

The MTBF Sucks. Here's Why.



Workbook

Presented by

Dr Chris Jackson

Reliability engineer



Copyright © 2025 Christopher Jackson

All rights reserved. No part of this publication may be reproduced, distributed or transmitted in any form or by any means, including photocopying or other electronic or mechanical methods without the prior written permission of the author (Christopher Jackson), except as permitted under Section 107 or 108 of the 1976 United States Copyright Act and certain other non-commercial uses permitted by copyright law.

For permission requests, write to Christopher Jackson using the details on the following pages.

Limit of Liability/Disclaimer of Warranty: While the author has used his best efforts in preparing this document, he makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. The author shall not be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

Who is your teacher? ... Dr Chris Jackson

Dr Jackson holds a Ph.D. and MSc in Reliability Engineering from the University of Maryland's Center of Risk and Reliability. He also holds a BE in Mechanical Engineering, a Masters of Military Science from the Australian National University (ANU) and has substantial experience in the field of leadership, management, risk, reliability and maintainability. He is a Certified Reliability Engineer (CRE) through the American Society of Quality (ASQ) and a Chartered Professional Engineer (CPEng) through Engineers Australia.

Dr Jackson is the cofounder of www.is4.com online learning, was the inaugural director of the University of California, Los Angeles (UCLA's) Center of Reliability and Resilience Engineering and the founder of its Center for the Safety and Reliability of Autonomous Systems (SARAS). He has had an extensive career as a Senior Reliability Engineer in the Australian Defence Force, and is the Director of Acuitas Reliability Pty Ltd.

He has supported many complex and material systems develop their reliability performance and assisted in providing reliability management frameworks that support business outcomes (that is, make more money). He has been a reliability management consultant to many companies across industries ranging from medical devices to small satellites. He has also led initiatives in complementary fields such as systems engineering and performance-based management frameworks (PBMFs). He is the author of two reliability engineering books, co-author of another, and has authored several journal articles and conference papers.





Table of Contents

what is RANDOM?	5
what is the MEAN?	8
how can I use the area under a PROBABILITY DENSITY CURVE?	11
so what is a RELIABILITY CURVE?	12
so how do RELIABILITY and PROBABILITY DENSITY CURVES ... relate?	14
so what is RELIABILITY?	15
... and the MEAN (MTTF) is NOT RELIABILITY	15


Failure is a **random process**. And this is something that our primitive human brains really struggle with.

what is RANDOM?

Something that is **random** (at least to us) is a process that gives different outputs or results for (seemingly) identical inputs. These inputs could be using lots of (seemingly) identical **ball bearings**. A **ball bearing** supports a rotating shaft with minimal friction. Many machines (cars, smart locks, kitchen appliances and lots more) rely on rotating shafts, meaning they also rely on **ball bearings**. A ball bearing is made up of an **inner race** (the ring that attaches to the shaft), an **outer race** (that is held in place) with several **steel balls** rolling between the two and minimizing friction. But because **failure** is a **random process** ...




But although these **ball bearings** are indistinguishable, they are not exactly the same.



A **dislocation** involves neighboring rows of metal atoms that don't line up perfectly. Both **impurities** and **dislocations** are invisible microscopic **defects** that create lots of tiny structural weaknesses that differ between **ball bearings** that look identical to us. So ball bearings will fail in slightly different ways at perhaps very different times.


There are lots of factors that drive how each **ball bearing fails** that we will represent below with our **random hand of failure**.



Experience has shown that manufacturing **defects** (like **impurities** and **dislocations**) play a small role in how long each **ball bearing** lasts. How we treat our **ball bearings** (lubrication, maintenance, correct alignment and so on) is much more important.


We usually use the phrase **Time to Failure (TTF)** to describe how long it takes for something to **fail**. **TTF** is not always calendar time and is better seen as a measure of the age of a product. The **TTF** of ball bearings is measured in terms of **cycles** or **revolutions**. This means that a two-year-old **ball bearing** that has been barely used might be effectively younger than a week-old **ball bearing** that has been heavily used.

Let's represent the **TTFs** of 100 ball bearings we test with a blue ball ...




But this is not reality. In practice, ball bearings will **fail** at different times. But just because something is **random**, doesn't mean its unpredictable.

For example, we can see patterns in 'real-world' blue dots ...




We often try to look for patterns in **random** outcomes. But it is difficult for us to 'see' patterns in simple representations of data points on an arrow. We need another way of looking at data ...



But the **histogram** above is not particularly smooth, even with 100 data points or blue dots (which is a LOT of data in the world of **reliability engineering**). One way of getting a 'smoother' **histogram** that better represents the underlying random process is to get more data. So if we test 10 000 **ball bearings** ...

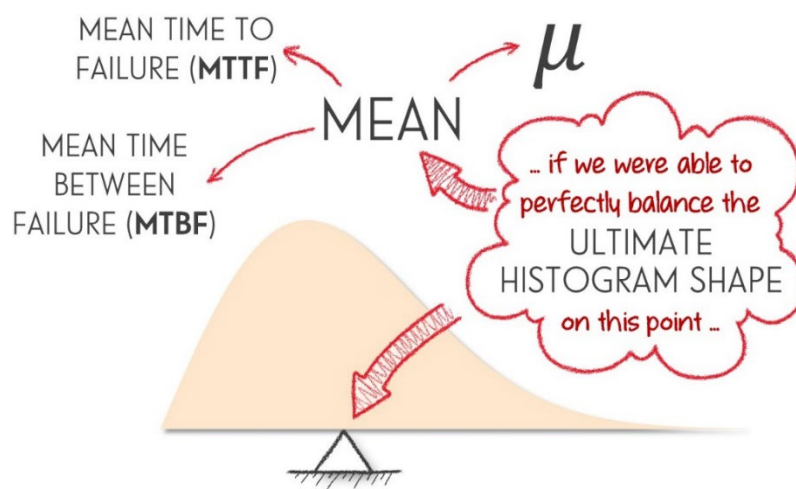


If we test an INCREDIBLE number of **ball bearings**, we would get a **histogram** with the bars are so small that they create a smooth shape (perhaps called the **ultimate histogram**). The curve on top of this shape has a special name ...




what is the MEAN?

The **mean** is a term that is used in the field of probability, statistics, and **reliability engineering** (often incorrectly). But what is it actually?




The symbol ' μ ' is a Greek lowercase letter that is pronounce 'mu.' Because the **mean** is just the balance point of the **ultimate histogram** shape, it has no physical meaning. But the **mean**, **Mean Time to Failure (MTTF)**, or **Mean Time between Failure (MTBF)** if our product is repairable is often used to 'represent' **reliability**. The **mean** is NOT **reliability**.

Let's look at the **failure** of a small satellite, which is often no bigger than a shoebox. Small satellites are not mass produced and are often riddled with **defects** that mass produced products don't have. So the **probability density curve** and its **ultimate histogram** shape that describes the **TTF** of small satellites has a particular shape (below) which is not what we 'think' **random** things look like.




Products with lots of **defects** that **wear out** have **ultimate histogram** shapes that hug the vertical axis as shown above. And these **defects** include the seemingly ever-present manufacturing problems, **assembly** or **installation error**, **interface issues** (between components), **damage** during transit or launch, and **software bugs**. This weird shape has an 'interesting' **mean** or balance point.



So even though the weird **ultimate histogram** shape above is perfectly balanced on the **mean** (**MTTF**), 80 % of its area is on the left. This is because the remaining 20 % on the right extends so far out that it creates more 'leverage' to create this perfect balance even though this only accounts for a small fraction of this area.


But this is not how it works for every product. The **TTF** for a rubber tire behaves very differently ...



This is how our human brains prefer visualizing **randomness**. We expect half of all tires to **fail** by the **mean (MTTF)**. And just to confirm, the balance points for the small satellite (on the previous page) and the rubber tire **ultimate histogram** shapes are in exactly the same place.


The reason the **ultimate histogram** shape above looks like this is that rubber tires **wear out**, which means they **accumulate damage**. Lots of different things **accumulate damage** through **mechanisms** like **diffusion** (where atoms move across the silicon crystal of an electronic component), **creep** (where solids slowly deform over time under load), **fatigue** (where cracks grow slowly over time after being repeatedly stressed), **wear** (involving solid matter being abraded away when it comes into contact with something else), **lubricant contamination**, **oxidation** (where substances react with oxygen to turn into something else) and **corrosion** (oxidation that only applies to metals).

If the **ultimate histogram** shape of a components **TTF** has a peak (as shown below), it tells us that the component is **accumulating damage** and **wearing out**.




how can I use the area under a PROBABILITY DENSITY CURVE?

The **probability density curve** summarizes everything we know about a **random process** (like **failure**). We often want to use the **probability density curve** to us make better **reliability engineering** decisions. We do this using the area, which allows us to (for example) work out when we expect the first 10 % of our **ball bearings** to **fail**.



We are almost always only interested in understanding when the first 'small fraction' of products or components will **fail**. We typically want to replace most things before they **fail**, as **failure** is much more expensive (and potentially damaging) then a scheduled replacement.


We often use what we call the ' L_{10} ' **life**, which is the time by which we expect 10 % of our products to **fail**. The bearing industry uses the ' B_{10} ' **life** where ' B ' represents **ball bearing life**.




Most (90 %) of the area under the **probability density curve** is to the right of the ' L_{10} ' or ' B_{10} ' **life**. This represents the fraction of **ball bearings** that are still working at the ' L_{10} ' or ' B_{10} ' **life**.

so what is a RELIABILITY CURVE?

The **probability density curve** summarizes the **random** nature of **failure**. But trying to measure areas under curves is difficult. We can represent the same **random process** with a **reliability curve**, like the one below for the **ball bearing** with the ' B_{10} ' life occurring when the **reliability curve** drops to 90 %.



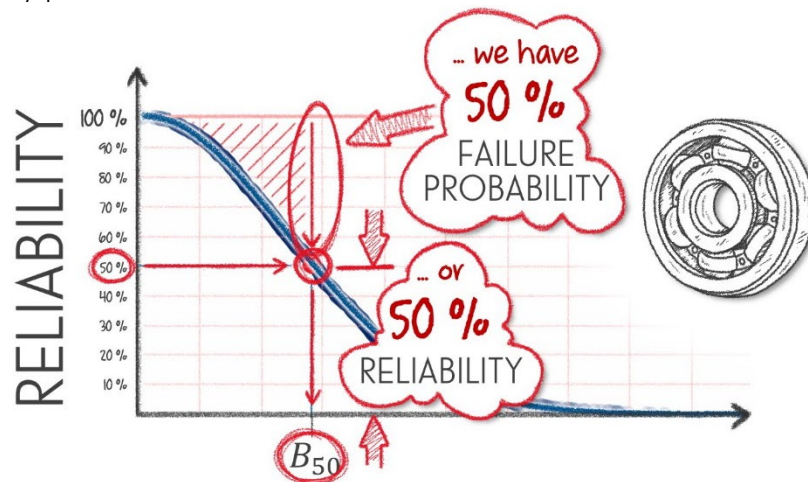
Not all **reliability curves** start at 100 %.




This is a common problem for products, machines and systems that deploy in arduous conditions. And 'arduous' conditions can include product mishandling by delivery truck drivers. So if 15 % of products (like small satellites) suffer **DOA failures**, then their **reliability curves** need to start at 85 % instead of 100 %.

But most **reliability curves** start at 100 % when **TTF** is equal to zero. This represents the premise (or assumption?) that everything works when we first use it. If this is not the case, we need to adjust our **reliability curve**.


The **reliability curve** can then tell us the **probability** that our product (like a **ball bearing**) is still working at any point in time.



We call the time when a product has **50 % reliability** (or when we expect half our products to have **failed**) the ' L_{50} ' life or ' B_{50} ' life if the component is a **ball bearing**. And all **reliability curves** 'head' in the same direction as products get older.




But we are typically ONLY interested in the first part of the **reliability curve** ...




so how do RELIABILITY and PROBABILITY DENSITY CURVES ... relate?

Every **probability density curve** has a specific **reliability curve**, and vice versa.



For example ...



The **probability density curve** aligns with how our brains like visualize **randomness**. The height of the **probability density curve** at a particular **TTF** represents how likely it is for our product to **fail** at that **TTF**.

The **reliability curve** makes it easier to work out how likely it is for something to still be working at specific points in time but are not how our brains like visualizing **randomness**.

so what is RELIABILITY?

RELIABILITY

... is the probability that an item will not fail for a designated period under specified conditions.

There are four key parts to this definition.

We need to define **failure**. We often think this is obvious, but it is not. There are lots of different ways things can **fail**, and some are much worse than others. And **failure** is all about what customers believe **failure** to be.

We need to then specify **conditions**. Many engineers think **specified conditions** are a 'get out of jail free card' in that we can specify 'perfect' conditions for the 'perfect' user or customer, using our product in the 'perfect' way, without having to worry about accidental misuse, or use in different ways. This thinking is not helpful and has been the root cause of many bankruptcies. **Specified conditions** need to be a genuine attempt to describe how customers and users will really use our product or component.

We now need to define a **designated period**. And this simply means that **reliability** has a curve (which we just looked at). The curve represents how **reliability** changes over time, so we can't just talk about **reliability** without specifying a time (or designate period).

And finally, we need to define **probability**. **Reliability** is always a **probability**. It won't tell us exactly when things will **fail**. But it does allow us to have a pretty good idea of how many will have **failed** by that **designated period**.

... and the MEAN (MTTF) is NOT RELIABILITY

Many people think **reliability** is the **MTTF**. It is not. The **probability density curves** for small satellites and rubber tires are very different even if they have the same **MTTF**. So this means that their **reliabilities** will be very different at different times, again, even when they have the same **MTTF**.

The MTTF is not a failure free period.

It is not even a largely failure free period.

The **MTTF** or **mean** is the time by which 50 – 80 % of your products have **failed**, depending on how they **fail**. And it is all because the **mean** is the balance point of the area underneath a **probability density curve**, which has no physical meaning.