2. BASIC EVENT PROBABILITY MODELS

2.1 Overview

This chapter introduces the models used for basic events and for initiating events. This first section is an overview, and the remaining sections of the chapter give more details.

Probabilistic risk assessment (PRA) considers various possible accident sequences. An accident sequence begins with some **initiating event**, which challenges the safety of the plant. Typically one or more standby safety systems are then demanded, and other, normally operating systems must continue operating to ensure that no serious undesirable consequences occur. For the systems to fail to bring the situation under control, several components must either fail or be unavailable. The logic events in the PRA model that represent these failures or modes of unavailability are called **basic events**.

It is not possible to predict precisely when an initiating event or a component failure will occur, because the processes that lead to their occurrences are complex. Therefore, the initiating events and basic events are modeled as resulting from random processes.

The first step in the data analysis task is, therefore, to determine the appropriate probability model to represent the initiating event or basic event. These probability models typically have one or more parameters. Thus, the next major step is to estimate the values of these parameters. This estimation is based on the most applicable and available data. This process of choosing data sources, extracting the data in an appropriate form, and using it to estimate the parameters is the main subject of this handbook.

Basic events are customarily divided into unavailability (because the equipment is undergoing testing or maintenance), failure to start or change state, and failure to run (after successfully starting) or maintain state to the end of the required mission time. Unavailability and failure to run are each modeled in a single way. On the other hand, two different probability models have been used to represent a failure to start or to change state. The first, and more commonly used, method is to model the failures as having a constant probability of **failure on a demand**. The second method is to model the failures as occurring, in an unrevealed way, randomly in time. The failed condition is then discovered at the time of the demand. This is usually called the **standby failure-rate model**. Both models will be discussed.

The above events are the typical ones considered in a PRA. In addition, however, one must occasionally analyze durations, such as the time to restore offsite power or time to recover a failed component. Although such an analysis is not needed for a typical accident sequence, it is discussed in this handbook.

In summary, five topics are considered in the rest of this chapter:

- initiating events
- failures to start or change state (modeled in two possible ways)
- failures to run or maintain state
- unavailability from being out of service
- durations

These topics are the subjects of Sections 2.2 through 2.6. Each section begins with examples of the data that might be analyzed. This is followed by a brief subsection presenting the assumptions of the usual model for the random process (the result of underlying physical mechanisms) and a describing the kind of data that can be observed. The next subsection summarizes the data required to estimate the model parameter(s). The example data sets are then examined in the light of the model assumptions. These examinations illustrate the kind of thinking necessary for the data analyst. Finally, the section may conclude with a short discussion of related issues.

As a preview, Table 2.1 indicates the models, the parameters, and the data needed for each of the topics in the above five bullets. The top line of the table also indicates which section of Chapter 2 treats the topic.

Table 2.1 Kinds of models considered.

2.2 Initiating Events	2.3 Failures to Start or Change State (2 models)		2.4 Failures to Run or Maintain State	2.5 Unavailability	2.6 Durations
Typical Event					
Event occurs initiating accident sequence	Standby system fails on demand		System in operation fails to run, or component changes state during mission	System is unavailable, intentionally out of service, when demanded	A condition persists for a random time period
Parameter(s) to Estimate					
λ, event frequency	For failure on demand: p, probability of failure on demand	For standby failure: λ, rate of occurrence of standby failures	λ, rate of failure to run	u, fraction of time when component will be out of service	Parameters of assumed probability distribution of duration time
Data Required to Estimate Parameters[a]					
Number of events, x, in total time, t	Number of failures, x, in total number of demands, n	Number of failures, x, in total standby time, t	Number of failures, x, in total running time, t	Observed fractions of time when component out of service; OR durations of observed out-of-service events	Depends on model, but typically the lengths of the observed durations
[a]. The data here are the minimal requirements to estimate the parameter. More detailed data are needed to check the model assumptions.					

The term **system** is used in Table 2.1 and from here on to denote the set of hardware for which data are collected; it may be an entire nuclear power plant (NPP), or a system in the traditional sense, such as the auxiliary feedwater (AFW) system, or a train, component, or even piece part. This reduces the need for phrases such as “system or component.”

The lengthiest part of each section below consists of examination of examples to see whether the assumptions of the probability model appear to be satisfied. Verifying model assumptions is an important part of good data analysis. Ways to investigate the appropriateness of assumptions are considered in Chapter 6, along with parameter estimation. The present chapter, however, only introduces the assumptions and illustrates their meanings through examples. If the assumptions are clearly not satisfied, some mention is given of ways to generalize the model, although such generalizations are not presented until Chapters 7 and 8 in this handbook.

Also, examples and extended discussion of examples are printed in Arial font, to distinguish them from the more general material.

2.2 Initiating Events

2.2.1 Examples

In the context of a nuclear power plant PRA, an initiating event is any event that perturbs the steady state operation of the plant thereby initiating an abnormal event such as a transient or loss-of-coolant accident within a plant. Initiating events begin sequences of events that challenge plant control and safety systems. Failure of these systems can lead to core damage and a release of radioactivity to the environment. However, the consideration of the potential plant response to initiating events, is irrelevant when estimating their frequencies.

Here are several examples of data sets counting such initiating events.

Example 2.1 Unplanned reactor trips

A U.S. commercial nuclear power plant had 34 unplanned reactor trips in 1987 through 1995. It had its initial criticality on Jan. 3, 1987, and experienced a total of 64651 critical hours, or 7.38 critical years (Poloski et al. 1999a).

Example 2.2 Shutdown loss of offsite power

In U.S. commercial nuclear power plants in 1980-1996, there were 80 plant-centered loss-of-offsite-power (LOSP) events during shutdown. In that period, the plants experienced 455.5 reactor-shutdown years (Atwood et al. 1998).

Example 2.3 Through-wall pipe leaks

In world-wide experience of western-style PWRs (3362 calendar years of operation), a single through-wall leak event has been reported in large-diameter piping (Poloski et al. 1999a, Appendix J).

Example 2.4 Temperature sensor/ transmitters

Eide et al. (1999) report that temperature sensor/ transmitters in the reactor protection system (RPS) of Westinghouse NPPs had 32 failures in 2264.1 component-years. These sensor/transmitters operate continuously, and when they fail they are repaired or replaced in a relatively short time. The number of failures is conservatively estimated from sometimes incomplete NPRDS data, and the number of component years is based on an estimated number of components per loop.

These examples have several elements in common. First, they involve a **number of events** that occurred, and an **exposure time**, or **time at risk**, when the events could have occurred. The next subsection will give a simple probability model for generating random events in time. In addition, in each of the above examples corrective action is taken after any event, so that the system then resumes operation (the system is **repairable**.) This means that the recorded operating history consists of a sequence of random event occurrences, which is summarized as a count of events in some fixed time. This type of data will direct us to a particular type of analysis, presented in Chapter 6.

The events may be the initiating events of an ordinary PRA (Example 2.1), initiating events of a shutdown PRA (Example 2.2), failures in a passive system (Example 2.3), which incidentally happen also to be initiating events in a PRA, or failures of a continuously running system (Example 2.4), which initiate quick repair action. A PRA analyst would distinguish among the examples based on their safety consequences. The present discussion, however, adopts the viewpoint of probability modeling, in which the important fact is not the consequence of the events, but the way that they occur randomly in time. Reactor trip initiators are the prototypical example of such events, but not the only example.

The exposure time is the length of time during which the events could possibly occur. In the Example 2.1, the exposure time is reactor-critical-years, because a reactor trip can only occur when the reactor is at power. Because only one plant is considered, "critical years" can be used as shorthand for "reactor-critical-years." In Example 2.2, the event of interest is LOSP during shutdown, so the exposure time must be the number of reactor-shutdown-years in the study period. In Example 2.3, reactor-calendar-years are used, primarily because more detailed worldwide data could not be easily obtained. The model therefore assumes that a crack in large-diameter piping could occur with equal probability during operation and during shutdown. The model also does not consider differences between plants, such as differences in the total length of large-diameter piping at a plant. In Example 2.4 the exposure time is the number of component-years, because the components operate constantly.

The possible examples are endless. The events could be unplanned demands for a safety system, forced outage events, or many other kinds of events that resemble initiating events.

The data given in the above examples are expressed in the crudest summary terms, a count of events in a total exposure time. This is sufficient for the simple model of this section. Section 2.6 will consider more sophisticated models using the exact event times.

The data could also be broken down into smaller pieces. For example, the initiating event data could be summarized for each calendar year, with an event count and an exposure time reported separately for each year from 1987 through 1995. This additional information allows one to look for trends or other patterns, as discussed in Chapter 7.

2.2.2 Probability Model

The assumptions concerning the physical process are given here, and a description of the kind of data that can be observed.

It is standard to assume that the event count has a Poisson distribution. As listed in Section A.6.2, the usual assumptions (following Thompson 1981) for a Poisson process are:

1. The probability that an event will occur in any specified short exposure time period is approximately proportional to the length of the time period. In other words, for an interval with exposure time Δt the probability of an occurrence in the interval is approximately $\lambda \times \Delta t$ for some $\lambda > 0$.

2. Exactly simultaneous events do not occur.

3. Occurrences of events in disjoint exposure time periods are statistically independent.

In addition, it is worthwhile to spell out the kind of data that can be observed.

- A random number of events occur in some prespecified, fixed time period. As a minimum, the total number of events and the corresponding time period are observed.

Under the above assumptions, the number of occurrences X in some fixed exposure time t is a Poisson distributed random variable with mean $\mu = \lambda t$,

$$\Pr(X = x) = e^{-\mu}\mu^{x} / x! \ . \qquad (2.1)$$

The **probability distribution function**, or p.d.f., is sometimes used to abbreviate this: $f(x) = \Pr(X = x)$. (Throughout this handbook, upper case letters are used for random variables and lower case letters are used for particular numbers.)

The parameter λ is a **rate** or **frequency**. To make things clearer, the kind of event is often stated, that is, "initiating event rate" in Example 2.1, "through-wall-crack occurrence frequency" in Example 2.3, and so forth. Because the count of events during a fixed period is a unitless quantity, the mean number of occurrences μ is also unitless. However, the rate λ depends on the units for measuring time. In other words, the units of λ are 1 per unit of time, such as 1/year or 1/reactor-critical-hour.

This model is called a **Poisson process**. It is extremely simple, because it is completely specified by the exposure time, t, and the one unknown parameter, λ. Assumption 1 implies that the rate λ does not change over time, neither with a monotonic trend, nor cyclically, nor in any other way. Assumption 2 says that exactly simultaneous events do not occur. The only way that they could occur (other than by incredible coincidence) is if some synchronizing mechanism exists, a common cause. Therefore, the operational interpretation of Assumption 2 is that common-cause events do not occur. Assumption 3 says that the past history does not affect the present. In particular, occurrence of an event yesterday does not make the probability of another event tomorrow either more or less likely. This says that the events do not tend to occur too much in clusters, nor do they tend to be systematically spaced and evenly separated.

2.2.3 Data Needed to Validate Model and Estimate λ

Suppose that the Poisson model holds. Then any reasonable estimator of λ needs only two pieces of information: the total exposure time, t, in the data period, and the number of events, x, that occurred then.

However, more information is needed to investigate whether the Poisson model is valid. For example, the data might cover a number of years or a number of plants, and λ might not be constant over time or the same at all plants. These possibilities are not allowed by the listed model assumptions. To study whether they occur, the times and locations of the initiating events should be recorded, or at least the data should be partitioned into subsets, for example corresponding to plants or years. Then the event count and exposure time, x_i and t_i, should be given for each subset.

2.2.4 Case Studies: Validity of Model Assumptions in Examples

Let us examine the reasonableness of the Poisson model assumptions for Examples 2.1 through 2.4. Chapter 6 will address this issue by performing data analysis. Here we will merely cite the results of published studies and use critical thinking.

Example 2.1 Initiating Events

To make this example more precise, assume that any time interval starts on some date at some time and ends on some date at some time, and that the length of the interval, Δt, is the number of critical years contained between the start and stop of the time interval. For example, if the time period is two 24-hour days and the reactor was critical for half of that time, then Δt = 1/365 critical years. An initiating event is an event with the reactor critical, causing an unplanned reactor trip.

Assumption 1 is violated in two ways. First, in the industry as a whole, and presumably in individual plants, the probability of an initiating event in an interval of length Δt (such as one critical day) has not been constant. Instead, the probability dropped substantially from 1987 to 1995. Equivalently, the event rate, λ, dropped from 1987 to 1995. This violation can be eliminated by considering only a short time period for the study, such as one calendar year instead of nine years. If, however, the whole nine-year period is of interest, a more complicated model must be used, such as one of the trend models described in Chapter 7.

A second violation of Assumption 1 arises because this particular plant was new at the start of the study period, with initial criticality on January 3, 1987, and commercial start on May 2, 1987. Many new plants seem to experience a learning period for initiating events, and this plant had 15 of its 34 initiating events during the first 6 months of 1987. After that initial period with a high event rate, the event rate dropped sharply. This violation of Assumption 1 can be resolved by eliminating data before the plant reached a certain age. That is, do not count either the operating time or the initiating events from the plant until it has reached a certain age — exclude that portion of the plant's history from the universe being studied.

Assumption 2 says that exactly simultaneous initiating events do not occur. This is quite reasonable for events at a single plant.

Assumption 3 says that the probability of an initiating event in one time period does not depend on the presence or absence of an initiating event in any earlier time period. This assumption may be challenged, if the plant personnel learn from the first event, thus reducing the probability of a second event. This kind of dependence of one event on another is not allowed by Assumption 3. Suppose, however, that the learning is modeled as a general kind of learning, so that the event rate decreases over time but not as a clear result of any particular events. This may justify using a Poisson model with a trend in the event rate, as considered in detail in Chapter 7.

One might worry about the finite length of time that a reactor is down after a reactor trip. During that time, no initiating events can occur, because the definition of initiating event requires that the reactor be at power. This is a false worry. During the time when the reactor is down, the plant has dropped out of the study. Its shutdown hours are not counted in the exposure time. Only when the reactor comes up again does it begin contributing hours of exposure time and possible initiating events.

Example 2.2 Shutdown LOSP

Just as with the previous example, consider the three assumptions of the Poisson model. In this case, because data come from the entire industry, λ is interpreted as the average rate for the entire industry.

Consider first Assumption 1. The report that studied this data (Atwood et al. 1998) found no evidence of a trend in the time period 1980 through 1996. It did find evidence of differences between plants, however. These differences can affect the industry average, as plants enter the study when they start up or leave the study when they are decommissioned. When a plant with an especially high or low event rate enters or leaves the study, this will affect the industry average somewhat. However, the event rate at the worst plant differed from the industry average by only a factor of about 3.4, and the best plant differed from the average by less than that. Many plants, 116, were considered. Therefore, the effect of a single plant's startup or decommissioning should be small. Therefore, it appears that the overall industry event rate was approximately constant, as required by Assumption 1.

Assumption 2 rules out exactly simultaneous events. This is not quite true for events at sister units at a single site, because a common cause can result in simultaneous LOSP at both units.

Of the 80 events in the data, two pairs of events occurred together at sister units, each pair from a common cause. Thus, simultaneous events do occur, but they are not frequent. The departure from Assumption 2 is probably not large enough to be serious. Chapter 6 will return to this issue, when it considers goodness of fit to a model.

Assumption 3 requires statistical independence of the number of events in disjoint time intervals. As with Example 2.1, there may be some learning, although the lack of trend indicates that any learning is minimal.

In summary, the assumptions for the Poisson model seem to be approximately satisfied.

Example 2.3 Through-Wall Leaks

This differs from the other examples in that the number of events is very small. Any departures from the Poisson assumptions cannot be seen in the data, because so few events have occurred. With no theoretical reason to postulate a trend or other nonconstancy, or a high rate of multiple events, or dependence between events, we accept the Poisson assumptions. The assumptions may not be perfectly true, and a different model may be more accurate, but the Poisson model is simple, and good enough for analyzing such a sparse data set.

Example 2.4 Temperature Sensor/Transmitters

The report (Eide et al. 1999) divides the total study time into two halves, and finds a difference between λ in 1984-1989 and λ in 1990-1995. The example here is for 1990-1995 only. Within this time period the report does not see strong evidence of a trend. That is, a small trend may be present, but the time period is too short, and the failures too few, for any trend to be clear. Further, because the components are regularly maintained, it is reasonable to assume that the failure rate, λ, is roughly constant, as required by Assumption 1.

Assumption 2 requires that common-cause failures be negligible. However, the report states that 14 of the 32 component failures occurred during 4 common-cause events. Thus, Assumption 2 is seriously violated.

Finally, Assumption 3 requires independence of the number of events in disjoint time intervals. The report does not address this issue, but independence appears plausible.

In summary, the example violates Assumption 2, but probably satisfies the other two assumptions. One way to deal with the violation of Assumption 2 would be to model the independent failures and the common-cause failures separately, although Eide et al. do not do this.

2.2.5 Discussion

2.2.5.1 More General Models

The model considered here is a **homogeneous Poisson process (HPP)**, which has a constant event occurrence rate, λ. The number of events in time t is a Poisson random variable with parameter $\mu = \lambda t$. A generalization is a **nonhomogeneous Poisson process (NHPP)**, in which λ is a function of t. Such a model is useful for analyzing trends. Chapter 6 includes ways to test the assumptions of a homogeneous Poisson process, and Chapter 7 includes ways to analyze data where a trend is present.

When data come from the industry, one may consider the differences between plants. Ways to model such differences are discussed in Section 8.2 of this handbook. The present chapter's interest is restricted to λ when no such variation is present. Of course, if the data come from only one plant, λ refers to that plant and the issue of differences typically does not arise.

Any mathematical model, such as the model for a homogeneous Poisson process given here, is an imperfect approximation of the true process that generated the data. When the data set is sparse (few events in the above examples), (a) it is difficult or impossible to see evidence of departures from the model, and (b) the data set is too small to allow realistic estimation of the parameters of a more complicated model. When the data set has many events, departures from the model become visible, and typically a more complicated model is appropriate. These statements have been illustrated by the small and large data sets given as examples.

2.2.5.2 Constancy of *t*

In the model considered here, the exposure time is treated as fixed, and the number of events is treated as random. This is the normal type of data found in PRA work. In reliability analysis, on the other hand, equipment is sometimes tested until it fails. That is, a predetermined number of items are tested, say *n* items. Each item is run until it fails, and the total running time of the items is random. This is an example of **duration** data, discussed in Section 2.6. The probability model is the Poisson process presented above, but the data collection, and resulting data analysis, are different.

2.3 Failure to Change State

2.3.1 Examples

Here are four examples of failure to change state, three with failure to start and one with failure to close.

Example 2.5 HPCI failures to start

> At 23 BWRs in the 1987-1993 time period, the high pressure coolant injection (HPCI) system had 59 unplanned attempts to start. The system failed to start on 5 of these demands (Grant et al. 1995). The failures were typically erratic starts, which the operator stabilized manually. These demands occurred during 113.94 reactor-critical-years.

Example 2.6 EDG failures to start

> Emergency diesel generators (EDGs) are sometimes demanded because of unplanned loss of power to a safety bus, and they are also tested periodically, with one set of tests during each operating cycle and another set of tests monthly. In addition, a return-to-service test is normally performed after maintenance of an EDG. At one plant over an 18-month time period, the number of such demands is counted, and the number of failures to start is counted.

Example 2.7 Steam binding in AFW

> Between demands, steam binding can develop in the AFW system, so that one or more pumps cannot function when demanded. This is mentioned by Wheeler et al. (1989), and by Nickolaus et al. (1992).

Example 2.8 Failures of isolation valves

> Nickolaus et al. (1992) review the causes of about 45 failures of air-operated and motor-operated isolation valves. Some of the principal causes are corrosion, instrument drift, and moisture in instrument and control circuits. Other causes include contamination and corrosion products in the instrument air system, and debris in the system. These are all conditions that can develop while the valves are not being used.

2.3.2 Failure on Demand

All these examples involve a number of demands and a number of failures, where the terms "demand" and "failure" can be defined according to the purposes of the study. Non-PRA contexts provide many other examples of failures on demand. An early example in elementary probability or statistics courses is tossing a (possibly biased) coin *n* times, and counting the number of heads. Count either a head or a tail as a "failure." Just as in the PRA examples, this example has a number of demands, with a random number of the demands resulting in failures.

2.3.2.1 Probability Model

The standard model for such data assumes that the number of failures has a binomial distribution. The assumptions are listed in Appendix A.6.1. The assumptions about the physical process can be boiled down to these two.

1. On each demand, the outcome is a failure with some probability *p*, and a success with probability 1 − *p*.

2. Occurrences of failures for different demands are statistically independent; that is, the probability of a failure on one demand is not affected by what happens on other demands.

The following kind of data can be observed.

- A random number of failures occur in some fixed, prespecified number of demands. As a minimum, the total number of failures and number of demands are observed.

Under these assumptions, the random number of failures, *X*, in some fixed number of demands, *n*, has a binomial(*n*, *p*) distribution.

$$\Pr(X = x) = \binom{n}{x} p^x (1 - p)^{n-x}, \quad x = 0, \ldots, n \tag{2.2}$$

where

$$\binom{n}{x} = \frac{n!}{x!(n-x)!} \; .$$

This distribution has two parameters, *n* and *p*, of which only the second is unknown. (Although *n* may not always be known exactly, it is treated as known in this handbook. Lack of perfect knowledge of *n*, and other uncertainties in the data, are discussed briefly in Section 6.1.2.)

2.3.2.2 Data Needed to Validate Model and Estimate *p*

Suppose that the binomial model holds. Then any reasonable estimator of *p* needs only two pieces of information: the number of demands, *n*, in the data period, and the number of failures, *x*, that occurred then.

However, more information is needed to investigate whether the binomial model is valid. For example, Assumption 1 assumes that *p* is the same on all demands. If the data cover a number of years or a number of plants, *p* might not be constant over time or the same at all plants. To study whether this is true, the times and locations of the demands and failures should be recorded, or at least the data should be partitioned into subsets, for example corresponding to plants or years. Then the failure and demand counts, x_i and n_i, should be given for each subset.

2.3.2.3 Case Studies: Validity of Model Assumptions in Selected Examples

Let us examine Examples 2.5 and 2.6 to see if the assumptions appear to be true.

Example 2.5 HPCI Failures to Start

Assumption 1 says that the probability of failure on demand is the same for every demand. If data are collected over a long time period, the assumption requires that the failure probability does not change. Likewise, if the data are collected from various plants, the assumption is that *p* is the same at all plants.

In the HPCI example, the five failures do not reveal any clear trend in time. However, one Licensee Event Report (LER) mentions that a better-designed switch had already been ordered before the HPCI failure. This gives some evidence of a gradual improvement in the HPCI system, which might be visible with more data.

As for differences between plants, it happens that 3 of the 5 failures occurred at a single plant. Therefore, it might be wise to analyze that one plant (3 failures in 9 demands) separately from the rest of the industry (2 failures in 50 demands). In fact, Grant et al. (1995) did not analyze the data that way, because they considered two types of failure to start, and they also considered additional data from full system tests performed once per operating cycle. However, the high failure probability for the one plant was recognized in the published analysis.

Assumption 2 says that the outcome of one demand does not influence the outcomes of later demands. Presumably, events at one plant have little effect on events at a different plant. However, the experience of one failure might cause a change in procedures or design that reduces the failure probability on later demands at the same plant. One of the five LERs mentions a permanent corrective action as a result of the HPCI failure, a change of piping to allow faster throttling. This shows some evidence of dependence of later outcomes on an earlier outcome at that plant.

Example 2.6 EDG Failures to Start

Assumption 1 says that every demand has the same probability, *p*, of failure. This is certainly not true for return-to-service tests, because such tests are guaranteed to result in success. If the EDG does not start on the test, maintenance is resumed and the test is regarded as a part of the maintenance, not as a return-to-service test. Therefore, any return-to-service tests should not be used with the rest of the data.

As for the other demands, one must decide whether the unplanned demands, operating cycle tests, and monthly tests are similar enough to have the same value of *p*. Can plant personnel warm up or otherwise prime the diesel before the test? Can an operator stop the test if the EDG is clearly having trouble, and then not consider the event as a test? If so, the different types of demands do not have the same *p*, and they should not be analyzed as one data set. For PRA purposes, one is interested in the failure probability on an actual unplanned demand. To estimate this, one should use only data from unplanned demands and from tests that closely mimic unplanned demands.

If the EDGs in the data set differ in some way, such as having different manufacturers, this may also lead to different values of *p* on different demands. Analyzing the data while ignoring differences between the individual EDGs will allow us to estimate the average *p*, corresponding to failure to start for a random EDG. However, this average *p* is not the same as the *p* for a particular EDG.

Assumption 2 says that the outcome on one demand does not affect the probability of failure on a different demand. When the plant is very new there may be some learning from individual failures, but when the plant is mature, failure or success on one demand should not change the chances of failure or success on later demands. The only way for such dependence to arise is if the first failure results from a common cause. If the plant is mature and common-cause failures are rare, then Assumption 2 is approximately satisfied.

Example 2.7 Steam binding in AFW

This is clearly a case of a latent failed condition that waits undetected until the demand. Therefore, it is better modeled as a standby failure, and is discussed in a later section.

Example 2.8 Failures of isolation valves

Here the causes listed are degradations, so that the probability of failure increases over time. If failures from such causes are rare, then the increase in failure probability may not be a problem. However, in that case there seems to be little reason for separating these failures from other failures of the isolation valves. This example is also discussed further in the section on standby failures.

2.3.2.4 Discussion

2.3.2.4.1 More General Models

The model considered here has a constant failure probability, *p*. A generalization would let *p* be a function of time. Such a model is useful for analyzing trends. Chapter 6 includes ways to test the assumptions of the model assumed here, and Chapter 7 includes ways to analyze data where a trend is present.

When data come from the industry, one might consider the differences between plants, just as for events in time. Ways to model such differences are discussed in Chapter 8. The present section's interest is restricted to *p* for the industry as a whole, the average of all the plants. Of course, if the data come from only one plant, *p* refers to that plant and the issue of differences typically does not arise.

As in all the sections of this chapter, any mathematical model is an imperfect approximation of the true process that generated the data. When the data set is sparse (few demands, or few or no failures, or few or no successes), (a) it is difficult or impossible to see evidence of departures from the model, and (b) the data set is too small to allow realistic estimation of the parameters of a more complicated model. When the data set has many events, departures from the model become visible, and a more complicated model may be appropriate.

2.3.2.4.2 Constancy of *n*

In the model considered here, the number of demands is treated as fixed, and the number of failures is treated as random. This is the normal type of data found in PRA work. It is possible to consider a model in which data are collected until x failures occur, and we count the number of demands, n, required to obtain x failures. In this case, the number of demands is random and follows a so-called negative binomial distribution. Such "waiting-time" models have little application in PRA, and are not considered in this handbook.

One could argue that the numbers of demands in the examples are not really fixed in advance. That is, no one decided in advance to look at the outcomes of 59 unplanned HPCI demands. Instead, Grant et al. decided to look at seven years of data from 23 plants, and they observed that 59 demands had taken place. The response to this argument is that we are actually conditioning on the number of demands, that is, dealing with conditional probabilities assuming that 59 demands take place. Conditioning on the number of demands enables us to focus on the quantity of interest, p. Treating both the number of failures and the number of demands as random is needlessly complicated, and yields essentially the same conclusions about p as do the simpler methods of this handbook.

2.3.3 Standby Failure

As stated in the introduction to this chapter, failure to change state can be modeled in two ways. One way was given in Section 2.3.2. The second way is given here, in which the system (typically a component) is assumed to transition to the failed state, at a constant transition rate, while the component is in standby. This latent failed condition ensures that the system will fail when it is next demanded, but the condition is not discovered until the next inspection, test, or actual demand.

2.3.3.1 Probability Model

The underlying assumption is that the transition to the failed condition occurs randomly in time. Two settings must be distinguished:

1. the *data*, the operational experiences in the past that allow us to estimate λ, and
2. the *application* to PRA, in which the estimate of λ is used to estimate the probability that a component will fail when demanded.

These two settings are discussed in the next two subsections.

2.3.3.1.1 Probability Model for the Data

It is customary to consider only the simplest model.

1. Assuming that the system is operable at time t, the probability that the system will become failed during a short time period from t to $t + \Delta t$ depends only on the length of the exposure period, Δt, not on the starting time of the period, t.

2. Failure of distinct systems, or of one system during distinct standby periods, are independent of each other.

The kind of observable data is spelled out here. It is obvious, but is written down here for later comparison with the data for similar models.

- At times unrelated to the state of the system, the condition of each system (failed or not) can be observed. As a minimum, the total number of failures and the corresponding total standby time are observed.

The times mentioned here can be scheduled tests or unplanned demands.

Assumption 1 is essentially the same as for a Poisson process in Section 2.2.2. It implies that there is a proportionality constant, λ, satisfying

$$\lambda\Delta t \approx \Pr(t < T \le t + \Delta t \mid T > t),$$

where T is the random time when the system becomes failed. Then the probability that the system is failed when observed at time t is

$$\Pr(\text{system is in failed state at time } t) = 1 - e^{-\lambda t} . \quad (2.3)$$

This follows from Equation 2.5, given in Section 2.6 for the exponential distribution. The parameter λ is called the **standby failure rate**.

2.3.3.1.2 Application of Model to PRA

The model is used to evaluate the probability of failure on demand by assuming that (a) the failure is revealed by a periodic test and the component is then returned to operation, and (b) the demand occurs at a random time within the testing cycle. With these assumptions, the probability of failure on demand is approximated by

$$p = \lambda t_{test}/2 \ , \quad (2.4)$$

where λ is the standby failure rate and t_{test} is the time interval between tests.

A more accurate expression is the average of terms from Equation 2.3, averaging over all the possible demand times in the test interval:

$$p = \frac{1}{t_{test}} \int_0^{t_{test}} \left(1 - e^{-\lambda s}\right) ds$$
$$= 1 - \left(1 - e^{-\lambda t_{test}}\right) / (\lambda t_{test})$$

2.3.3.2 Data Needed to Validate Model and Estimate λ

Suppose that the standby failure rate model holds. If the standby times are all similar, then an estimator of λ needs only two pieces of information: the number of failures, x, in the data period, and the corresponding total standby time, t. If, instead, the standby times vary substantially, then the total standby times should be recorded separately for the failures and the successes, as stated in Section 6.4.

To validate the model, the data could be partitioned. As with initiating events, if the data come from various years of plants, the data could be partitioned by year and/or by plant, and the above information should be given for each subset.

2.3.3.3 Case Studies: Validity of Model Assumptions in Examples

Let us now examine the applicability of the model assumptions in the examples given above.

Examples 2.5 and 2.6

Which is more appropriate, the failure-on-demand model or the standby-failure model? In these two examples, either model could be used.

Under the failure-on-demand model, the expected number of failures is proportional to the number of demands. Under the standby-failure model, the expected number of failures is proportional to the standby time. The statistician looks for data with variation in number of demands and standby time, to see which is more closely correlated with the observed failure counts. The engineer, on the other hand, looks at the apparent mechanisms of the observed failures. Do the mechanisms suggest that the number of failures is proportional to the number of demands or to the standby exposure time?

Examples 2.7 and 2.8

Assumption 1 says that the failed-condition event is as likely to hit the system in one time interval as in another of the same length. In Example 2.7, if the steam comes from backflow through a check valve, it will build up, and become more of a problem when the AFW system has been unattended longer. In this case, Assumption 1 is violated. Many of the causes mentioned for the valves in Example 2.8 also involve gradual buildup, not a sudden transition. For such causes, the random failed-condition shock is inappropriate as a model, although it has sometimes been used. A model of cumulative degradation would be more appropriate.

Assumption 2 says that the systems (AFW pumps or isolation valves) act independently. However, many of the causes in the examples can act as common-cause mechanisms. Steam-binding of the AFW system was a recognized common-cause mechanism in the 1970s and 1980s. This means that Assumption 2 may be plausible if interest is in the performance of a single AFW pump or a single isolation valve, but not if interest is in an interconnected set of such components.

Section D-1 of Poloski et al. (1998) says that steam binding has not been seen in 1987-1995 AFW experience. Therefore, Example 2.7 is probably no longer relevant, although it received great attention at one time.

2.3.3.4 Discussion

The discussion of Examples 2.7 and 2.8 has shown that some modeled standby-failure mechanisms involve gradual degradations, not instantaneous transitions. A wear-out model, as described in Section 2.6, would mimic the physics of failure more closely, and so be more appropriate. To estimate the rate of wear-out, data are needed for systems at various times. In the present context, this means knowing the state (failed or not) of AFW systems or isolation valves at *various times* after the last test or inspection. Such data are rarely available from operating plants, which have fixed testing schedules. This lack of data undoubtedly contributes to the use of the simple, but inappropriate, standby-failure model.

Fortunately, degradation mechanisms have become minor contributors to risk. When a degradation mechanism is recognized as important, the natural response is not to collect data to better estimate the rate of degradation. Instead, the natural response is to shorten the interval between preventive maintenance activities, and so to identify and correct incipient degradation, or to modify the plant to mitigate or eliminate the problem. Examples are the apparent elimination of steam-binding in AFW pumps, mentioned above, and of intergranular stress corrosion cracking (IGSCC) in BWR piping (Poloski et al. 1999a, Appendix J).

2.3.4 Comparison of the Two Models for Failure to Change State

One great appeal of the standby-failure model is that the parameter estimator does not need knowledge of the number of demands. Standby time is normally much easier to obtain than a count of demands. To validate the use of the standby-failure model, reason as in the above discussion of Examples 2.5 and 2.6, trying to decide if the number of failures is proportional to the standby time or to the demand count.

The choice of model makes a difference! For example, suppose that an EDG is tested monthly by starting it. In 100 monthly tests, 2 failures have been seen. A simple estimate of p, the probability of failure on demand, is 2/100 = 0.02. A simple estimate of λ, the standby failure rate, is 0.02/month. Now suppose that a basic event in a PRA is that the EDG fails to start, when demanded at a random time. Based on the estimate of p, the estimated probability of the basic event is 0.02. Based on the estimate of λ and Equation 2.4, the estimated probability of the basic event is

$$(0.02/\text{month})\times(1\text{ month})/2 = 0.01\ .$$

The two models give estimates that differ by a factor of 2! The reason is simple. The failure-on-demand model says that all demands have the same probability of failure. The standby-failure model says that demands soon after a successful test have smaller probability of failure.

2.4 Failure to Run during Mission

Aspects of this type of failure closely resemble the initiating events of Section 2.2. One important difference is in the kind of data normally present. The difference is summarized here.

Example 2.4 of Section 2.2 is an example of continuously running components, temperature sensor/transmitters, that occasionally failed to run. When a component failed, it was repaired or replaced in a relatively short time, and resumed operation. That is, the component was **repairable**. The present section considers components or systems that do not run continuously. Instead, they are occasionally demanded to start, and then to run for some mission time. If they fail during the mission, they are **nonrepairable**, that is, they cannot be repaired or replaced quickly. Three points deserve clarification:

- Some failures may be recoverable. They would not be modeled as failures in the sense of causing mission failure. Unrecoverable failures cause mission failure, however.

- Given enough time, almost any system can be repaired. During a mission, however, time is not available. Because the component or system cannot be repaired *within the time constraints*, it is called "nonrepairable."
- From the viewpoint of data analysis, the terms "repairable" and "nonrepairable" refer to how data are reported. If data collection stops after the first failure, at least for that particular application of that component or system, the system is called nonrepairable. If, instead, a sequence of failure times is recorded, the system is called repairable.

As stated earlier, the word **system** is used generally for any piece of hardware for which data are taken. In particular, components and trains are kinds of systems.

2.4.1 Examples

Here are two examples of failures to run during missions.

Example 2.9 EDG failures to run

Grant et al. (1996) report that in 844 demands of 30 minutes or more for emergency diesel generators (EDGs) to run, there were approximately 11 unrecovered failures to run in the first 30 minutes. The count is approximate because a few failure times were not given and had to be inferred.

Example 2.10 AFW turbine train failures to run

Poloski et al. (1998) report that in 583 unplanned demands of auxiliary feedwater (AFW) system turbine trains, the train failed to run 2 times, and the total running time was 371 train-hours. The information is taken from LERs, only 17% of which report running times for the train. The total running time of 371 hours is an extrapolation from the LERs with reported run times.

These examples are typical, in that hardly any of the demands to run resulted in a failure. Therefore, for most demands the time when failure would eventually have occurred is unknown.

2.4.2 Probability Model

In principle, the times to failure are **durations**. Section 2.6 deals with duration data, in the context of recovery times. That section mentions various possible distributions of time to failure, of which the simplest is the exponential distribution. In the present setting, however, nearly all of the observed times are truncated before failure. This is illustrated by the above examples. Therefore, the full distribution of the time to failure cannot be observed. In Example 2.9, no information is given about the distribution of failures times after the first 30 minutes. In Example 2.10, the average run time was only 38 minutes, and most AFW missions lasted for less than one hour. In such cases the exponential distribution, restricted to the observed time period, is a simple, reasonable approximation of the observable portion of the distribution.

Two assumptions are made concerning the physical process.

1. Assuming that no failure has occurred by time t, the probability that the failure will occur in a short time period t to $t + \Delta t$ depends only on the length of the exposure period, Δt, not on the starting time of the period, t.

2. Failures of distinct systems, or of one system during distinct missions, are independent of each other.

The kind of observable data is as follows.

- For each observed mission, the run time is observable, and whether the run terminated in failure or in successful completion of the mission. As a minimum, the total run time and the number of failures to run are observed.

Assumption 1 implies that the time to failure is exponentially distributed with parameter λ. The interpretation of λ is that if the system is running, the probability of failure in the next short interval of length Δt is approximately $\lambda\Delta t$. That is

$$\lambda\Delta t \approx \Pr(t < T \le t + \Delta t \mid T > t),$$

where T is the random time until failure. When defined this way, λ is sometimes called the **failure rate**, or **rate of failure to run**. Many authors use the term **hazard rate**, denoted by *h*, and discussed in Appendix A.4.4.

Note, the definition of λ is different for repairable systems (Section 2.2) and nonrepairable systems (the present section), even though it is represented by the same Greek letter and is called "failure rate" in both cases. See Thompson (1981) for a reasonably clear discussion of the subtle differences, and the glossary of this handbook for a summary of the definitions. The topic is discussed further in Appendix A.4.4.

It is instructive to compare the models for failure to run and standby failure. The physical process is essentially identical, but the observable data differs in the two models.

2.4.3 Data Needed to Validate Model and Estimate λ

Suppose that the time to failure has an exponential distribution. Then any reasonable estimator of λ needs only two pieces of information: the total running time, *t*, in the data period, and the number of failures to run, *x*, that occurred then.

However, more information is needed to investigate whether the exponential distribution is valid. Assumption 1 says that λ is constant during the mission. To investigate this, the analyst should know the failure times, that is, how long the failed pumps ran before failing. The analyst should also know the mission times, that is, how long the system ran when it did not fail; often, however, this information is not recorded and can only be estimated or approximated.

Implicit in Assumption 1 is that λ is the same over all the years of data, at all the plants where the data were collected. To investigate this, the data should be divided into subsets, corresponding to the different plants and years. Then the failure count and running time, x_i and t_i, should be given for each subset. This is the exact analogue of what was said in Section 2.2.3 for initiating events.

2.4.4 Case Studies: Validity of Model Assumptions in Examples

Consider now whether the assumption of the model is plausible for the two examples.

Example 2.9 EDG Failures to Run

Assumption 1 says that a running EDG is as likely to fail in one short time interval as in any other time interval of the same length. That is, the EDG does not experience burn-in or wear-out failures. The reference report (Grant et al. 1996) says that this is not true over a 24-hr mission. Indeed, that report divides the EDG mission into three time periods (first half hour, from ½ hour to 14 hours, and from 14 to 24 hours) to account for different failure rates during different time periods. Within the first half hour, however, the data do not give reason for believing that any short time interval is more likely to have a failure than any other time interval. Therefore, Assumption 1 can be accepted.

Assumption 2 is violated by common-cause failures. It is also violated if a failure's root cause is incorrectly diagnosed, and persists on the next demand. If these two conditions are rare the assumption may be an adequate approximation. More subtle dependencies are difficult to detect from data.

Example 2.10 AFW Turbine Train Failures to Run

Assumption 1 says that a running turbine train is as likely to fail in one short time interval as in any other time interval of the same length. The data are too sparse — only 2 observed failures — to confirm or refute this assumption.

2.4.5 Discussion

The exponential time to failure can also be derived as the time to *first* failure in a Poisson process of Section 2.2. The present context is simpler, however, because the process ends after the first event, failure to run. The Poisson-process assumptions about hypothetical additional failures are irrelevant.

2.5 Unavailability

Unavailability is the fraction of time when a system is not able to perform its function. Sometimes this is qualified. For example, test-and-maintenance unavailability is the unavailability caused by testing and maintenance of the system.

2.5.1 Example

The example here is typical.

Example 2.11 HPCI unavailability for test and maintenance

> The HPCI system is typically available when demanded unexpectedly, but sometimes it is out of service for testing or maintenance, scheduled or unscheduled.

The system has two states, up (in service) and down (out of service), and a long-term fraction of time when the system is down. The fraction of time when a standby system is up is called the system **availability**.

2.5.2 Probability Model

The simplest data consist of measurements in a number of time periods, called **reporting periods** here. Each reporting period corresponds to an **exposure time**, or **time at risk**, the time when the system should be available. The system's **down time**, or **outage time**, is the time when the system was out of service. One set of mathematical assumptions for unavailability is given here.

1. The down times during different reporting periods are statistically independent of each other.
2. The system has the same probability of being down at a random time in one reporting period as in any other reporting period.

Two kinds of data can be considered.

- For each reporting period, the outage time and exposure time are observed.
- Alternatively, the number of individual outages and the duration each outage are observed.

With data described under the second bullet, the number of outages can be considered a Poisson count (Section 2.2), and the durations can be modeled as in Section 2.6. Therefore, such data will be considered only briefly, in Section 2.5.4. Data described under the first bullet will be used in most of the treatment of this handbook.

In Example 2.11 unexpected down times occur randomly, but planned down times also occur at regular intervals, scheduled maintenance for the system. This must be recognized when the reporting periods are defined. For example, if a motor-driven pump has most of its scheduled maintenance during the plant's refueling outages, and the pump's availability during shutdown is of interest, then separate analyses should probably be performed for the time when the reactor is up and when the reactor is down, to keep Assumption 2 from being violated.

If a system is maintained periodically, one must define long enough reporting periods so that each reporting period includes the same number of regularly scheduled lengthy down times.

A mathematical way to model the status of a repairable system uses a **state variable**, defined as $S(t) = 1$ if the system is up at time t, and $S(t) = 0$ if it is down at time t. The **availability** of a system is the probability that the system is up at a random time T,

$$A = \Pr[S(T) = 1],$$

and the **unavailability** is

$$\overline{A} = \Pr[S(T) = 0] = 1 - A\,.$$

Reliability textbooks consider availability as a function of time. In typical PRA applications, however, the relevant availability and unavailability are at the random time of an unplanned demand for a system, and therefore are the constants defined above. From the perspective of this data analysis handbook, availability and unavailability are parameters to be estimated. In keeping with the notation of lower-case letters for parameters, this handbook often denotes the unavailability by u instead of by $\overline{A}$.

A particular system history is illustrated in Figure 2.1, from Engelhardt (1996). This figure shows when a particular system was operating ($S = 1$) or shutdown ($S = 0$). A nominally identical system would have a somewhat different history for the same period, or the same system would have a different history over a different time period of the same length.

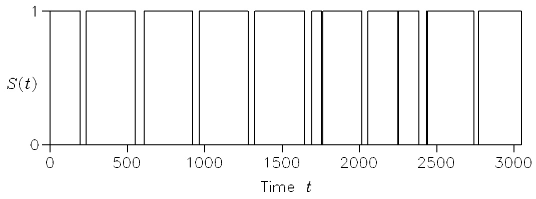

Figure 2.2 Uptime and downtime status for one system.

Actually, this figure refers to uptime for a nuclear power plant, not the kind of "system" normally analyzed by a PRA. However, the figure is shown only as an illustration of the up-and-down behavior that must be analyzed when estimating unavailability.

2.5.3 Data Needed to Validate Model and Estimate *u*

The unavailability, *u*, can be estimated from the total exposure time and corresponding outage time in a single reporting period.

However, every data reporting period will give a somewhat different estimate of *u*. To quantify the uncertainty in the estimate, data from a number of reporting periods are needed, for example, data from separate time periods or from separate, nominally identical hardware systems. Moreover, the methods given in Chapter 6 aggregate, that is, combine, reporting periods, at the very least so that the aggregated reporting periods do not contain outage times of zero.

In summary, a large enough data set is needed so that it consists of more than two reporting periods (as a bare minimum), and each reporting period is long enough so that outage times of zero do not occur.

Assumption 1 is normally validated by careful thinking, not by data analysis.

To validate Assumption 2 of the model, the data should be partitioned into subsets, corresponding, for example, to different years or safety systems or plants. Then each subset must be large enough to give data as described in the previous paragraph.

2.5.4 Discussion

Suppose now that the observed data are described by the second bullet in Section 2.5.2, so that the individual durations of up times and down times are reported. Then the **mean time to failure** (**MTTF**) can be estimated from the durations of the up times, and the **mean time to repair** (**MTTR**) can be estimated from the durations of the down times. A strong result can be shown under the following assumptions.

1. The times to failure are identically distributed and the times to repair are identically distributed.
2. All the durations are statistically independent of each other.

Then (see Ross 1983, pp. 66-67) it can be shown that the unavailability equals

$$\overline{A} = \frac{MTTR}{MTTF + MTTR} \ .$$

This provides a basis for estimating availability and unavailability when the full duration data are available.

Upon reflection, this striking result is not so surprising. If the data consist of *n* up times with *n* down times interspersed, MTTR would be estimated by (total down time)/*n* and MTTF would be estimated by (total up time)/*n*. Then the natural estimate of unavailability would be

[down-time/*n*]/[up-time/*n* + down-time/*n*]
= down-time/(up-time + down-time) ,

which is just the observed fraction of time when the system is down.

2.6 Recovery Times and Other Random Duration Times

This section is about modeling of time data. Often, a measurement of interest is a random duration time, such as the time required to return a failed system to service or the lifetime of a piece of hardware. The distinction between random duration times here and events in time in Sections 2.2 and 2.4 is that here the individual times are measured on a continuous scale with units such as minutes or hours, while the earlier data sets involve discrete counts of the number of events occurring in a total length of time.

2.6.1 Examples

Here are some examples involving random duration times. They are only summarized here. Actual examples, with lists of durations times, will be analyzed in Chapter 6.

Example 2.12 Recovery times from loss of offsite power

A plant occasionally loses offsite power. When this happens, the plant reports the time until power is restored. Atwood et al. (1998) present such durations for LOSP events in 1980-1996.

Example 2.13 Repair times for turbine-driven pumps

A turbine-driven pump must occasionally be taken out of service for unplanned maintenance. The duration of time out of service for maintenance may be extractable from maintenance records.

Example 2.14 Time to failure of a component

A typical power plant will have many individual components such as compressors. When a component is put into service, it operates intermittently until it fails to perform its required function for some reason. Høyland and Rausand (1994) give an example of such data.

Example 2.15 Times to suppress fires

When a fire occurs in a nuclear power plant, the time until the fire is suppressed is of interest. Nowlen et al. (2002) report on analysis of such suppression times. One difficulty is that the time of fire onset often is not known exactly.

Example 2.16 Gradual degradation until failure

Examples 2.7 (steam binding) and 2.8 (failure of isolation valves) involve gradual degradation, which builds up until the system is inoperable. The time until the system is inoperable can be modeled as a duration time.

The common element in these examples is a duration time that varies in an unpredictable way. In Examples 2.12 and 2.13, the recovery time is composed of several factors such as time to diagnose, perform and test repairs, and the time to complete documentation required before returning the plant to normal operating conditions. Example 2.14 is a failure-to-run example, similar to those of Section 2.4. This example differs from that of Section 2.4, however, because here it is assumed that virtually all of the times to failure are recorded. In Section 2.4, on the other hand, most of the systems did not fail during the test period or operational mission. The severe truncation of the data in Section 2.4 meant that only a simple model could be considered. The more complete data here allows analysis of a more complex model. Example 2.15 is complicated by the lack of exact knowledge of the duration time. Finally, Example 2.16 gives a realistic conceptual way to model the gradual degradations encountered in Section 2.3.1, although good data are unobtainable.

All five examples involve a duration time that is uncertain due to random factors. Consequently the duration times are modeled as continuous random variables.

2.6.2 Duration-Time Models

The duration, *T*, is random, following some probability distribution. Two assumptions are made about the process.

1. Each duration is statistically independent of the others.
2. All the random durations come from the same

probability distribution.

The data description is simple:

3. The individual durations are observable. As a bare minimum, the number of durations and the total duration time are observed.

Assumptions 1 and 2 can be summarized by saying that the durations are **independent** and **identically distributed**. Independence means that one duration does not influence the probability of any other duration. The assumption of identical distributions means that each random duration is as likely as any other to be long or short. If the durations are from distinct systems, the systems are assumed to be identical and to act independently. If the durations are in sequence, as for a system that alternates being up and down, the assumption implies that no learning or long-term aging takes place, and that each repair restores the system to a condition as good as new. Such a process is called a **renewal process**.

The assumptions do not require a particular distribution for the time between events. The most important such distributions in PRA applications are:

- lognormal
- exponential
- Weibull
- gamma

These distributions are all summarized in Appendix A.7. An important part of the data analysis consists of deciding on the form (or several plausible forms) of the distribution. This will be discussed in Chapter 6. For now, we simply note that these and other distributions are possible.

There are different ways to specify a probability distribution, and the next material summarizes some of the concepts: their definitions, how to interpret them, and how they are related to each other. The data-analysis techniques of Chapter 6 will use these ways of characterizing distributions. The usual convention is to denote the random variables using capital letters, *T*, and observed times as lower case, *t*. The letter *T* is used, rather than some other letter such as *X*, because the random quantities are times. As seen from the examples, the durations may be times to repair, times to failure, or other times. However, the concepts and formulas are valid for any application.

The **cumulative distribution function** (c.d.f.) of a real-valued random variable *T* is defined as

$$F(t) = \Pr(T \le t)$$

for all real numbers *t*. The name is sometimes abbreviated to **distribution function**. The c.d.f. is the probability that the random variable *T* will assume a value which is either less than or equal to *t*. The c.d.f. is a monotonic increasing function of *t*, with the limiting properties $F(0) = 0$ and $F(+\infty) = 1$. [For random variables that, unlike durations, can take negative values, the limiting properties are $F(-\infty) = 0$ and $F(+\infty) = 1$. That general case has few applications in this handbook.]

The distribution is commonly used to characterize the lifetimes, or recovery times, or some other kind of durations, of a whole population of systems. The population might be a large set of identical systems that are operating in similar applications and with durations that vary due to random influences. $F(t)$ is the fraction of items that have durations *t* or less, in a hypothetical infinite population.

A related function, denoted by $f(t)$, is called a **probability density function** (p.d.f.) for a continuously distributed positive-valued random variable *T*. It is related to the c.d.f. by

$$f(t) = \frac{d}{dt}F(t) \quad \text{and}$$

$$F(t) = \int_0^t f(u)du \quad .$$

The variable *u* is a dummy variable of integration, and *t* is the upper limit of the integral. An example of a p.d.f. and the associated c.d.f. are shown in Figure 2.3.

It follows that probabilities corresponding to occurrences in a small interval of time are approximately proportional to the p.d.f.,

$$\Pr(t < T \le t + \Delta t) \approx f(t)\Delta t.$$

Therefore, the ordinate of a p.d.f. has units of "probability density" and not probability (as for a c.d.f.). Thus, a p.d.f. determines how to assign probability over small intervals of time. Now consider an arbitrary interval from *a* to *b*. In this case we have

$$\Pr(a < T \le b) = \int_a^b f(t)dt \ .$$

This simplest distribution is the **exponential** distribution. It arises when the assumption of Section 2.4.2 is satisfied. (That assumption is phrased as if *T* is a time until failure.) In that case, the probability distribution is exponential, and determined by a single parameter, λ. The p.d.f. and c.d.f. are given by

$$f(t) = \lambda e^{-\lambda t}$$

$$F(t) = 1 - e^{-\lambda t} \ . \qquad (2.5)$$

2.6.3 Data Needed to Estimate Distribution of Durations and Validate Model

In general, a sample of observed durations is needed to estimate the distribution of duration times. These durations must independent and identically distributed, that is, they must be generated by a process satisfying the two assumptions given at the beginning of Section 2.6.2.

The special case when the times are assumed to have an exponential(λ) distribution is simpler. Only the number of durations and the total duration time are needed to estimate λ,. However, the individual durations are still needed to investigate whether the distribution is exponential or some other form. Incidentally, when the distribution is assumed to be exponential, the model given here differs from the standby-failure model (Section 2.3.3.1) and from the failure-to-run model (Section 2.4.2) *only* by the kind of data that can be observed.

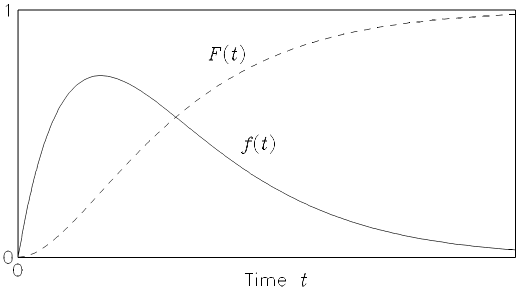

Figure 2.3 Probability density function (p.d.f.) and cumulative distribution function (c.d.f.).

To validate whether the distribution is the same for all the data, extra information should be recorded for each duration, the relevant circumstances of each duration. The circumstances of interest are those that might affect the durations, such as time of the event, system location, and system condition just before the event.

2.6.4 Validity of Model Assumptions in the Examples

Examples 2.12 through 2.14 all appear to satisfy the assumptions of Section 2.6.2. Example 2.15 also does, except that the durations are not observed exactly.

In each case, all the distributions come from some distribution. Discovering the form of that distribution is a task for the data analyst.

One might ask whether the durations are statistically independent. For example, does a long repair time for a turbine-driven pump add an extra benefit to the pump, so that the next few repair times will be short?

One might also ask, for each example, whether the durations all come from the same probability distribution. For example, if the data cover a period of years, has there been any long-term learning, so that recovery times, repair times, or refueling times tend to be shorter than at the start of the data period? Are different durations associated with different systems for the turbine-driven pumps, with different causes of loss of offsite power, or with different kinds of fires?

The above are questions that could be investigated during the data analysis, if enough durations have been observed.

Example 2.15 is complicated by lack of exact measurements of the durations. Bounds can be given, and the analysis must be based on these upper and lower bounds rather than on exact times.

Example 2.16 is different because the durations are not observable at all. It might be theoretically interesting to model the time until the system is in a failed condition as a duration. But there is no monitor on the pump or valve that says, “At this time the system just became inoperable.” Therefore, the durations are not directly observable, not even in principle. Therefore the methods of this handbook are not applicable to this example.