n} + \frac{\sigma_2^2}{m}}, \bar{x} - \bar{y} + z_{\alpha/2} \sqrt{\frac{\sigma_1^2}{n} + \frac{\sigma_2^2}{m}} \right); \\ & \left(-\infty, \bar{x} - \bar{y} + z_{\alpha/2} \sqrt{\frac{\sigma_1^2}{n} + \frac{\sigma_2^2}{m}} \right); \text{ and} \\ & \left(\bar{x} - \bar{y} - z_{\alpha/2} \sqrt{\frac{\sigma_1^2}{n} + \frac{\sigma_2^2}{m}}, \infty \right); \end{aligned}$$

where \bar{x} and \bar{y} are the observed sample means.

If the variances are unknown, then we can try to use the same idea of estimating them through the sample variances

$$\begin{aligned} S_1^2 &= \sum_{i=1}^n (\mathcal{X}_i - \bar{\mathcal{X}})^2 / (n-1) \\ S_2^2 &= \sum_{i=1}^m (\mathcal{Y}_i - \bar{\mathcal{Y}})^2 / (m-1), \end{aligned}$$

and use $\frac{\bar{\mathcal{X}} - \bar{\mathcal{Y}} - (\mu_1 - \mu_2)}{\sqrt{S_1^2/n + S_2^2/m}}$ to construct a confidence interval. However, we do not know how this random variable is distributed. Unfortunately, this distribution is very difficult to understand in general. So we focus on the special case where $\sigma_1 = \sigma_2$.¹⁵¹

¹⁴⁹ For example, if we want to know if a system is faster than another.

¹⁵⁰ Remember that the sum of independent normals is a normal.

¹⁵¹ This is, in fact, the only case that we can feasibly handle.

Suppose that the population variances are both equal to an unknown σ^2 . In that case, we know¹⁵²

$$(n-1) \frac{S_1^2}{\sigma^2} \sim \chi_{n-1}^2,$$

$$(m-1) \frac{S_2^2}{\sigma^2} \sim \chi_{m-1}^2.$$

Independence of the samples implies that these two chi-square distributions are also independent, and hence¹⁵³

$$(n-1) \frac{S_1^2}{\sigma^2} + (m-1) \frac{S_2^2}{\sigma^2} \sim \chi_{n+m-2}^2.$$

We have also already established that

$$\frac{\bar{\mathcal{X}} - \bar{\mathcal{Y}} - (\mu_1 - \mu_2)}{\sqrt{\frac{\sigma^2}{n} + \frac{\sigma^2}{m}}} \sim \mathcal{N}(0, 1),$$

and that $\bar{\mathcal{X}}, \bar{\mathcal{Y}}, S_1$, and S_2 are independent random variables. Define now

$$S_p^2 = \frac{(n-1)S_1^2 + (m-1)S_2^2}{n+m-2}.$$

Then it follows that¹⁵⁴

$$\frac{\bar{\mathcal{X}} - \bar{\mathcal{Y}} - (\mu_1 - \mu_2)}{\sqrt{S_p^2(1/n + 1/m)}} = \frac{\bar{\mathcal{X}} - \bar{\mathcal{Y}} - (\mu_1 - \mu_2)}{\sqrt{\sigma^2(1/n + 1/m)}} \bigg/ \sqrt{S_p^2/\sigma^2}$$

is a t -distribution with $n+m-2$ degrees of freedom. Thus, the $100(1-\alpha)$ confidence interval for $\mu_1 - \mu_2$ is¹⁵⁵

$$\left(\bar{x} - \bar{y} - t_{\alpha/2, n+m-2} s_p \sqrt{1/n + 1/m}, \bar{x} - \bar{y} - t_{\alpha/2, n+m-2} s_p \sqrt{1/n + 1/m} \right).$$

WE NEEDED TO ESTIMATE the unknown variance σ^2 to compute this confidence interval. In this case, we had two different samples that provided two estimates for this variance. The estimator S_p^2 used is the weighted average of the two sample variances where the weights are proportional to their degrees of freedom. In general, this is called the *pooled estimator*.

Approximate Confidence Interval for the Mean of a Bernoulli

SUPPOSE THAT WE WANT to estimate the parameter p of a Bernoulli. If \mathcal{X} denotes the number of successes observed in a sample of size n , then \mathcal{X} is a binomial (n, p) . If the sample is large enough, we know that \mathcal{X} is approximately a normal with mean np and variance $np(1-p)$. Thus, for $\alpha \in (0, 1)$

$$P \left[-z_{\alpha/2} < \frac{\mathcal{X} - np}{\sqrt{np(1-p)}} < z_{\alpha/2} \right] \approx 1 - \alpha.$$

Then an approximate confidence *region* for p is

$$\left\{ p \left| -z_{\alpha/2} < \frac{\mathcal{X} - np}{\sqrt{np(1-p)}} < z_{\alpha/2} \right. \right\}.$$

¹⁵² See again Theorem 67.

¹⁵³ Recall that the sum of two independent chi-squares is a chi-square with degrees of freedom equal to the sum of the degrees of freedom.

¹⁵⁴ $(n+m-2)S_p^2/\sigma^2 = (n-1)\frac{S_1^2}{\sigma^2} + (m-1)\frac{S_2^2}{\sigma^2}$.

¹⁵⁵ The one-sided intervals can be easily derived in the same way.

However, this is *not* an interval.^o In order to find an interval, consider $\hat{p} = \mathcal{X}/n$. Since \hat{p} is the maximum likelihood estimator of p , it should be approximately equal to p , and $\sqrt{n\hat{p}(1-\hat{p})} \approx \sqrt{np(1-p)}$. Thus, $\frac{\mathcal{X}-np}{\sqrt{n\hat{p}(1-\hat{p})}}$ is still approximately a standard normal RV. This implies then that

^oDo you see why?

$$P\left[\hat{p} - z_{\alpha/2}\sqrt{\hat{p}(1-\hat{p})/n} < p < \hat{p} + z_{\alpha/2}\sqrt{\hat{p}(1-\hat{p})/n}\right] \approx 1 - \alpha,$$

which yields the (approximate) confidence interval for p .

Example 75. A poll result expresses that 52% of the population favours a candidate with a margin of error of $\pm 4\%$. Knowing that polls are commonly expressed as 95% confidence intervals, what does this statement mean, and what can we say about the size of the poll?

Solution. A 95% confidence interval for the proportion p of people in favour of the candidate given a sample of size n is

$$\hat{p} \pm z_{0.025}\sqrt{\hat{p}(1-\hat{p})/n} = 0.52 \pm 1.96\sqrt{0.52 \cdot 0.48/n}.$$

Since the margin of error is $\pm 4\%$, we can conclude that

$$1.96\sqrt{0.52 \cdot 0.48/n} = 0.04.$$

Which, solving for n yields $n = \frac{(1.96)^2(0.52)(0.48)}{(0.04)^2} = 599.29$. That is, approximately 599 people were polled. \triangle

AS BEFORE, WE MAY be interested in finding a confidence interval of a bounded size b , and need to find out the size of the sample needed to achieve this. Notice that the length of this interval is $2z_{\alpha/2}\sqrt{\hat{p}(1-\hat{p})/n}$. Unfortunately, we do not know \hat{p} (or even p) in advance, so we cannot use this bound to determine the right size n . The trick in this case is to start with a partial sample. With the first sample, we compute an estimate p^* of p . Then, using p^* as an approximation for p (and \hat{p}), we compute the desired sample size by solving

$$\begin{aligned} b &= 2z_{\alpha/2}\sqrt{p^*(1-p^*)/n} \\ b^2 &= (2z_{\alpha/2})^2 p^*(1-p^*)/n \\ n &= \frac{(2z_{\alpha/2})^2 p^*(1-p^*)}{b^2}. \end{aligned}$$

If the original partial sample had k elements, then we need to find an additional sample of $n - k$ elements.

Alternatively, notice that the $100(1-\alpha)\%$ confidence interval of p will have length b if n is at least $\frac{(2z_{\alpha/2})^2}{b^2} p(1-p)$, which is maximised when $p = 1/2$. Thus, any sample size $n \geq \frac{z_{\alpha/2}^2}{b^2}$ will be such that the confidence interval has length at most b .¹⁵⁶

¹⁵⁶ This sample will often be too large in comparison to the bound found before, but it requires no preliminary sampling, and no additional computations.

Hypothesis Testing

HYPOTHESIS TESTING IS A close relative to parameter estimation. The difference is that, rather than trying to find out the unknown parameters of a distribution, hypothesis testing tries to verify—or refute—a statement (hypothesis) about these parameters.¹⁵⁷ For example, we can try to verify whether the mean height of unibz students is 200cm by taking a sample and computing their height. We want to develop a procedure that verifies whether the observed values are consistent with the hypothesis. If the value derived from the sample is inconsistent with our hypothesis (e.g. we get a sample mean of 170cm), then we can *reject* the hypothesis. Otherwise (e.g., if the sample mean is 198cm), we cannot.¹⁵⁸ Usually, we will want to test the opposite of our hypothesis, to be able to reject it.

Significance Levels

SUPPOSE THAT WE WANT to test a hypothesis H_0 about a parameter θ . We often call H_0 the *null hypothesis*.¹⁵⁹ For example, if θ is the mean of a normal distribution, one could generate hypotheses such as

$$H_0 : \theta = 1 \quad \text{or} \quad H_0 : \theta \leq 1.$$

Notice that the latter, when true, will not specify the distribution of the population. A hypothesis of this kind is called *composite* as it refers to a set of possible values. A hypothesis that fully specifies the distribution, like the former, is *simple*.

Given a sample of size n , we need to decide when to reject the hypothesis. Formally, a test for H_0 is defined through a region in \mathbb{R}^n , called the *critical region*. The idea is that C defines the region where H_0 will be rejected; i.e., we reject H_0 if the sample $(\mathcal{X}_1, \dots, \mathcal{X}_n) \in C$. One common test for the hypothesis that the mean θ of a normal with variance 1 is 1 is the critical region¹⁶⁰

$$C = \{(\mathcal{X}_1, \dots, \mathcal{X}_n) \mid |\bar{\mathcal{X}} - 1| > 1.96\sqrt{n}\}.$$

That is, we reject the null hypothesis $\theta = 1$ if the sample mean differs from 1 by at least $1.96\sqrt{n}$.

Recall that the idea is to check whether H_0 is consistent with the observations in the data. So we should only reject it if the data is very unlikely when H_0 is true. We do this by specifying a value α , called the *significance*

¹⁵⁷ It is a *hypothesis* because it is not known whether it is true or not.

¹⁵⁸ In the literature, people often speak of *accepting* a hypothesis. This is a confusing use of the word that lead to misinterpretations. When one *rejects* a hypothesis, it means that there is strong evidence against it. In that sense *accepting* a hypothesis would mean only that there is no evidence against it (but there may also be no strong evidence in its favour).

¹⁵⁹ As mentioned before, the null hypothesis is often the opposite of our claim and we want to reject it.

¹⁶⁰ Compare this with 95% confidence intervals.

level such that H_0 has probability at most α of being rejected when it is true.¹⁶¹ It is good practice to decide this α in advance; in life sciences this is often set to 0.1, 0.05, or 0.005.

The usual approach to hypothesis testing is the following. Given a hypothesis $H_0 : \theta \in \omega$, where ω is a set of parameter values, we find a point estimator $\tilde{\theta}$ for θ , and reject H_0 if $\tilde{\theta}$ is *far* from ω . Obviously, the specific distance required depends on the distribution of $\tilde{\theta}$ under the assumption of H_0 being true, and on the desired significance level. For example, the critical region computed before rejects the hypothesis if the sample mean¹⁶² is further than $1.96/\sqrt{n}$ from 1. This was chosen to meet the significance level $\alpha = 0.05$.

The Mean of a Normal Population

CONSIDER A SAMPLE FROM a normal population with an unknown mean μ and a known variance σ^2 . Suppose that we are interested in testing the null hypothesis $H_0 : \mu = \mu_0$ against the alternative hypothesis $H_1 : \mu \neq \mu_0$, where μ_0 is a given constant. We would like to reject H_0 if the sample mean $\bar{\mathcal{X}}$ is far from μ ; that is, we want to define a critical region of the form

$$C = \{\mathcal{X}_1, \dots, \mathcal{X}_n \mid |\bar{\mathcal{X}} - \mu| > c\},$$

where c is some suitably chosen constant.

If we are interested in a given significance level α , we need to find the value of c that makes the probability of type I error to be α ; that is, we want a c such that¹⁶³

$$P[|\bar{\mathcal{X}} - \mu_0| > c \mid \mu = \mu_0] = \alpha.$$

If $\mu = \mu_0$, then $\bar{\mathcal{X}} \sim \mathcal{N}(\mu_0, \sigma^2/n)$; equivalently, $\frac{\sqrt{n}(\bar{\mathcal{X}} - \mu_0)}{\sigma}$ is a standard normal distribution. Thus, we are searching for a c such that

$$\alpha = P\left[|Z| > \frac{c\sqrt{n}}{\sigma}\right] = 2P\left[Z > \frac{c\sqrt{n}}{\sigma}\right],$$

or equivalently, $P\left[Z > \frac{c\sqrt{n}}{\sigma}\right] = \alpha/2$. This is solved by setting $\frac{c\sqrt{n}}{\sigma} = z_{\alpha/2}$; i.e., $c = \frac{z_{\alpha/2}\sigma}{\sqrt{n}}$. This means that the significance level α test is to reject H_0 iff $|\bar{\mathcal{X}} - \mu_0| > \frac{z_{\alpha/2}\sigma}{\sqrt{n}}$ (see Figure 21); or equivalently, if $\frac{\sqrt{n}}{\sigma}|\bar{\mathcal{X}} - \mu_0| > z_{\alpha/2}$.

Example 76. We want to test the hypothesis that a normal random variable \mathcal{X} with variance 4 has mean 8.¹⁶⁴ From a sample of size 9, we obtain $\bar{\mathcal{X}} = 9.2$. Given a significance level $\alpha = 0.05$, we compute the test statistic

$$\frac{\sqrt{n}}{\sigma}|\bar{\mathcal{X}} - \mu_0| = \frac{3}{2}(1.2) = 1.8.$$

Since this number is not greater than $z_{0.025} = 1.96$, we cannot reject the hypothesis.¹⁶⁵ \triangle

One important question is what is the *right* significance level to use. The specific choice depends on the application, but to be fair and correct it should be decided *before* the test is made, and not adapted to fit our desired conclusions. Importantly, a lower significance level makes it harder to reject the hypothesis, decreasing our chance of error in that case.

¹⁶¹ That is, we bound the probability of error of *type I*—rejecting H_0 when it is correct—by α .

¹⁶² That is, the point estimate for the mean.

¹⁶³ Assuming that $\mu = \mu_0$ (that is, H_0 is true), the probability of rejecting H_0 should be α .




Figure 21: Critical region for a test.

¹⁶⁴ That is, $H_0 : \mu = 8$.

¹⁶⁵ Notice that if we had chosen a more liberal significance level (like 0.1) then we could have rejected H_0 , but under the understanding that the probability of error of type I would increase (from 5 to 10% in this case).

The hypothesis test computes first the test statistic $v := \frac{\sqrt{n}}{\sigma} |\bar{\mathcal{X}} - \mu_0|$, and rejects the null hypothesis $H_0 : \mu = \mu_0$ if $\alpha \geq P[|Z| \geq v]$.¹⁶⁶ The probability $P[|Z| \geq v]$ is called the *p-value* of the test. It provides the critical significance level in the sense that H_0 will be rejected iff the significance level α is greater than or equal to the p-value.

¹⁶⁶ The probability that a standard normal distribution is farther away than the test statistic.

Example 77. In Example 76 we have computed the test statistic to be

$$\frac{\sqrt{n}}{\sigma} |\bar{\mathcal{X}} - \mu_0| = 1.8.$$

Then, the p-value is $P[|Z| > 1.8] = 2P[Z > 1.8] = 2 \cdot 0.036 = 0.072$. This means that the null hypothesis $H_0 : \mu = 8$ will be rejected for any significance level $\alpha > 0.072$.¹⁶⁷

¹⁶⁷ Recall from Example 76 that the test was not rejected with $\alpha = 0.05$, but would have been rejected with $\alpha = 0.1$.

Suppose that instead of 9.2 the sample yielded $\bar{\mathcal{X}} = 10.4$. Then, the test statistic is $\frac{\sqrt{n}}{\sigma} |\bar{\mathcal{X}} - \mu_0| = 3.6$, which gives the p-value $2P[Z > 3.6] = 0.0003$. The hypothesis would be rejected even for very low significance levels. \triangle

ANOTHER TYPE OF ERROR is called *type II*: not rejecting a wrong hypothesis.¹⁶⁸ In the case of hypotheses for the mean μ of a normal, we define the function

$$\begin{aligned} \beta(\mu) &:= P[\text{not rejecting } H_0 \mid \mu] = P\left[\left|\frac{\bar{\mathcal{X}} - \mu_0}{\sigma/\sqrt{n}}\right| \leq z_{\alpha/2} \mid \mu\right] \\ &= P\left[-z_{\alpha/2} \leq \frac{\bar{\mathcal{X}} - \mu_0}{\sigma/\sqrt{n}} \leq z_{\alpha/2} \mid \mu\right] \end{aligned}$$

¹⁶⁸ In our case, failing to reject it when it is wrong.

The function $\beta(\mu)$, called the *operating characteristic (OC) curve*, describes the probability of *not* rejecting H_0 when the true mean is μ .

Recall that $\bar{\mathcal{X}}$ is a normal with mean μ and variance σ^2/n . In other words, $\frac{\bar{\mathcal{X}} - \mu}{\sigma/\sqrt{n}} \sim \mathcal{N}(0, 1)$. This means that¹⁶⁹

¹⁶⁹ We use P_μ to say that this probability depends on the actual value of μ .

$$\begin{aligned} \beta(\mu) &= P_\mu \left[-z_{\alpha/2} \leq \frac{\bar{\mathcal{X}} - \mu_0}{\sigma/\sqrt{n}} \leq z_{\alpha/2} \right] \\ &= P_\mu \left[-z_{\alpha/2} - \frac{\mu}{\sigma/\sqrt{n}} \leq \frac{\bar{\mathcal{X}} - \mu_0 - \mu}{\sigma/\sqrt{n}} \leq z_{\alpha/2} - \frac{\mu}{\sigma/\sqrt{n}} \right] \\ &= P_\mu \left[-z_{\alpha/2} - \frac{\mu}{\sigma/\sqrt{n}} \leq Z - \frac{\mu_0}{\sigma/\sqrt{n}} \leq z_{\alpha/2} - \frac{\mu}{\sigma/\sqrt{n}} \right] \\ &= P_\mu \left[\frac{\mu_0 - \mu}{\sigma/\sqrt{n}} - z_{\alpha/2} \leq Z \leq \frac{\mu_0 - \mu}{\sigma/\sqrt{n}} + z_{\alpha/2} \right] \\ &= P_\mu \left[Z \leq \frac{\mu_0 - \mu}{\sigma/\sqrt{n}} + z_{\alpha/2} \right] - P_\mu \left[Z \leq \frac{\mu_0 - \mu}{\sigma/\sqrt{n}} - z_{\alpha/2} \right] \\ &= \Phi \left(\frac{\mu_0 - \mu}{\sigma/\sqrt{n}} + z_{\alpha/2} \right) - \Phi \left(\frac{\mu_0 - \mu}{\sigma/\sqrt{n}} - z_{\alpha/2} \right). \end{aligned}$$

The precise values of this function obviously depend on the significance level α . For a fixed significance level, the OC curve is symmetric around μ_0 and depends on μ , and its relation to the hypothesis μ_0 . It reaches its maximum $1 - \alpha$ at $\mu = \mu_0$ and decreases as μ gets farther away from this point.^o

^oDo you see why?

Example 78. Continuing Example 76, suppose that we want to compute the probability of not rejecting the null hypothesis $H_0 : \mu = 8$ when the true

value of the mean is 10.¹⁷⁰ Then we compute $\frac{\sqrt{n}}{\sigma}(\mu_0 - \mu) = \frac{3}{2}(8 - 10) = -3$. Using the significance level $\alpha = 0.05$, since $z_{0.025} = 1.96$ the probability is

$$\beta(10) = \Phi(-3 + 1.96) - \Phi(-3 - 1.96) = \Phi(-1.04) - \Phi(-4.96) = 0.145. \quad \Delta$$

The complement of the OC curve, $1 - \beta(\mu)$ is called the *power function* of the test. It expresses the probability of rejecting the hypothesis when the true value is μ .

THE OC FUNCTION IS used to determine the size of a sample needed to guarantee properties of type II errors. Suppose that we want to guarantee that the probability of not rejecting $H_0 : \mu = \mu_0$ is approximately β . For each constant μ_1 , we want to find an n such that $\beta(\mu_1) \approx \beta$.¹⁷¹ That is,

$$\beta \approx \Phi\left(\frac{\mu_0 - \mu_1}{\sigma/\sqrt{n}} + z_{\alpha/2}\right) - \Phi\left(\frac{\mu_0 - \mu_1}{\sigma/\sqrt{n}} - z_{\alpha/2}\right).$$

Unfortunately, we cannot solve this problem analytically. To approximate n , we can use the following reasoning. Suppose (w.l.o.g.) that $\mu_1 > \mu_0$. Then $\mu_0 - \mu_1 < 0$ and hence

$$\frac{\sqrt{n}(\mu_0 - \mu_1)}{\sigma} - z_{\alpha/2} < -z_{\alpha/2}.$$

Since the cumulative density function Φ is always increasing, we get

$$\Phi\left(\frac{\sqrt{n}(\mu_0 - \mu_1)}{\sigma} - z_{\alpha/2}\right) \leq \Phi(-z_{\alpha/2}) = P[Z \leq -z_{\alpha/2}] = P[Z \geq z_{\alpha/2}] = \alpha/2.$$

Assuming that α is small enough—as usually done when dealing with confidence levels—it follows that $\Phi\left(\frac{\sqrt{n}(\mu_0 - \mu_1)}{\sigma} - z_{\alpha/2}\right) \approx 0$, which means that

$$\beta \approx \Phi\left(\frac{\sqrt{n}(\mu_0 - \mu_1)}{\sigma} + z_{\alpha/2}\right). \quad (\dagger)$$

Notice that $\beta \in (0, 1)$, which means that

$$\beta = P[Z > z_\beta] = P[Z < -z_\beta] = \Phi(-z_\beta).$$

Together with (\dagger) , it follows that we want some n such that

$$\begin{aligned} \Phi(-z_\beta) &\approx \Phi\left(\frac{\sqrt{n}(\mu_0 - \mu_1)}{\sigma} + z_{\alpha/2}\right); \quad \text{i.e.,} \\ -z_\beta &\approx \frac{\sqrt{n}(\mu_0 - \mu_1)}{\sigma} + z_{\alpha/2}. \end{aligned}$$

Solving for n then yields^o

$$n \approx \frac{(z_\beta + z_{\alpha/2})^2 \sigma^2}{(\mu_0 - \mu_1)^2}$$

Example 79. Continuing Example 76, if we want the 0.05 level test for the hypothesis $H_0 : \mu = 8$ to have at least a 75% probability of rejection when $\mu = 9.2$, then we need a sample of size¹⁷²

$$n \approx \frac{(z_{0.25} + z_{0.025})^2 \cdot 4}{(8 - 9.2)^2} = \frac{4(1.96 + 0.67)^2}{1.2^2} = 19.21.$$

Thus we need a sample of size 20. Let us confirm this using the OC function. If the true mean is 9.2, and $\mu_0 = 8$ then¹⁷³

¹⁷⁰ Recall from Example 76 that the sample size is 9 and the variance is 4.

¹⁷¹ Recall that the OC function depends on μ .

^oTry to verify that the same approximation would hold if $\mu_1 < \mu_0$.

¹⁷² Recall that β is the probability of not rejecting; in this case $\beta = 0.25$. Recall also that $z_{0.025} = 1.96$ and $z_{0.25} = 0.67$.

¹⁷³ Recall that $\beta(\mu)$ is

$$\Phi\left(\frac{\mu_0 - \mu}{\sigma/\sqrt{n}} + z_{\alpha/2}\right) - \Phi\left(\frac{\mu_0 - \mu}{\sigma/\sqrt{n}} - z_{\alpha/2}\right).$$

$$\begin{aligned}\beta(9.2) &= \Phi\left(\frac{\sqrt{20}(8-9.2)}{2} + 1.96\right) - \Phi\left(\frac{\sqrt{20}(8-9.2)}{2} - 1.96\right) \\ &= \Phi(-0.723) - \Phi(-4.643) \approx \Phi(-0.723) = 1 - \Phi(0.723) = 0.235.\end{aligned}$$

Thus, in this case, if the true mean is 9.2, there will be a 76.5% chance of rejecting the null hypothesis $H_0 : \mu = 8$. \triangle

THROUGHOUT THE PREVIOUS ANALYSIS, we considered a test that would reject whenever the observed sample mean is far away from the hypothesis mean, in any direction. Sometimes, a more meaningful hypothesis is to *bound* the mean e.g. from above. In this case, our null hypothesis will be of the form $H_0 : \mu \leq \mu_0$ and should be rejected only if the estimate of μ is much greater than μ_0 , but not if it is much smaller.¹⁷⁴ In other words, the critical region becomes

$$C = \{(\mathcal{X}_1, \dots, \mathcal{X}_n) \mid \bar{\mathcal{X}} - \mu_0 > c\},$$

for a suitably chosen c .¹⁷⁵ For the test to have a significance level α , we need that¹⁷⁶

$$P[\bar{\mathcal{X}} - \mu_0 > c \mid \mu = \mu_0] = \alpha.$$

As usual, we use the fact that $\bar{\mathcal{X}} \sim \mathcal{N}(\mu, \sigma^2/n)$; thus if the true mean is μ_0 then $\frac{\sqrt{n}(\bar{\mathcal{X}} - \mu_0)}{\sigma} \sim \mathcal{N}(0, 1)$. We want a c such that $P\left[Z > \frac{\sqrt{nc}}{\sigma}\right] = \alpha$; that is $\frac{\sqrt{nc}}{\sigma} = z_\alpha$, or¹⁷⁷

$$c = \frac{z_\alpha \sigma}{\sqrt{n}}.$$

Hence, the test is to reject H_0 iff $\bar{\mathcal{X}} - \mu_0 > z_\alpha \sigma / \sqrt{n}$. This is the *one-sided critical region*.

To compute p-values in a one-side test like this one, we proceed as before: first use the data to compute the statistic $\sqrt{n}(\bar{\mathcal{X}} - \mu_0)/\sigma$. The p-value is the probability that a standard normal distribution is at least as large as this number.¹⁷⁸

Example 80. Consider again Example 76, where we want to test the hypothesis $H_0 : \mu \leq 8$. The test statistic is $\sqrt{n}(\bar{\mathcal{X}} - \mu_0)/\sigma = 3(1.2)/2 = 1.8$. Then, the p-value is¹⁷⁹ $1 - \Phi(1.8) = 0.036$. This means that the test will reject H_0 for any significance level above 0.036. \triangle

The OC function of the one-sided test $\beta(\mu) = P[\text{not rejecting } H_0 \mid \mu]$ can be computed accordingly, dependent on the significance level α , by¹⁸⁰

$$\begin{aligned}\beta(\mu) &= P[\bar{\mathcal{X}} \leq \mu_0 + z_\alpha \sigma / \sqrt{n}] = P\left[\frac{\sqrt{n}(\bar{\mathcal{X}} - \mu)}{\sigma} \leq \frac{\sqrt{n}(\mu_0 - \mu)}{\sigma} + z_\alpha\right] \\ &= P\left[Z \leq \frac{\sqrt{n}(\mu_0 - \mu)}{\sigma} + z_\alpha\right] = \Phi\left(\frac{\sqrt{n}(\mu_0 - \mu)}{\sigma} + z_\alpha\right)\end{aligned}$$

Notice that β decreases as μ increases.[◦] Recall that $\Phi(z_\alpha) = 1 - \alpha$. Thus, $\beta(\mu_0) = 1 - \alpha$.

In the construction, we used the condition $\mu = \mu_0$ to test the hypothesis $H_0 : \mu \leq \mu_0$. We need to verify that this test does preserve the significance level α ; i.e., that when H_0 is true, the probability of rejecting is bounded by α ; equivalently, we may verify that $1 - \beta(\mu) \leq \alpha$ holds for all $\mu \leq \mu_0$,¹⁸¹ or

¹⁷⁴ The latter case would fall inside our hypothesis.

¹⁷⁵ Remember that the critical region is the area of the input variables where the null hypothesis will be rejected. Compare this to the two-sided critical region from Page 60.

¹⁷⁶ The significance level is the probability of error of type I: rejecting the hypothesis when it is true. Notice that the conditioning is still w.r.t. $\mu = \mu_0$. This refers to the extreme case; if the actual mean is smaller than μ_0 it will be harder to reject the hypothesis.

¹⁷⁷ Compare again with two-sided confidence intervals.

¹⁷⁸ If $v := \sqrt{n}(\bar{\mathcal{X}} - \mu_0)/\sigma$ is the test statistic, then the p-value is $P[Z \geq v]$.

¹⁷⁹ The p-value in this case is $P[Z > 1.8]$; see note 178.

¹⁸⁰ Under the assumption that the true mean is μ , $\frac{\sqrt{n}(\bar{\mathcal{X}} - \mu)}{\sigma} \sim \mathcal{N}(0, 1)$.

[◦] Does this make sense?

¹⁸¹ Remember that the OC function $\beta(\mu)$ expresses the probability of not rejecting H_0 if the true mean is μ .

alternatively $\beta(\mu) \geq 1 - \alpha$. By the previous discussion, we know already that for all $\mu \leq \mu_0$, $\beta(\mu) \geq \beta(\mu_0) = 1 - \alpha$. Thus, this test does have a significance level of α for the hypothesis $H_0 : \mu \leq \mu_0$.

JUST AS WE HAVE done for $\mu \leq \mu_0$, it is also possible to introduce a one-sided test that bounds the mean by below. In this case, we would reject the null hypothesis $H_0 : \mu \geq \mu_0$ with significance level α iff $\frac{\sqrt{n}}{\sigma}(\bar{\mathcal{X}} - \mu_0) < -z_\alpha$, and the p-value is equal to the probability that Z gets a value under this test statistic.[◦]

◦Try to write this in symbols, to verify that you understand the notions.

Example 81. A bottled water company claims that its water contains in average less than 8mg of sodium per litre. A sample of 25 1-litre bottles yields an average of 7.5mg of sodium. What can we conclude, with a significance level of 5%, if we know that the standard deviation of sodium content is 1.8mg?

Solution. Remember that the only relevant thing that can be done through hypothesis testing is to *reject* the null hypothesis. Thus, in order to support the claim that the average is less than 8mg, we try to reject the opposite claim; that is, we intend to reject $H_0 : \mu \geq 8$. The test statistic is

$$\sqrt{n}(\bar{\mathcal{X}} - \mu_0)/\sigma = 5(7.5 - 8)/1.8 = -1.389,$$

and hence the p-value is $P[Z < -1.389] = 0.082$. Since this number is greater than 0.05, we cannot reject at 5% significance level the null hypothesis: the evidence is not strong enough to support the claim. \triangle

SO FAR WE HAVE considered a known population variance σ^2 . If the variance is unknown, but we still want to test the hypothesis $H_0 : \mu = \mu_0$ for some constant μ_0 ,¹⁸² we want to reject H_0 if the sample mean is too far from μ_0 , but we need an adequate notion of being far. In the previous case, we used the fact that $\bar{\mathcal{X}}$ is a normal to reject whenever $\left| \frac{\sqrt{n}(\bar{\mathcal{X}} - \mu_0)}{\sigma} \right| > z_{\alpha/2}$. As σ is not known anymore, we use its maximum likelihood estimator S .

¹⁸² This is not a simple hypothesis anymore, because it covers a large space depending on the values of σ .

To achieve a significance level α , recall that if $\mu = \mu_0$, the RV $\frac{\sqrt{n}(\bar{\mathcal{X}} - \mu_0)}{S}$ has a t-distribution with $n - 1$ degrees of freedom. In particular,

$$P \left[-t_{\alpha/2, n-1} < \frac{\sqrt{n}(\bar{\mathcal{X}} - \mu_0)}{S} < t_{\alpha/2, n-1} \right] = 1 - \alpha.$$

Thus, we reject H_0 (with significance level α) iff $\left| \frac{\sqrt{n}(\bar{\mathcal{X}} - \mu_0)}{S} \right| > t_{\alpha/2, n-1}$.

If t is the observed value of the test statistic $T = \frac{\sqrt{n}(\bar{\mathcal{X}} - \mu_0)}{S}$, then the p-value of this test is the probability that $|T|$ exceeds $|t|$.¹⁸³ This p-value tells us the significance levels at which the null hypothesis will be rejected.

¹⁸³ The probability that the absolute value of a t random variable with $n - 1$ degrees of freedom is larger than $|t|$.

Example 82. A worried neighbour claims that students drink an average of 3 litres of beer every night. To investigate this claim, 25 randomly selected students are observed. The observations yield a sample mean of 2.91l and a sample standard deviation of 0.47l. To verify the claim, we test the hypothesis $H_0 : \mu = 3$. The test statistic is

$$T = \frac{\sqrt{n}(\bar{\mathcal{X}} - \mu_0)}{S} = -\frac{5 \cdot 0.09}{0.47} = 0.9574.$$

The p-value for this test data is

$$P[|T_{24}| > 0.9574] = 2P[T_{24} > 0.9574] = 0.3479.$$

In other words, the only way to reject this hypothesis is to incur in a very high significance level (that is, increase the probability of error of type I). This experiment is consistent with the hypothesis. \triangle

One-sided hypothesis tests are built in the obvious way. Hence, the hypothesis $H_0 : \mu \leq \mu_0$ will be rejected iff $\frac{\sqrt{n}(\bar{\mathcal{X}} - \mu_0)}{S} > t_{\alpha, n-1}$.

Equality of Means of Normal Populations

OFTEN, WE WANT TO compare two approaches; for example, whether two software systems are equally efficient. This is usually verified by testing whether two normal populations have the same mean.

Let $\mathcal{X}_1, \dots, \mathcal{X}_n$ and $\mathcal{Y}_1, \dots, \mathcal{Y}_m$ be two independent samples from normal RVs having unknown means $\mu_{\mathcal{X}}, \mu_{\mathcal{Y}}$, and known variances $\sigma_{\mathcal{X}}^2, \sigma_{\mathcal{Y}}^2$. We want to test the hypothesis $H_0 : \mu_{\mathcal{X}} = \mu_{\mathcal{Y}}$. Obviously, $\bar{\mathcal{X}} - \bar{\mathcal{Y}}$ is an estimator for $\mu_{\mathcal{X}} - \mu_{\mathcal{Y}}$. Rewriting the null hypothesis as $H_0 : \mu_{\mathcal{X}} - \mu_{\mathcal{Y}} = 0$, we would like to reject H_0 if $\bar{\mathcal{X}} - \bar{\mathcal{Y}}$ is far from zero; i.e., if $|\bar{\mathcal{X}} - \bar{\mathcal{Y}}| > c$ from some suitable c .

Since \mathcal{X} and \mathcal{Y} are independent normal distributions, we know that $\bar{\mathcal{X}} - \bar{\mathcal{Y}} \sim \mathcal{N}\left(\mu_{\mathcal{X}} - \mu_{\mathcal{Y}}, \frac{\sigma_{\mathcal{X}}^2}{n} + \frac{\sigma_{\mathcal{Y}}^2}{m}\right)$; ¹⁸⁴ or equivalently,

$$\frac{\bar{\mathcal{X}} - \bar{\mathcal{Y}} - (\mu_{\mathcal{X}} - \mu_{\mathcal{Y}})}{\sqrt{\frac{\sigma_{\mathcal{X}}^2}{n} + \frac{\sigma_{\mathcal{Y}}^2}{m}}} \sim \mathcal{N}(0, 1).$$

If H_0 is true, then $\mu_{\mathcal{X}} - \mu_{\mathcal{Y}} = 0$ and so $(\bar{\mathcal{X}} - \bar{\mathcal{Y}}) / \sqrt{\frac{\sigma_{\mathcal{X}}^2}{n} + \frac{\sigma_{\mathcal{Y}}^2}{m}} \sim Z$. Then,

$$P\left[-z_{\alpha/2} \leq \frac{\bar{\mathcal{X}} - \bar{\mathcal{Y}}}{\sqrt{\frac{\sigma_{\mathcal{X}}^2}{n} + \frac{\sigma_{\mathcal{Y}}^2}{m}}} \leq z_{\alpha/2}\right] = 1 - \alpha.$$

In other words, for a significance level α , we reject $H_0 : \mu_{\mathcal{X}} = \mu_{\mathcal{Y}}$ iff ¹⁸⁵

$$\frac{|\bar{\mathcal{X}} - \bar{\mathcal{Y}}|}{\sqrt{\frac{\sigma_{\mathcal{X}}^2}{n} + \frac{\sigma_{\mathcal{Y}}^2}{m}}} \geq z_{\alpha/2}.$$

Following the same idea, the hypothesis $H_0 : \mu_{\mathcal{X}} \leq \mu_{\mathcal{Y}}$ will be rejected with significance level α iff

$$\bar{\mathcal{X}} - \bar{\mathcal{Y}} \geq z_{\alpha} \sqrt{\frac{\sigma_{\mathcal{X}}^2}{n} + \frac{\sigma_{\mathcal{Y}}^2}{m}}.$$

SUPPOSE NOW THAT THE variances of the two normals \mathcal{X} and \mathcal{Y} are unknown but equal to a value σ^2 . ¹⁸⁶ To find an adequate critical region, we

¹⁸⁴ Remember that the sample of \mathcal{X} has size n and that for \mathcal{Y} has size m .

¹⁸⁵ The test statistic (which is useful for computing the p-value) is the left-hand side of this inequality.

¹⁸⁶ That is, $\sigma^2 = \sigma_{\mathcal{X}}^2 = \sigma_{\mathcal{Y}}^2$; see note 151.

first estimate the value of σ^2 . In this case, we can use the sample variances

$$S_{\mathcal{X}}^2 = \frac{\sum_{i=1}^n (\mathcal{X}_i - \overline{\mathcal{X}})^2}{n-1}$$

$$S_{\mathcal{Y}}^2 = \frac{\sum_{i=1}^m (\mathcal{Y}_i - \overline{\mathcal{Y}})^2}{m-1},$$

and the fact that¹⁸⁷

$$\frac{\overline{\mathcal{X}} - \overline{\mathcal{Y}} - (\mu_{\mathcal{X}} - \mu_{\mathcal{Y}})}{\sqrt{S_p^2(1/n + 1/m)}} \sim t_{n+m-2},$$

where S_p^2 is the pooled estimator of σ^2 .¹⁸⁸ Then, if the null hypothesis $H_0 : \mu_{\mathcal{X}} = \mu_{\mathcal{Y}}$ is true, the statistic $\frac{\overline{\mathcal{X}} - \overline{\mathcal{Y}}}{\sqrt{S_p^2(1/n + 1/m)}}$ has a t-distribution with $n + m - 2$ degrees of freedom. Thus, we reject this hypothesis iff

$$\frac{|\overline{\mathcal{X}} - \overline{\mathcal{Y}}|}{\sqrt{S_p^2(1/n + 1/m)}} \geq t_{\alpha/2, n+m-2}.$$

The p-values and the one-sided hypothesis tests can be derived in an analogous manner.

¹⁸⁷ See page 57.

¹⁸⁸ $S_p^2 = \frac{(n-1)S_{\mathcal{X}}^2 + (m-1)S_{\mathcal{Y}}^2}{n+m-2}$.

Regression

IN HYPOTHESIS TESTING AND parameter estimation, we are concerned with understanding a single random variable.¹⁸⁹ In practical applications, we often want to understand the relationship between two or more variables. For instance, the relationship between age, type of work, and the number of accidents they have in a year.

In regression, we consider a single *response* variable (also known as *dependent* variable) \mathcal{Y} that depends on a set of *input* (or *independent*) variables $\mathcal{X}_1, \dots, \mathcal{X}_n$.¹⁹⁰ Our goal is to understand this dependency; e.g., to predict, given the age and job description of a person, their likelihood of having an accident. The simplest such relationship is linear. That is, where $\mathcal{Y} = \beta_0 + \beta_1\mathcal{X}_1 + \dots + \beta_n\mathcal{X}_n$ for some $\beta_i \in \mathbb{R}, 0 \leq i \leq n$. If this relationship was precise, then from $n+1$ data points we could compute the exact values of the β_i s, and from them predict the precise response given the values of the input variables. In practice, we expect the response to be affected by a random error. That is, we have

$$\mathcal{Y} = \beta_0 + \beta_1\mathcal{X}_1 + \dots + \beta_n\mathcal{X}_n + \varepsilon,$$

where ε is a RV with mean 0. That is, $E[\mathcal{Y} \mid \vec{\mathcal{X}}] = \beta_0 + \beta_1\mathcal{X}_1 + \dots + \beta_n\mathcal{X}_n$.¹⁹¹

This defines the *linear regression* of \mathcal{Y} on the independent variables \mathcal{X}_i . In this case, the values $\beta_i, 0 \leq i \leq n$ are called the *regression coefficients*.¹⁹² The question now is how to obtain these coefficients.

Least Squares Method

WE ARE INTERESTED IN estimating the parameters of a linear regression model. For a simple regression problem with estimated coefficients $\alpha = a$, $\beta = b$, and an observation x_i of the input variable, the response should be $y = a + bx_i$. Due to the random error, the response will in fact be a value y_i that may differ from this prediction. Given a set of n data points, we want to choose the estimator that minimises the sum of the squared errors. That is, we want to minimise $SS = \sum_{i=1}^n (y_i - a - bx_i)^2$.¹⁹³ To achieve this, we equate the partial derivatives of SS w.r.t. the variables a and b to 0.

$$\begin{aligned} 0 &= \frac{\partial SS}{\partial a} = -2 \sum_{i=1}^n (y_i - a - bx_i) \\ 0 &= \frac{\partial SS}{\partial b} = -2 \sum_{i=1}^n x_i (y_i - a - bx_i), \end{aligned}$$

¹⁸⁹ The only exception we have seen is in testing the equality of means of two RVs. However, even in that case, we only check whether the means are equal or not, without trying to understand the relation between the two.

¹⁹⁰ Notice the abuse on the naming. Independence in this case refers to the response, and has nothing to do with probabilistic independence.

¹⁹¹ $\vec{\mathcal{X}}$ denotes the whole set $\mathcal{X}_1, \dots, \mathcal{X}_n$.

¹⁹² If there is only one independent variable $\mathcal{Y} = \alpha + \beta\mathcal{X} + \varepsilon$, we speak of a *simple regression*.

¹⁹³ The name SS comes from “Sum of Squares.”

or equivalently,

$$\sum_{i=1}^n y_i = na + b \sum_{i=1}^n x_i, \quad \sum_{i=1}^n x_i y_i = a \sum_{i=1}^n x_i + b \sum_{i=1}^n x_i^2.$$

These are known as the *normal equations*. Substituting $\bar{y} = \sum_{i=1}^n y_i / n$, and $\bar{x} = \sum_{i=1}^n x_i / n$, the first equation becomes $a = \bar{y} - b\bar{x}$. We can now substitute this value in the second equation to get

$$\sum_{i=1}^n x_i y_i = (\bar{y} - b\bar{x}) \sum_{i=1}^n x_i + b \sum_{i=1}^n x_i^2 = \bar{y}n\bar{x} - b\bar{x}n\bar{x} + b \sum_{i=1}^n x_i^2,$$

and hence

$$b = \frac{\sum_{i=1}^n x_i y_i - n\bar{x}\bar{y}}{\sum_{i=1}^n x_i^2 - n\bar{x}^2}.$$

The line $a + bx$, where a, b are computed from data using these equations is called the *estimated regression line*.

Estimator Distribution

SO FAR WE HAVE only assumed that the random error has mean 0. To study the properties of the estimators, we further assume that these errors are independent normal RVs with variance σ^2 . This means that given the input values \mathcal{X}_i the responses \mathcal{Y}_i are independent normal RV with mean $\alpha + \beta\mathcal{X}_i$ and variance σ^2 .¹⁹⁴

The least squares estimator b of β can be expressed as

$$b = \frac{\sum_{i=1}^n (x_i - \bar{x}) y_i}{\sum_{i=1}^n x_i^2 - n\bar{x}^2},$$

which is a linear transformation of independent normal random variables \mathcal{Y}_i , and hence also a normal RV. We compute its mean and variance.¹⁹⁵

$$\begin{aligned} E[b] &= \frac{\sum_{i=1}^n (x_i - \bar{x}) E[y_i]}{\sum_{i=1}^n x_i^2 - n\bar{x}^2} = \frac{\sum_{i=1}^n (x_i - \bar{x}) (\alpha + \beta x_i)}{\sum_{i=1}^n x_i^2 - n\bar{x}^2} \\ &= \frac{\alpha \sum_{i=1}^n (x_i - \bar{x}) + \beta \sum_{i=1}^n x_i (x_i - \bar{x})}{\sum_{i=1}^n x_i^2 - n\bar{x}^2} \\ &= \beta \frac{\sum_{i=1}^n x_i^2 - \bar{x} \sum_{i=1}^n x_i}{\sum_{i=1}^n x_i^2 - n\bar{x}^2} \\ &= \beta \end{aligned} \quad (\dagger)$$

$$\begin{aligned} Var(b) &= \frac{Var(\sum_{i=1}^n (x_i - \bar{x}) y_i)}{(\sum_{i=1}^n x_i^2 - n\bar{x}^2)^2} = \frac{\sum_{i=1}^n (x_i - \bar{x})^2 Var(y_i)}{(\sum_{i=1}^n x_i^2 - n\bar{x}^2)^2} \\ &= \frac{\sigma^2 \sum_{i=1}^n (x_i - \bar{x})^2}{(\sum_{i=1}^n x_i^2 - n\bar{x}^2)^2} \\ &= \frac{\sigma^2}{\sum_{i=1}^n x_i^2 - n\bar{x}^2} \end{aligned} \quad (\ddagger)$$

In particular this means that b can serve as an estimator of β .¹⁹⁶




Figure 22: Some linearly related data, with its estimated regression line.

¹⁹⁴ Notice that the variance σ^2 does not depend on the input variable, but is a constant associated to the error distribution. This variance is not necessarily known, and we must estimate it from the data.

¹⁹⁵ Recall that $\sum_{i=1}^n (x_i - \bar{x}) = 0$ (\dagger). In addition, $\sum_{i=1}^n (x_i - \bar{x})^2 = \sum_{i=1}^n x_i^2 - n\bar{x}^2$ (\ddagger).

¹⁹⁶ An estimator whose expected value corresponds to true value being estimated is called *unbiased*.

Recall that $a = \sum_{i=1}^n y_i / n - b\bar{x}$. Then a is also a linear combination of independent normal random variables, and hence is also a normal. As we did with b , we analyse the parameters of this RV.[◦]

◦ Check out the variance by yourself.

$$E[a] = \sum_{i=1}^n E[y_i] / n - \bar{x}E[b] = \sum_{i=1}^n (\alpha + \beta x_i) / n - \beta \bar{x} = \alpha + \beta \bar{x} - \beta \bar{x} = \alpha,$$

$$Var(a) = \frac{\sigma^2 \sum_{i=1}^n x_i^2}{n(\sum_{i=1}^n x_i^2 - n\bar{x}^2)}.$$

THE DIFFERENCES BETWEEN THE responses and their least square estimators $y_i - a - bx_i$, $1 \leq i \leq n$, are called *residuals*. To estimate the error variance σ^2 we use the sum of squares of the residuals

$$SS_R = \sum_{i=1}^n (y_i - a - bx_i)^2.$$

It can be shown that SS_R is independent from A and B and that SS_R/σ^2 is a chi-square with $n-2$ degrees of freedom.¹⁹⁷ Thus $E[SS_R/\sigma^2] = n-2$, or equivalently $E[SS_R/(n-2)] = \sigma^2$.

¹⁹⁷ The actual proof of these facts is outside the scope of this course.

To summarise, if the responses y_i , $1 \leq i \leq n$ are normally distributed with mean $\alpha + \beta x_i$ and (common) variance σ^2 , then the least squares estimates for α and β are $a = \bar{y} - b\bar{x}$ and $b = \frac{\sum_{i=1}^n (x_i - \bar{x})y_i}{\sum_{i=1}^n x_i^2 - n\bar{x}^2}$, respectively. These estimators are normally distributed as well:

$$a \sim \mathcal{N}\left(\alpha, \frac{\sigma^2 \sum_{i=1}^n x_i^2}{n(\sum_{i=1}^n x_i^2 - n\bar{x}^2)}\right),$$

$$b \sim \mathcal{N}\left(\beta, \frac{\sigma^2}{\sum_{i=1}^n x_i^2 - n\bar{x}^2}\right).$$

Moreover, the sum of squares of residuals is a chi-square with $n-2$ degrees of freedom, and is independent from a and b .

Statistical Inferences

WHEN CONSIDERING THE REGRESSION model $\mathcal{Y} = \alpha + \beta\mathcal{X} + \varepsilon$, an important question to ask is whether $\beta = 0$.¹⁹⁸ Thus, we want to test the null hypothesis $H_0 : \beta = 0$ vs. the alternative hypothesis $H_1 : \beta \neq 0$. Recall from our previous discussion that $(b-\beta)/\sqrt{\sigma^2/\sum_{i=1}^n x_i^2 - n\bar{x}^2} \sim \mathcal{N}(0, 1)$ ¹⁹⁹ and is independent from $SS_R/\sigma^2 \sim \chi_{n-2}^2$. Thus, we get that the variable

¹⁹⁸ If this is the case, then the response does not depend on the input variable.

¹⁹⁹ $b \sim \mathcal{N}\left(\beta, \frac{\sigma^2}{\sum_{i=1}^n x_i^2 - n\bar{x}^2}\right)$.

$$\frac{\sqrt{\sum_{i=1}^n x_i^2 - n\bar{x}^2}(b-\beta)/\sigma}{\sqrt{\frac{SS_R}{\sigma^2(n-2)}}} = \sqrt{\frac{(n-2)\sum_{i=1}^n x_i^2 - n\bar{x}^2}{SS_R}}(b-\beta)$$

has a t -distribution with $n-2$ degrees of freedom.²⁰⁰

If the null hypothesis is true (i.e., $\beta = 0$) then

$$\sqrt{\frac{(n-2)\sum_{i=1}^n x_i^2 - n\bar{x}^2}{SS_R}}b \sim t_{n-2}.$$

²⁰⁰ Notice that we could not use the normal distribution observed before because it depends on the unknown σ^2 . Instead, we estimate this variance through $SS_R/(n-2)$, but this requires the use of a t -distribution.

Thus, we reject H_0 with significance γ ²⁰¹ if $\sqrt{\frac{(n-2)S_{xx}}{SS_R}}|b| > t_{\gamma/2, n-2}$. The

²⁰¹ We use here γ to avoid confusions with the regression parameter α .

p -value can be computed as usual: from the test statistic $v = \sqrt{\frac{(n-2)S_{xx}}{SS_R}}|b|$, the p -value is $P[|T_{n-2}| > v] = 2P[T_{n-2} > v]$.

To get a confidence interval estimator for β we proceed as usual. Given $0 < \gamma < 1$, we know that

$$P\left[-t_{\gamma/2, n-2} \leq \sqrt{\frac{(n-2)S_{xx}}{SS_R}}(b - \beta) \leq t_{\gamma/2, n-2}\right] = 1 - \gamma.$$

Solving for β yields the $100(1 - \gamma)\%$ confidence interval estimator of β [◦]

$$\left(B - \sqrt{\frac{SS_R}{(n-2)S_{xx}}} t_{\gamma/2, n-2}, B + \sqrt{\frac{SS_R}{(n-2)S_{xx}}} t_{\gamma/2, n-2}\right).$$

◦Try the computations yourself.

IF ONE IS INSTEAD interested in inferences concerning the parameter α , one can proceed in the exact same manner. Indeed, we observe that

$$\sqrt{\frac{n(n-2)S_{xx}}{SS_R \sum_{i=1}^n x_i^2}}(a - \alpha) \sim t_{n-2}.$$

From this fact, we can derive confidence intervals, and hypothesis tests, alongside their p -values.[◦]

◦Try it yourself!

SUPPOSE NOW THAT WE have a given input value x_0 and we want to produce a confidence interval or test a hypothesis about the mean response w.r.t. this value. Since we do not know the parameters α, β , we can estimate them through a and b as before. But to obtain a meaningful confidence interval, we must understand the distribution of $a + bx_0$.

Note that $a + bx_0 = \bar{y} - b(\bar{x} - x_0) = \sum_{i=1}^n y_i (1/n - c(x_i - \bar{x})(\bar{x} - x_0))$, where c is a constant.²⁰² Each y_i is an independent observation of a normal random variable, and all the other parameters in the expression are constants given the data. Hence, $a + bx_0$ is a linear combination of independent normal random variables, which means that it is also a normal itself. Now we just need to know its mean and its variance.

²⁰² Recall that

$$a = \bar{y} + b\bar{x}$$

$$b = \frac{\sum_{i=1}^n (x_i - \bar{x})y_i}{\sum_{i=1}^n x_i^2 - n\bar{x}^2}.$$

The denominator of b is a constant that we will call c in the following analysis.

$$E[a + bx_0] = E[a] + x_0 E[b] = \alpha + \beta x_0$$

$$\begin{aligned} Var(a + bx_0) &= \sum_{i=1}^n (1/n - c(x_i - \bar{x})(\bar{x} - x_0))^2 Var(y_i) \\ &= \sigma^2 \sum_{i=1}^n (1/n^2 - 2c(x_i - \bar{x})(\bar{x} - x_0)/n + c^2(x_i - \bar{x})^2(\bar{x} - x_0)^2) \\ &= \sigma^2 \left(1/n - 2c(\bar{x} - x_0)/n \sum_{i=1}^n (x_i - \bar{x}) + c^2(\bar{x} - x_0)^2 \sum_{i=1}^n (x_i - \bar{x})^2 \right). \end{aligned}$$

Since $\sum_{i=1}^n (x_i - \bar{x}) = 0$ and $\sum_{i=1}^n (x_i - \bar{x})^2 = \sum_{i=1}^n x_i^2 - n\bar{x}^2 = 1/c$, it then follows that

$$Var(a + bx_0) = \sigma^2 (1/n + c(\bar{x} - x_0)^2).$$

In brief, $a + bx_0 \sim \mathcal{N}(\alpha + \beta x_0, \sigma^2 (1/n + c(\bar{x} - x_0)^2/S_{xx}))$. To handle the unknown σ^2 , recall that $SS_R/\sigma^2 \sim \chi_{n-2}^2$. Thus

$$\frac{a + bx_0 - (\alpha + \beta x_0)}{\sqrt{1/n + c(\bar{x} - x_0)^2/S_{xx}} \sqrt{SS_R/n-2}} \sim t_{n-2}.$$

Using this fact, we can then build confidence intervals and hypothesis tests for $\alpha + \beta x_0$ as usual.[◦]

◦Exercise!

Quality of the Model

WHEN TRYING TO MEASURE the variability of the response variable, one can use the squared error $S_{yy} = \sum_{i=1}^n (y_i - \bar{y})^2$. In this case, however, there are two different factors that influence the variability of the response. On the one hand, each value y_i depends on the input value x_i that may be different for each i . On the other hand, the response is a random variable with variance σ^2 .²⁰³

An important question is how much of this variation is caused by the different input values, and how much depends on the inherent uncertainty of the response. To answer this question, observe that the value $SS_R = \sum_{i=1}^n (y_i - a - bx_i)^2$ can be interpreted as the remaining squared variation once that the influence of the input values has been taken into account. Thus $S_{yy} - SS_R$ measures the amount of variation that is explained by the input values. We can then define the *coefficient of determination* $R^2 := \frac{S_{yy} - SS_R}{S_{yy}}$, which expresses the proportion of variation explained by the input values.

Obviously, R^2 is always between 0 and 1. Intuitively, a high coefficient of determination (close to 1) indicates that most of the variation results from the different input values, and the opposite holds for a low R^2 . This value is used as an indicator of the fitness of the model: the higher R^2 is, the regression model is considered to fit well the data.

NOTICE THAT LINEAR REGRESSION can be applied to any piece of data, regardless of its shape. However, the quality of the model requires a linear dependency between the variables. Sometimes a visual exploration suffices to rule out the possibility of a linear model (see Figure 23). However, a direct visualisation cannot deal with more veiled cases of non-linearity.

As a means to assess the quality of the model, one can analyse the residuals $y_i - (a + bx_i)$. Remember that each y_i is a normal; $a + bx_i$ estimates its mean, and $SS_R/(n-2)$ estimates its variance.²⁰⁴ Then, the *standardised residuals* $\frac{y_i - (a + bx_i)}{\sqrt{SS_R/(n-2)}}$ should all be approximately distributed as a standard normal RV, and they are all independent. Then, approximately 95% of these residuals should lie in $(-2, 2)$,²⁰⁵ and they should show a *random* behaviour without any obvious patterns. Figure 24 plots the standardised residuals from the first two data sets of Figure 23. The second plot shows

²⁰³ Even with identical values on the input variable, the responses may be different.




Figure 23: Three data samples with their linear regression fitting model. For the third plot, the linear model is obviously inadequate.

²⁰⁴ Recall the assumption that the variance of the error is constant for all input values.

²⁰⁵ $P[-1.96 < Z < 1.96] = 0.95$.




Figure 24: Residuals for the first two plots in Figure 23. The first one appears random, and hence is a good fit for the linear model. The second shows a pattern that suggests dependencies between the residuals.

a pattern that suggests that the linear model is not adequate for this data.

Non-linear Responses

EVEN WHEN THE RESPONSE does not depend linearly on the input variable, if the type of dependency is known one can sometimes transform the variables to obtain a linear model. The usual scenario is when the response variable grows exponentially on the input, that is, when the relationship is of the form $w = c \cdot d^x$, where c and d are two constants that we want to estimate.²⁰⁶ By taking the logarithm in both sides, we get $\log(w) = \log(c) + x \log(d)$. This now looks like the linear model $y = \alpha + \beta x$.

At this point, we can use a standard linear regression process (that is, least squares) to estimate the parameters α and β with a and b , respectively. Then, $w = \exp\{a + bx\}$.

FOR OUR REGRESSION MODEL, we have assumed that the variance of the the response is fixed for any input value. Often, a more realistic assumption is that these variances are dependent on the input variable, but known up to a proportionality constant. More formally, we know that for each piece of data, $\text{Var}(y_i) = \sigma^2/w_i$. As before, to obtain a linear model, we want to minimise the sum of squares, but in this case the weight of each error should be proportional to the variance.²⁰⁷ That is, we choose the estimators a, b that minimise the sum

$$\sum_{i=1}^n \frac{(y_i - (a + bx_i))^2}{\text{Var}(y_i)} = \frac{1}{\sigma^2} \sum_{i=1}^n w_i (y_i - a - bx_i)^2.$$

To find these values, we differentiate with respect to a and b and equate the derivatives to 0, as usual. Hence, we obtain

$$\begin{aligned} \sum_{i=1}^n w_i y_i &= a \sum_{i=1}^n w_i + b \sum_{i=1}^n w_i x_i, \\ \sum_{i=1}^n w_i x_i y_i &= a \sum_{i=1}^n w_i x_i + b \sum_{i=1}^n w_i x_i^2. \end{aligned}$$

Solving these equations, we obtain the *weighted* least square estimators.

In general, it is not always obvious to know how does the variance depend on the input variable. Still, there are many cases where this can be reasonably expected. A simple example is if one is trying to understand the relationship of water used in a house, dependent on the number of people living in it. Assuming that the amount of water used by each person is independent with a fixed variance, then the variance of use in a household will be proportional to the number of people in the house. Sometimes, the variance dependency can be understood by an analysis of the scatter plot of the data and its standardised residuals, as depicted in Figure 25.

AS WE HAVE ALREADY seen, the relationship between the independent and the response variables may not always be linear. Still, the dependency may follow a polynomial of the form $y = \beta_0 + \beta_1 x + \beta_2 x^2 + \dots + \beta_r x^r + \varepsilon$. In this case, we want to estimate the regression coefficients β_i . Generalising the idea of linear regression, given a data sample of size n , we compute the least square estimators b_0, \dots, b_r of β_0, \dots, β_r , which are the values that

²⁰⁶ Of course, one of them might be known, but the more general setting is where they are both unknown.

²⁰⁷ It should be clear that for the data points where the variance of y is larger, we should expect to observe larger errors. To handle all the data uniformly, we weight each point proportionally to their variance.




Figure 25: The variance of the data (observed by the margin of error) seems to increase with x .

minimise


$$\sum_{i=1}^n (y_i - \beta_0 - \beta_1 x - \dots - \beta_r x^r)^2.$$

Following the usual approach,²⁰⁸ we obtain the so-called *normal equations*

$$\begin{aligned} \sum_{i=1}^n y_i &= b_0 n + b_1 \sum_{i=1}^n x_i + b_2 \sum_{i=1}^n x_i^2 + \dots + b_r \sum_{i=1}^n x_i^r \\ \sum_{i=1}^n x_i y_i &= b_0 \sum_{i=1}^n x_i + b_1 \sum_{i=1}^n x_i^2 + b_2 \sum_{i=1}^n x_i^3 + \dots + b_r \sum_{i=1}^n x_i^{r+1} \\ \sum_{i=1}^n x_i^2 y_i &= b_0 \sum_{i=1}^n x_i^2 + b_1 \sum_{i=1}^n x_i^3 + b_2 \sum_{i=1}^n x_i^4 + \dots + b_r \sum_{i=1}^n x_i^{r+2} \\ &\vdots \\ \sum_{i=1}^n x_i^r y_i &= b_0 \sum_{i=1}^n x_i^r + b_1 \sum_{i=1}^n x_i^{r+1} + b_2 \sum_{i=1}^n x_i^{r+2} + \dots + b_r \sum_{i=1}^n x_i^{2r} \end{aligned}$$

The most important question to answer in polynomial regression is what degree of a polynomial to use. This requires a careful analysis regarding the trade-off between fitting and prediction. On the one hand, a polynomial with a higher degree will be able to fit the data better, producing smaller residuals.²⁰⁹ However, recall that the values of y are subject to a random noise. A better fitness in this case will mean fitting the noise as well, which may reduce the predictive power of the model.²¹⁰ As a rule of thumb, you want to have the smallest degree that matches the general shape observed in the scatter plot of the data.

To understand this, consider again the third plot from Figure 23. A visual inspection suggests that the data follows a quadratic pattern. Thus, we may try to fit it through a quadratic regression method. The resulting curve, on the left of Figure 26, provides a good approximation to the observed data. A polynomial of degree 15 (Figure 26, right), fits the observed data much better, but makes some unreasonable predictions around the limit values.



²⁰⁸ Equating the partial derivatives to 0.

²⁰⁹ In principle, a polynomial of degree n can perfectly fit a data sample of size n if all the values of the independent variable are different.

²¹⁰ This is known as *overfitting*.




Figure 26: Data from Figure 23 fitted through polynomial regression of degree 2 (left) and 15 (right). The fitting curve is represented by a dashed line. Notice that the polynomial of degree 15 approximates the data better, but predicts the response for values of the independent variable around 1.5 and around 19.5 to be far away from all observations. This is a clear case of over-fitting.