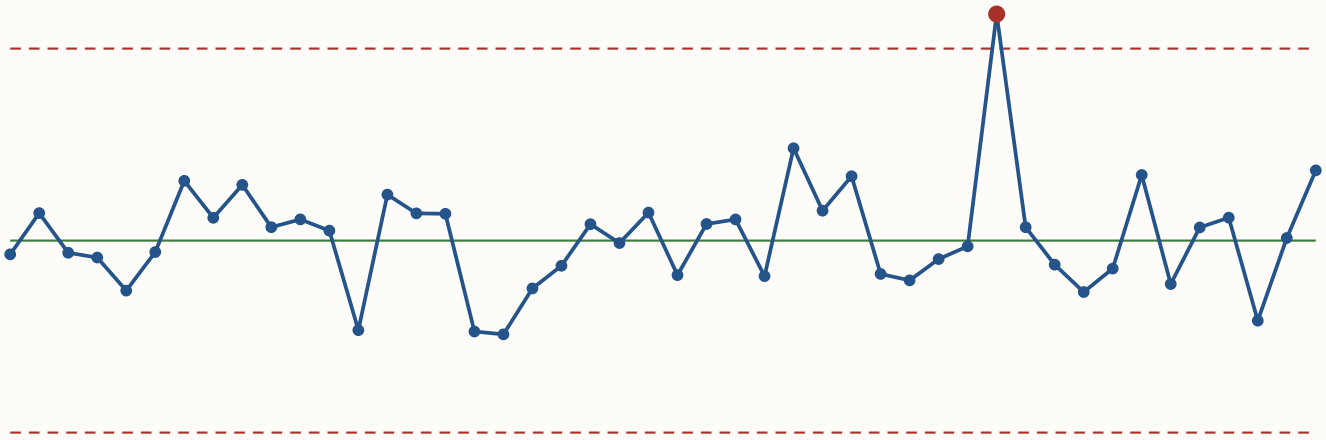
Reliability Solutions

RS-1

User Guide



Statistical process control, capability, measurement systems,
reliability and automated quality reports - step by step,
with every chart explained and tied to real shop-floor decisions.

For RS-1 version 1.0.0 · November 2026
Every picture in this guide was taken from the program itself, running the steps described.

Contents

1. Welcome to RS-1 9

What the program is for 9

Who this guide is for 9

How the guide is laid out 9

Every picture is real 10

2. Installing and signing in 11

Installing 11

Activating your license 11

The first account 12

Signing in every day 13

If you forget your password 14

3. A tour of the window 16

The start page 16

The window with data open 17

The ribbon tabs 18

Reading a result 20

The columns panel and column details 22

4. Getting your data in 23

Ways to bring data in 23

The shape that works best 23

Spec limits: tell the program once 24

Check the data before you trust it 24

5. Your first study in fifteen minutes 26

Step 1: open the data 26

Step 2: set the spec limits 26

Step 3: is the process stable? 26

Step 4: is it capable? 28

Step 5: the report 29

6. Variation: the idea behind every chart . . . 32

Two kinds of variation . . . 32

Seeing the difference . . . 32

The shape of the scatter . . . 34

Three numbers you will meet everywhere . . . 35

7. Choosing the right control chart . . . 36

The decision table . . . 36

Rational subgroups . . . 37

8. The I-MR chart: one reading at a time . . . 38

Running it . . . 38

Reading the chart . . . 39

Stable is not the same as good . . . 40

What to do about a red point . . . 41

A second example: patterns instead of points . . . 41

9. Xbar-R and Xbar-S: charts for subgroups . . . 43

Reading the Xbar-R chart . . . 44

Why the two charts disagree here . . . 45

Xbar-S . . . 45

10. The eight alarm rules . . . 47

Choosing which rules apply . . . 47

The zones . . . 47

Rule 1: one point beyond the limits . . . 48

Rule 2: nine points in a row on one side . . . 48

Rule 3: six points steadily rising or falling . . . 49

Rule 4: fourteen points alternating up and down . . . 49

Rule 5: two out of three points beyond 2 sigma . . . 50

Rule 6: four out of five points beyond 1 sigma . . . 50

Rule 7: fifteen points hugging the center . . . 51

Rule 8: eight points avoiding the center . . . 51

A summary to pin by the machine . . . 52

11. Limits that last: baselines, phases, locking and causes . . . 53

The four ways to set limits . . . 53

Before and after an improvement . . . 53

Fix the limits and lock them . . . 54

Recording causes . . . 55

12. Attribute charts: counting defects and defectives . . . 57

Defectives or defects? . . . 57

p chart . . . 57

np chart . . . 58

u and c charts . . . 59

Laney p': when there are thousands per sample . . . 60

G chart: rare events . . . 60

13. Special-purpose charts . . . 62

EWMA and CUSUM: catching small shifts . . . 62

Short-run charts: many part numbers, one process . . . 65

Pre-control: green, yellow, red . . . 65

14. SPC on the shop floor . . . 67

Data entry . . . 67

The SPC dashboard . . . 68

SPC alarms . . . 69

15. Process capability: Cp, Cpk, Pp, Ppk . . . 71

The idea in one picture . . . 71

Within and overall: Cpk or Ppk? . . . 71

How big is big enough? . . . 72

Four processes, four stories . . . 72

Every number on the capability page . . . 74

How many parts do you need? . . . 76

16. The capability report and non-normal data . . . 77

The capability report (sixpack) . . . 77

When the process is not stable . . . 79

Normal or not? The probability plot . . . 81
Non-normal capability . . . 82
Other capability tools . . . 83

17. Can you trust the gauge? Measurement systems . . . 84

Gage R&R; . . . 84
Attribute agreement: pass/fail inspection . . . 87
Bias and linearity . . . 89

18. Which problem first? Pareto and the fishbone . . . 92

The Pareto chart . . . 92
The fishbone (cause and effect) diagram . . . 93

19. Comparing groups: machines, shifts, suppliers . . . 95

Look first: box plots and mean plots . . . 95
Then test: ANOVA . . . 96
Two groups: the 2-sample t-test . . . 98
Significant is not the same as important . . . 99
Other comparisons . . . 99

20. Relationships: what drives the result? . . . 100

The scatter plot . . . 100
The correlation matrix . . . 101
Regression: how much, and which inputs matter . . . 102

21. Calculators: Z-scores and probabilities . . . 105

Z-score . . . 105
Probability calculator . . . 106

22. Reliability and designed experiments . . . 108

Weibull life analysis . . . 108
Designed experiments (DOE) . . . 109

23. Building a report, step by step . . . 113

Step 1: choose the report . . . 113
Step 2: Quick question! . . . 114

Step 3: review and build . . . 117
What you get . . . 118
The 23 reports . . . 120

24. Exporting results . . . 122
Export what is on screen . . . 122
Report packs . . . 123
Single charts . . . 124
A statement of the standard . . . 124

25. The Excel SPC calculator . . . 125

26. Accounts, settings and the audit trail . . . 132
Accounts . . . 132
The audit trail . . . 132
Settings and layout . . . 132
Saving your work . . . 133

27. Troubleshooting . . . 134
When something really is broken . . . 134

28. Glossary . . . 137

29. Command finder . . . 139
The Home tab . . . 139
The Explore tab . . . 140
The Quality tab . . . 141
The SPC tab . . . 142
The Compare Groups tab . . . 143
The Predict tab . . . 144
The Reliability tab . . . 145
The DOE tab . . . 146
The Visualize tab . . . 146
The Machine Learning tab . . . 147
The Data tab . . . 148
The Reports tab . . . 150

The View tab 150

PART I

Getting started

Install, sign in, find your way around the window, bring your data in, and run a complete study in fifteen minutes.

CHAPTER 1

Welcome to RS-1

RS-1 turns measurements into decisions: is the process stable, is it capable, can the gauge be trusted, what is causing the scrap, and what goes in the report to the customer.

What the program is for

Every manufacturing process varies. Two shafts turned on the same lathe a minute apart are never exactly the same diameter. Most of that variation is harmless noise, but some of it is a message: a tool is wearing, a fixture has loosened, a new lot of material behaves differently, the night shift sets the machine up another way. **Statistical process control (SPC)** is the discipline of telling the noise from the messages, quickly and with evidence, so that people act when something has changed and leave the process alone when it has not.

RS-1 does the arithmetic, draws the charts, applies the rules that flag real changes, writes a plain-English verdict under every result, and builds the reports your customers and auditors ask for. Your job is to bring the data, ask the right question, and act on the answer. This guide teaches all three.

Who this guide is for

- **Operators and inspectors** who enter readings and need to know what a red point on a chart means and what to do about it (Chapters 8, 10 and 14).
- **Quality engineers and technicians** who run capability studies, gauge studies and control charts, and answer to customers (Parts II to IV).
- **Supervisors and managers** who need one page that says which processes are fine and which need attention (Chapters 14 and 23).
- **Anyone new to statistics.** No formula knowledge is assumed. Where a number matters, the guide explains what it measures in words first and shows the formula only as a reference.

How the guide is laid out

The guide follows the order in which a real quality problem is usually worked:

Part	What it covers	Read it when
I. Getting started	Signing in, the window, loading data, a first complete study	You are new to the program
II. Control charts	Variation, choosing a chart, I-MR, Xbar, the eight rules, limits, attribute charts, special charts, the shop floor	You need to know whether a process is stable
III. Capability and measurement	Cp, Cpk, Pp, Ppk, non-normal data, Gage R&R, attribute agreement, bias and linearity	A customer asks for Cpk, or you doubt a gauge

Part	What it covers	Read it when
IV. Finding causes	Pareto, fishbone, comparing groups, relationships, calculators, reliability and experiments	Something is wrong and you need to know why
V. Reports and records	The report wizard, exports, the Excel calculator, accounts, audit trail, troubleshooting	You need to hand the results to someone
Reference	Glossary and command finder	You need a definition or a button

The boxes you will see

> **In real life**
> Shows how a chart or number is used on a real shop floor: what decision it supports and what people actually do next.

> **Tip**
> A shortcut or a better way to do something.

> **Watch out**
> A common mistake, or a result that is easy to misread.

> **Try it yourself**
> A short exercise on the example data that ships with the program. Doing these is the fastest way to learn.

Ribbon paths are written like this:

Where: SPC > Control Charts > Xbar-R / S

That means: click the **SPC** tab at the top of the window, find the **Control Charts** group, and click **Xbar-R / S**. You can also press **Ctrl+E** and type the command's name; see Chapter 2.

Every picture is real

Every screen and chart in this guide was produced by RS-1 itself, running the exact steps described, on the example data that ships with the program or on small teaching datasets built for a particular lesson. When you repeat a step on the same data you will see the same numbers. Your colors and text size may differ if you have changed them in Settings.

CHAPTER 2

Installing and signing in

A one-time setup: install, create the first account, and keep the recovery code somewhere safe.

Installing

1. Run **ArthurRS1-Setup.exe**. Windows may show a blue Windows protected your PC box for a new program; click **More info**, then **Run anyway**.
2. Accept the license, choose the install folder (the default is fine) and click **Install**. The installer adds a Start menu entry and, if you tick the box, a desktop shortcut.
3. Click **Finish** to start RS-1. On first start it checks that all of its parts are present and working; this takes a few seconds once.

> **Tip**
> Updates arrive the same way: the program checks for a newer version, shows what changed, and installs it only after verifying that it was signed by the publisher. Your data, settings and accounts are kept. To check now, click the small arrow beside the **?** button on the top row and choose Check for updates.

Activating your license

Each computer uses one license from your company's bundle. On the first start RS-1 asks for the license key you received by email, which looks like **RS1-7KQ2M-9XD4H-PAW3T-ZB8CN-5RVEJ**. Type or paste it and click **Activate**; the computer needs an internet connection for this. The key is not needed again on that computer, and reinstalling keeps the same license.

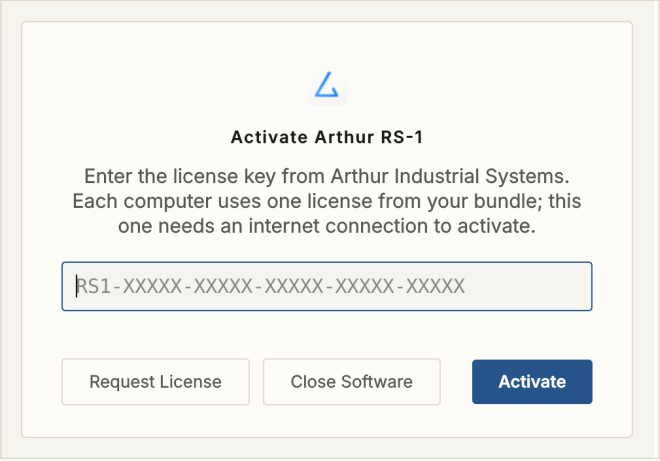


Figure 2.1 First start: enter the license key once.

RS-1 confirms the license when it starts and every few hours. Without internet it keeps working for up to **7 days** since the last confirmation, then asks you to connect. If your license is stopped or has ended, the program locks and shows who to contact; once it is restored, click **Check Again**.

> **Tip**
> Moving to a new computer? On the old one open **Help > Support** and click **Release This Computer's License**. Its license goes back to your bundle, ready for the new computer.

The first account

The first time RS-1 starts on a computer there are no accounts yet, so the sign-in window asks you to create the **administrator** account. The administrator can add other users later. Choose a user name and a password of at least eight characters; the bar under the password shows how strong it is.

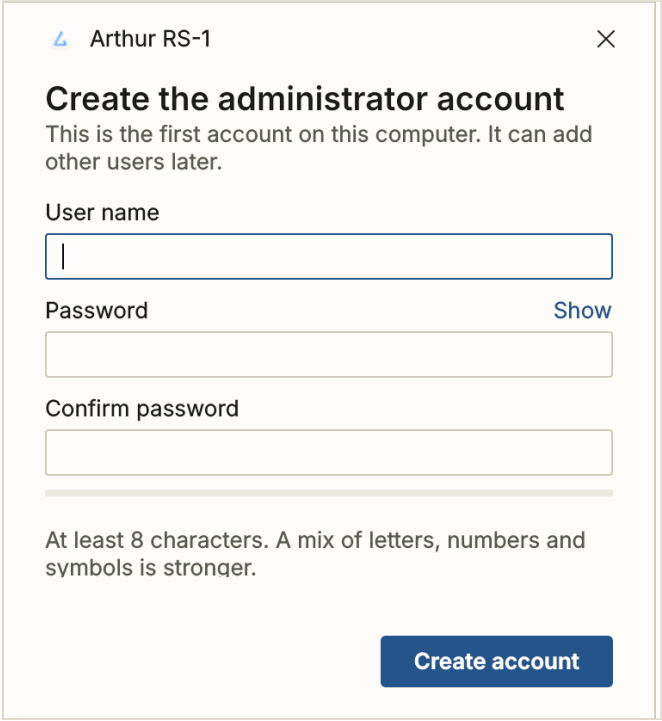


Figure 2.2 First start on a new computer: create the administrator account.

Each account gets a **recovery code** the moment it is created. It looks like **7KQ2-M9XD-4HPA-W3TZ** and is shown **once**. Write it down or store it in your password manager, tick **I have written this code down**, and click **Continue**.

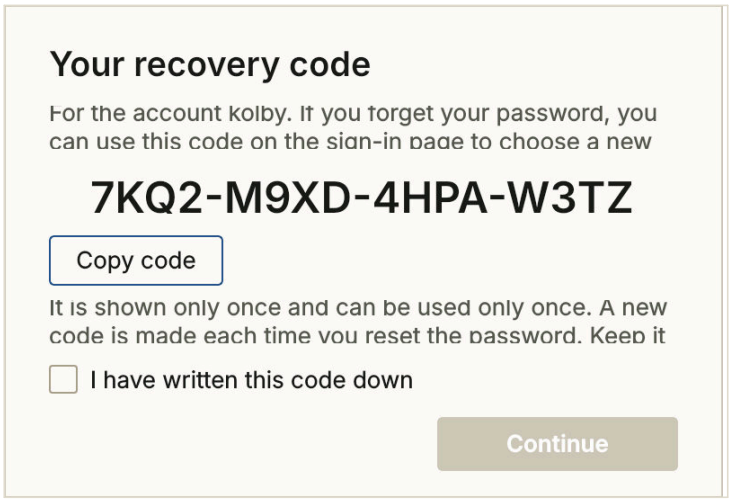


Figure 2.3 The recovery code. It works once; resetting your password gives you a new one.

Watch out

Nobody can read your password or recovery code back out of the program, including the administrator and Reliability Solutions. They are stored only as one-way fingerprints. Lose both and the account cannot be recovered; an administrator has to make you a new one.

Signing in every day

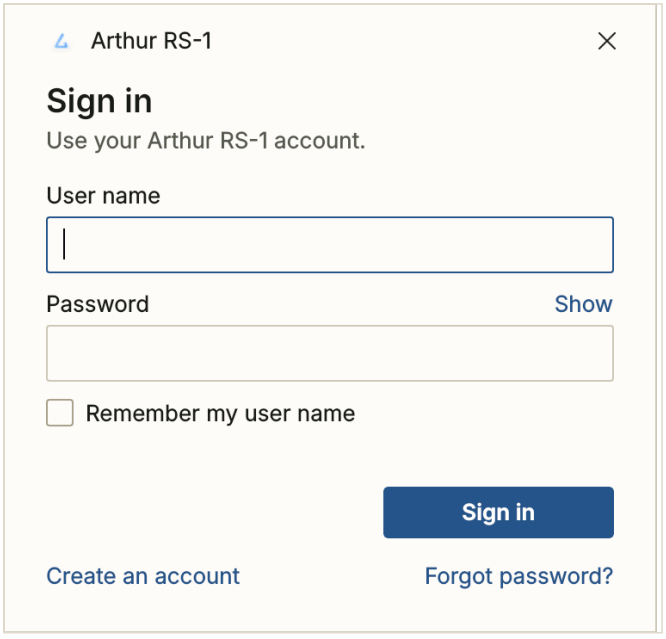


Figure 2.4 The everyday sign-in window.

- **Remember my user name** fills your name in next time; you still type the password.
- **Create an account** opens the sign-up form for another person on this computer.
- **Forgot password?** opens the reset form described next.

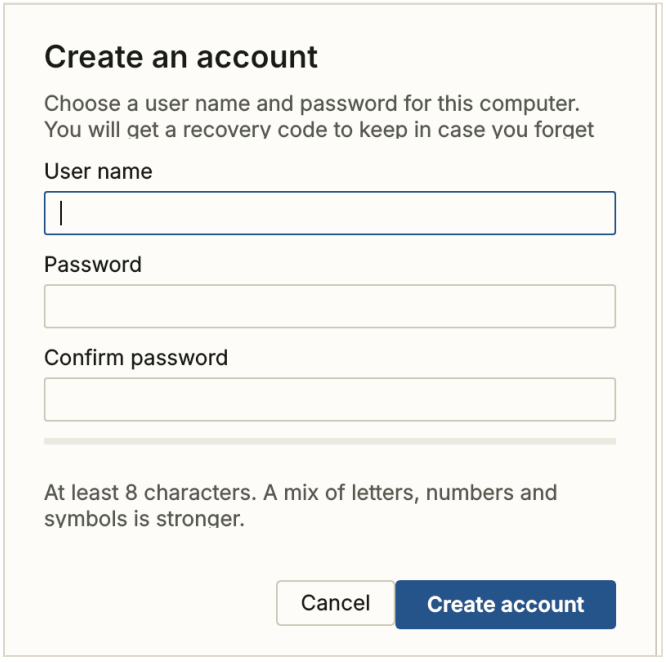


Figure 2.5 Creating another account from the sign-in window.

If you forget your password

1. On the sign-in window click **Forgot password?**
2. Enter your user name and your recovery code (the dashes are optional).
3. Choose a new password, type it twice, and click **Reset password**.
4. A **new** recovery code is shown, because the old one has now been used. Write it down the same way as before.

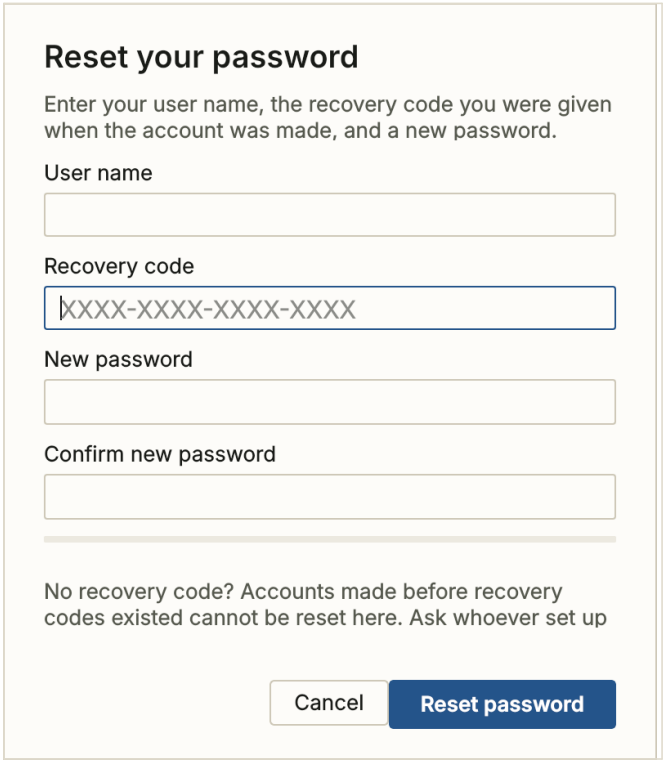


Figure 2.6 Resetting a password with the recovery code.

> **In real life**
> On a shared shop-floor PC, give each inspector an account of their own rather than sharing one. Every chart, cause note and report then carries the name of the person who made it, which is exactly what an auditor asks to see (Chapter 26).

CHAPTER 3

A tour of the window

Five areas: the top row, the ribbon, the columns panel, the data sheet and the results panel. Learn them once and every command works the same way.

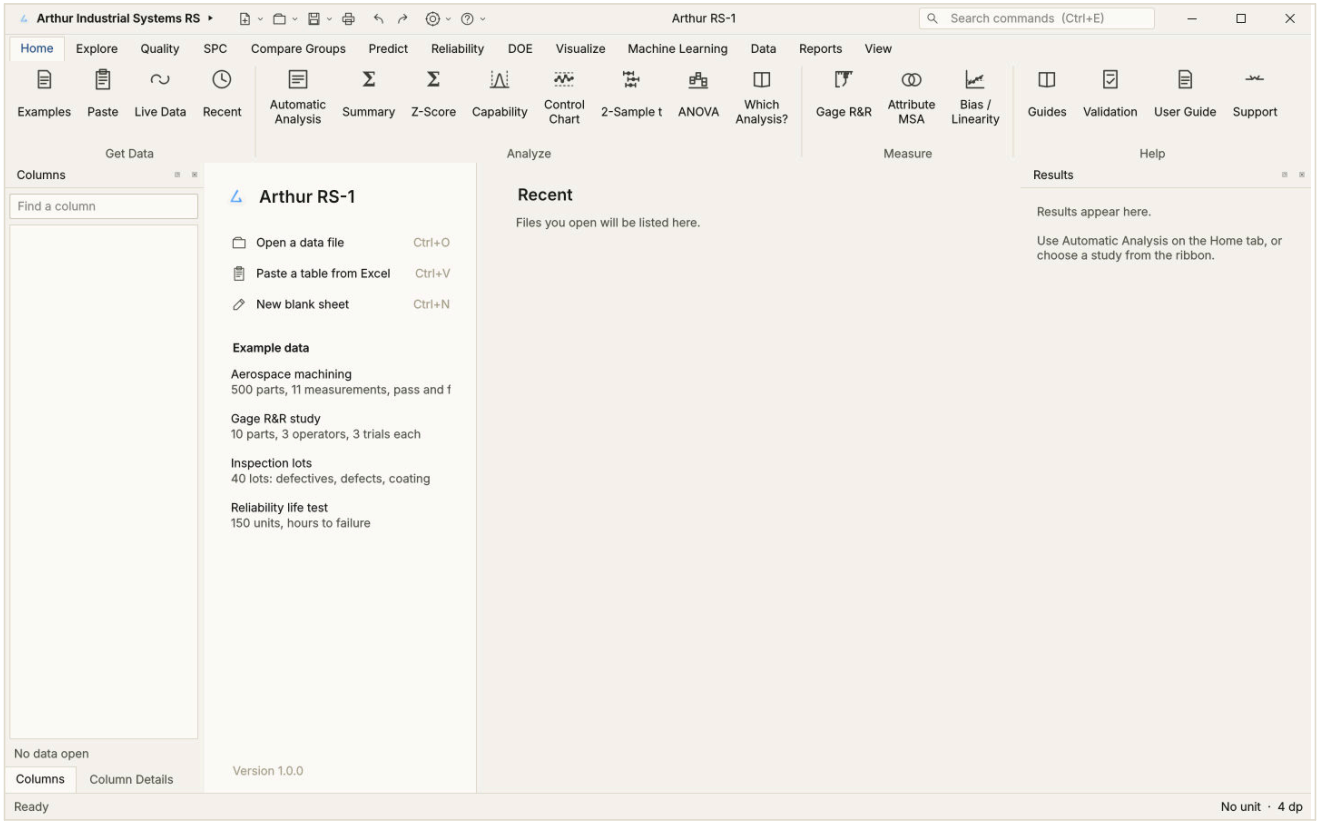
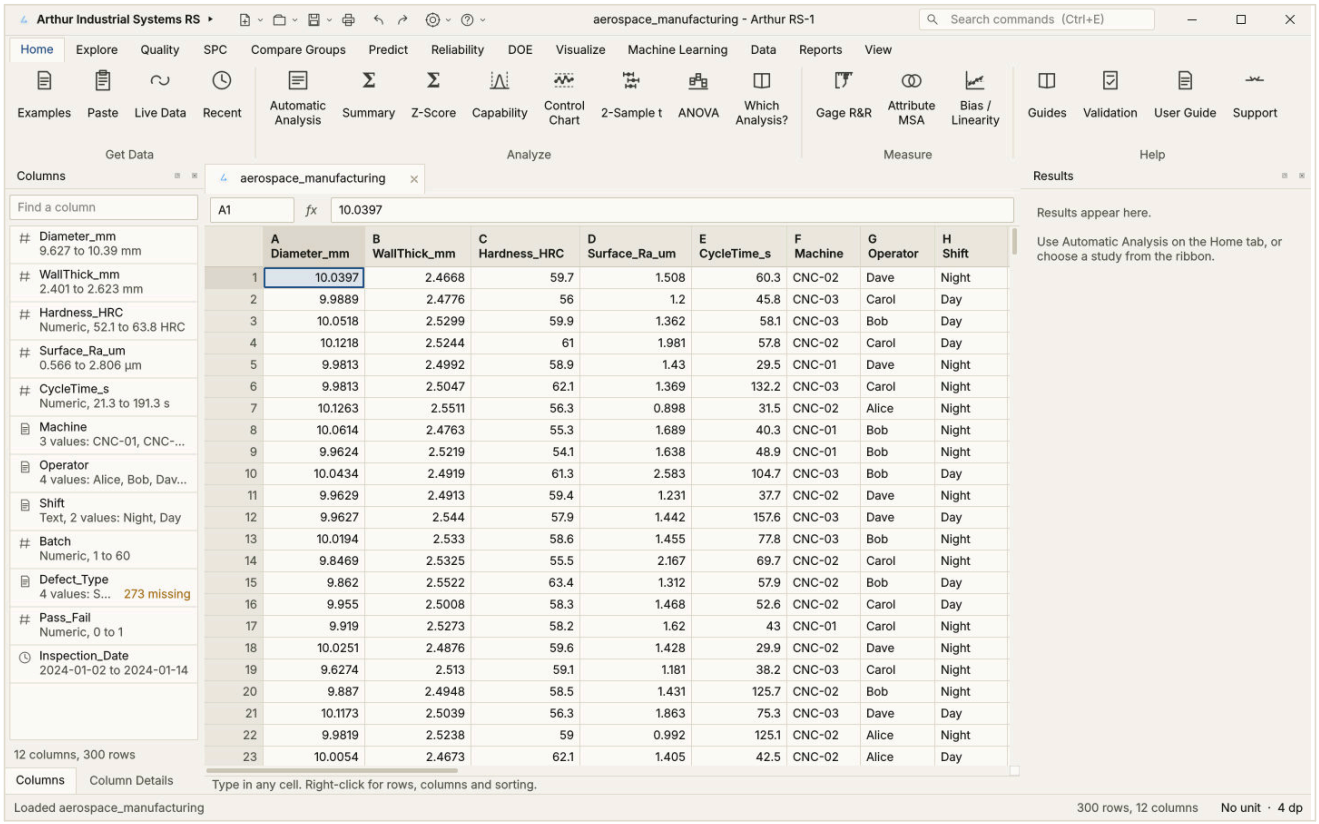


Figure 3.1 RS-1 just after signing in, before any data is open.

The start page

Before you open data, the middle of the window shows the start page. The three links at the top open a data file (**Ctrl+O**), paste a table you copied from Excel (**Ctrl+V**), or start a blank sheet you can type into (**Ctrl+N**). Under Example data are ready-made datasets used throughout this guide; click one to open it. Recent on the right lists files you opened before.

The window with data open



	A Diameter_mm	B WallThick_mm	C Hardness_HRC	D Surface_Ra_um	E CycleTime_s	F Machine	G Operator	H Shift
1	10.0397	2.4668	59.7	1.508	60.3	CNC-02	Dave	Night
2	9.9889	2.4776	56	1.2	45.8	CNC-03	Carol	Day
3	10.0518	2.5299	59.9	1.362	58.1	CNC-03	Bob	Day
4	10.1218	2.5244	61	1.981	57.8	CNC-02	Carol	Day
5	9.9813	2.4992	58.9	1.43	29.5	CNC-01	Dave	Night
6	9.9813	2.5047	62.1	1.369	132.2	CNC-03	Carol	Night
7	10.1263	2.5511	56.3	0.898	31.5	CNC-02	Alice	Night
8	10.0614	2.4763	55.3	1.689	40.3	CNC-01	Bob	Night
9	9.9624	2.5219	54.1	1.638	48.9	CNC-01	Bob	Night
10	10.0434	2.4919	61.3	2.583	104.7	CNC-03	Bob	Day
11	9.9629	2.4913	59.4	1.231	37.7	CNC-02	Dave	Night
12	9.9627	2.544	57.9	1.442	157.6	CNC-03	Dave	Day
13	10.0194	2.533	58.6	1.455	77.8	CNC-03	Bob	Night
14	9.8469	2.5325	55.5	2.167	69.7	CNC-02	Carol	Night
15	9.862	2.5522	63.4	1.312	57.9	CNC-02	Bob	Day
16	9.955	2.5008	58.3	1.468	52.6	CNC-02	Carol	Day
17	9.919	2.5273	58.2	1.62	43	CNC-01	Carol	Night
18	10.0251	2.4876	59.6	1.428	29.9	CNC-02	Dave	Night
19	9.6274	2.513	59.1	1.181	38.2	CNC-03	Carol	Night
20	9.887	2.4948	58.5	1.431	125.7	CNC-02	Bob	Night
21	10.1173	2.5039	56.3	1.863	75.3	CNC-03	Dave	Day
22	9.9819	2.5238	59	0.992	125.1	CNC-02	Alice	Night
23	10.0054	2.4673	62.1	1.405	42.5	CNC-02	Alice	Day

Figure 3.2 The aerospace machining example open: columns panel on the left, data sheet in the middle, results panel on the right.

Area	What it does
Top row	The program menu (the name at the far left), quick buttons for new, open, save, print, undo and redo, the settings gear, the help button, and the command search box (**Ctrl+E**).
Ribbon	The tabs (Home, Explore, Quality, SPC ...) and their buttons. Buttons are grouped by job: Get Data, Analyze, Measure, Help on Home, and so on.
Columns panel	Every column in the open data, with its type and a one-line summary: the range for numbers, the distinct values for text, and how many cells are missing. Type in Find a column to filter a long list.
Data sheet	The data itself, like a spreadsheet. Click a cell to see it in the box above the sheet; type to change it. Right-click for rows, columns and sorting. Select several cells and the bottom bar shows their count, sum, mean, min, max and standard deviation.
Results panel	Every analysis opens as a tab here: headings, numbers, charts and a plain-English verdict. Old results stay until you close their tabs.
Status bar	Messages on the left; the row and column count, the unit and the number of decimals shown on the right. Click the unit box to set the unit and decimals for the open data.

The ribbon tabs

Each tab collects the commands for one kind of question. The most used ones are also on **Home**, so a new user can do a great deal without leaving it.

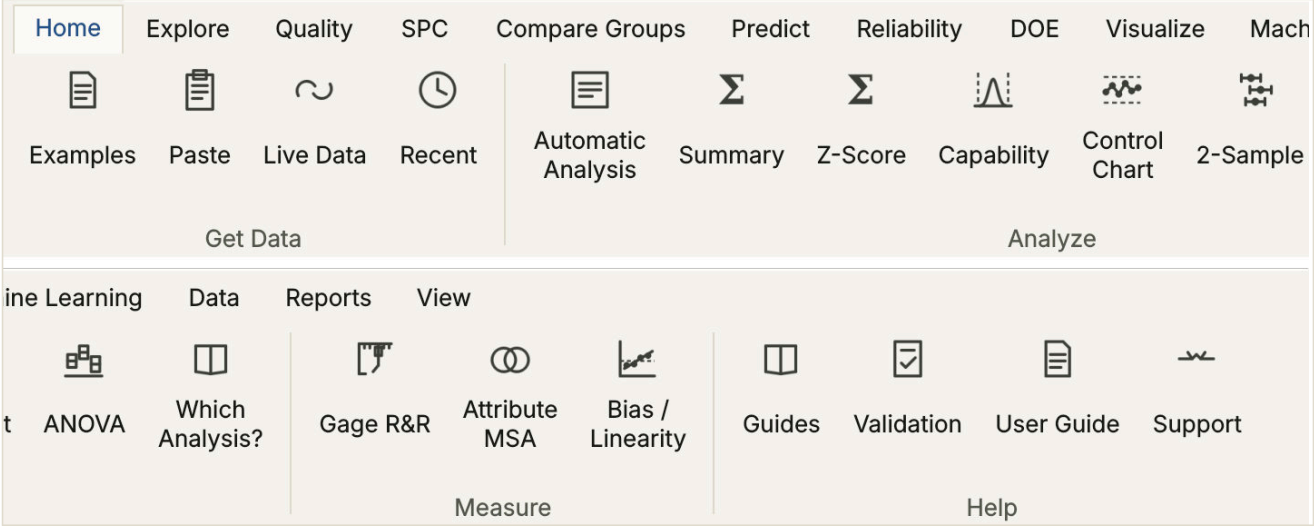


Figure 3.3 Home: get data, the everyday analyses, measurement-system studies and help.

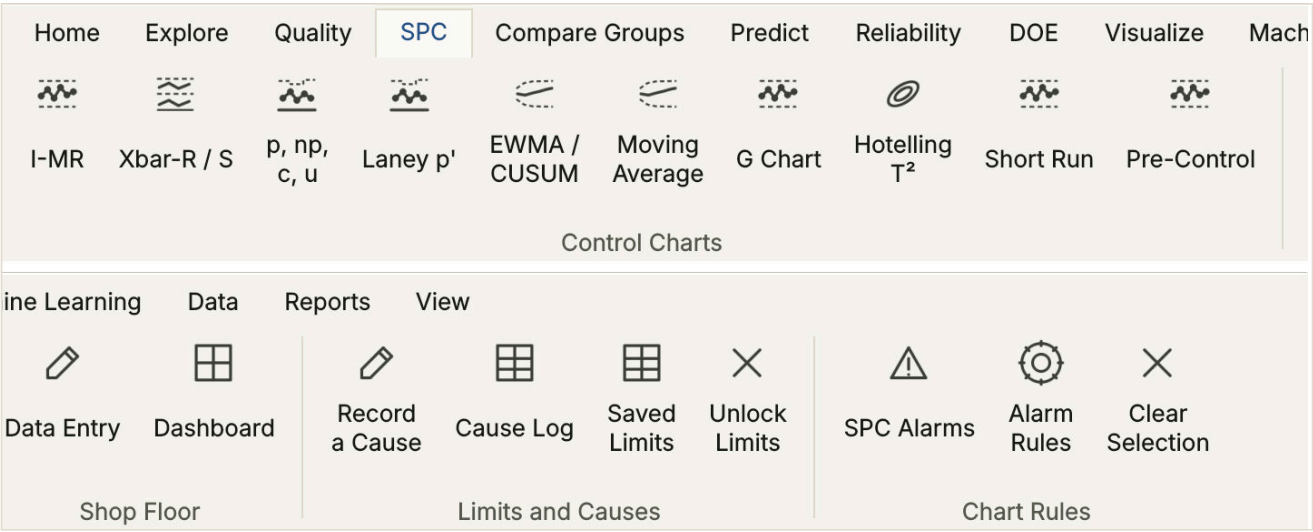


Figure 3.4 SPC: every control chart, the shop-floor tools, saved limits, causes and the alarm rules.

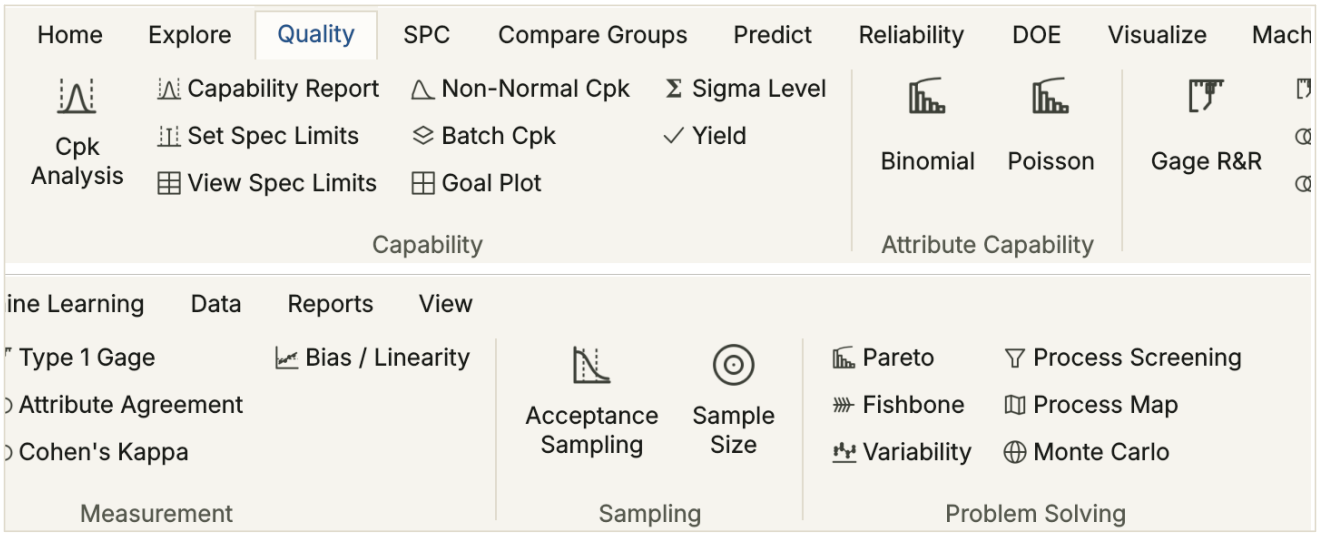


Figure 3.5 Quality: capability, spec limits, measurement systems, sampling and problem-solving tools.

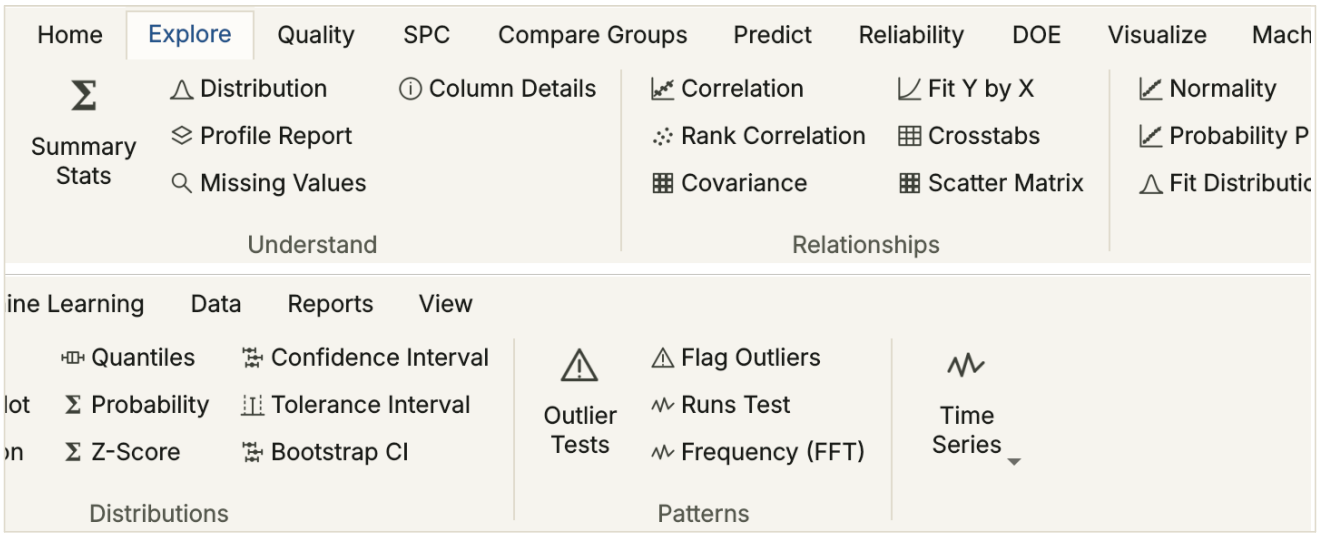


Figure 3.6 Explore: understand the data, distributions, relationships and patterns.

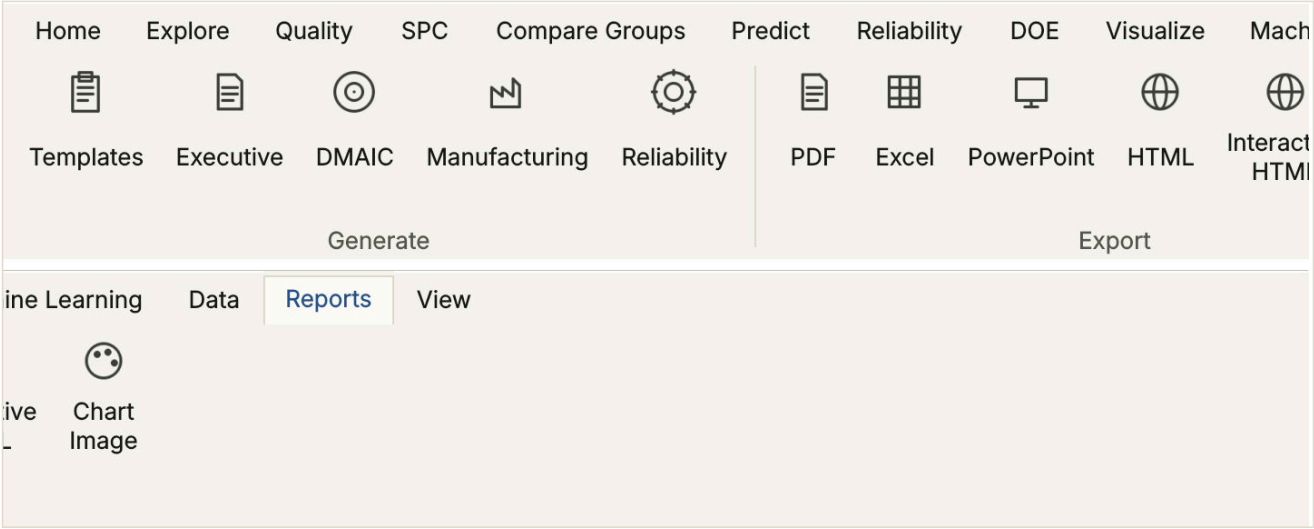


Figure 3.7 Reports: build reports and export results.

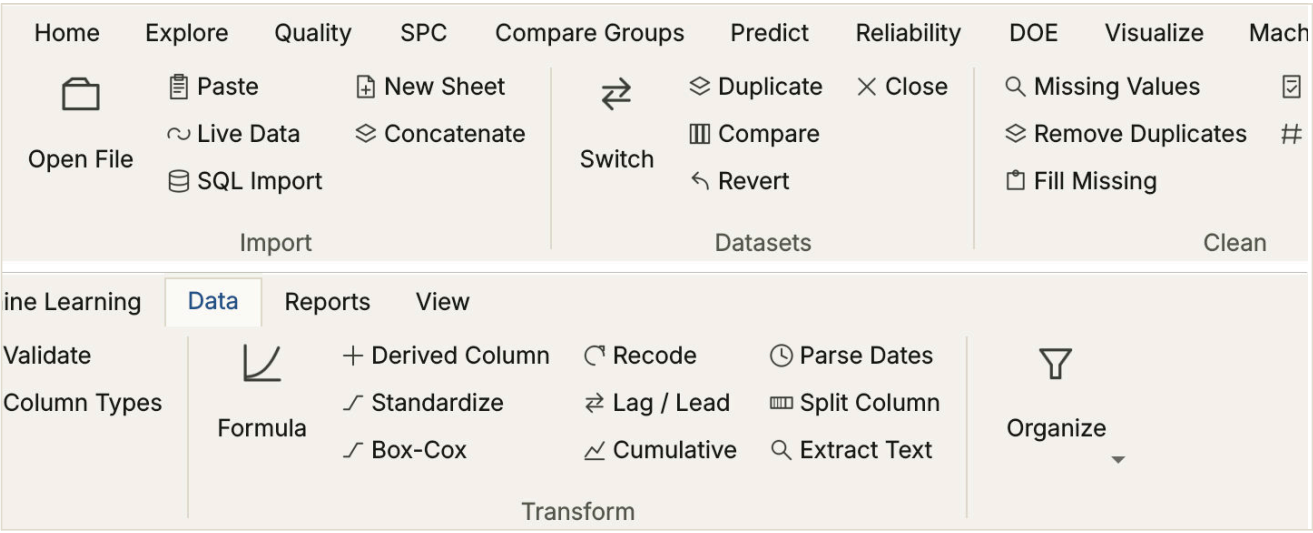


Figure 3.8 Data: import, clean, reshape and calculate columns.

> **Tip**
> Don't hunt for a button. Press **Ctrl+E**, type a few letters of what you want ("cpk", "gage", "pareto"), and press Enter. The search also understands plain words such as "is my process capable".

Reading a result

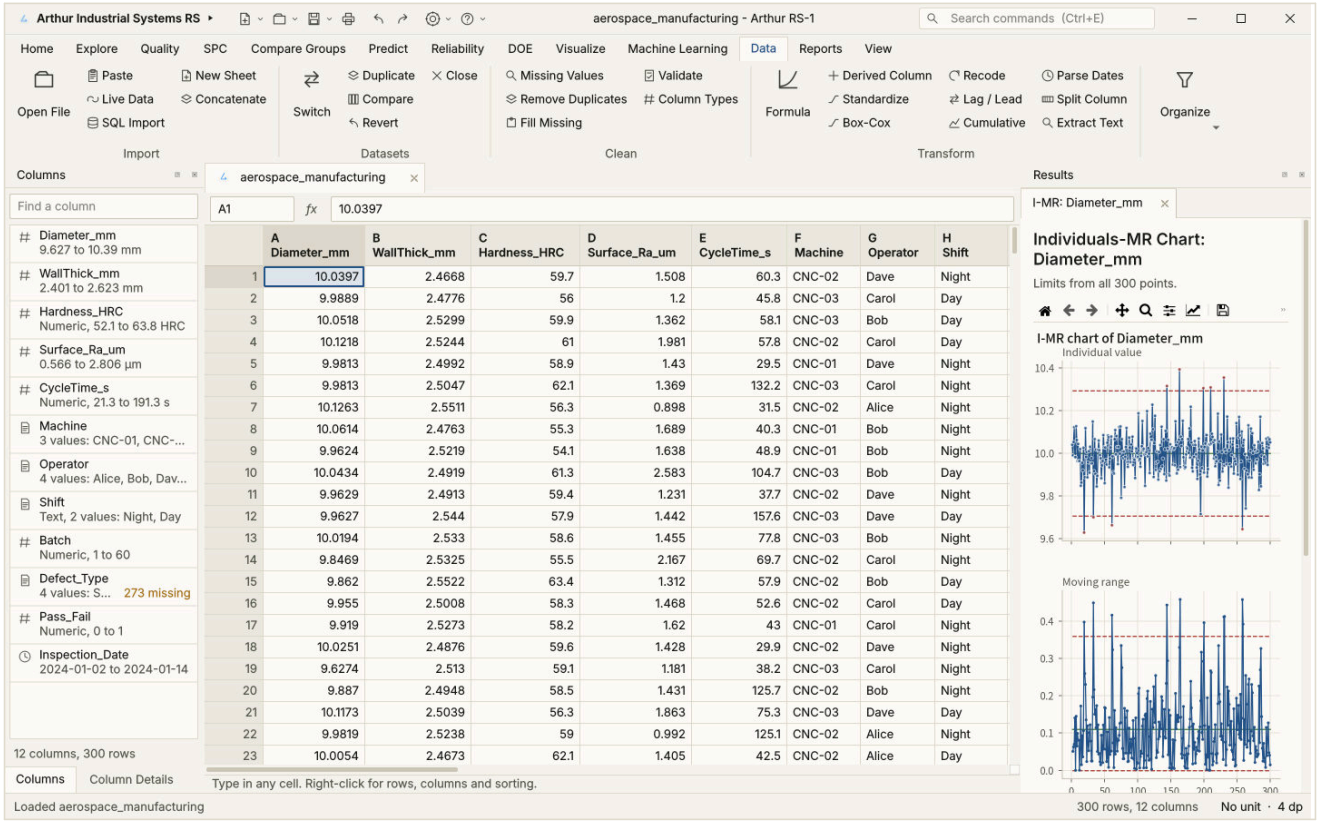


Figure 3.9 An I-MR chart opened in the results panel.

Every result follows the same pattern, from the top down:

1. A **heading** naming the analysis and the column.
2. A one-line **summary** of what was done (for example, which points the limits came from).
3. **Sections** with the numbers. Click a section's title to fold it away.
4. **Charts**, each with a small toolbar: reset view, back, forward, pan, zoom, settings, save the picture, and **Copy** to paste it into an e-mail or slide.
5. **Notes** with a symbol: a green tick for good news, a blue i for information, an amber triangle for something to look at. These are the program's verdict in plain words. Read them first.

> **Tip**
> Press **F1** at any time to open the on-screen guide page for the last command you used.

The columns panel and column details

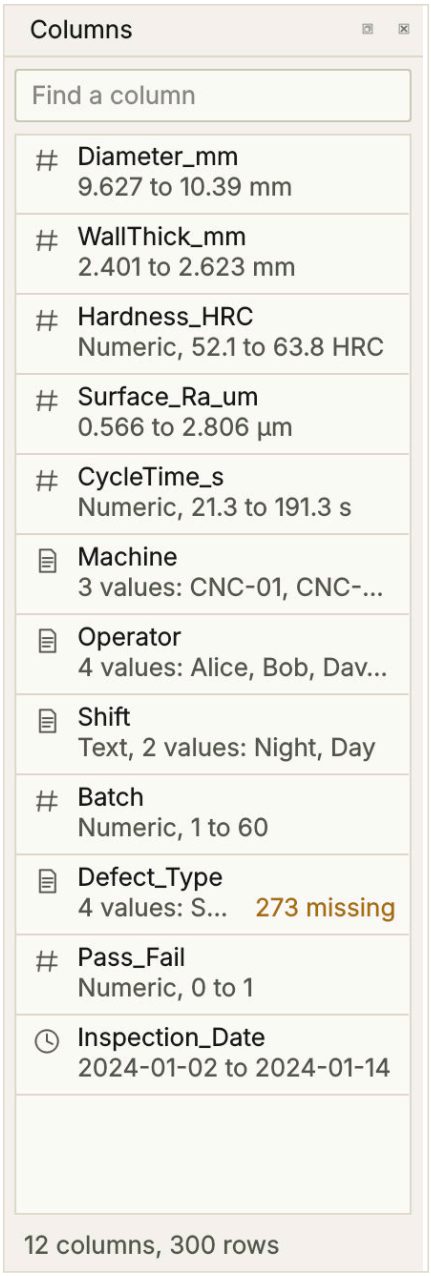


Figure 3.10 The columns panel. # marks numbers, the page icon marks text, the clock marks dates. 273 missing warns that most rows have no defect type.

Click **Column Details** under the panel to see the selected column in depth: a small histogram, its statistics, and a box for notes about the column (where the data came from, the gauge used). Notes travel with the data when you save a session.

CHAPTER 4

Getting your data in

RS-1 reads the files your plant already has. A few minutes spent on the shape of the data saves hours later.

Ways to bring data in

Source	How	Notes
Excel or CSV file	**Data > Import > Open File**, or Ctrl+O	Choose the sheet if the workbook has several. The first row is read as column names.
Copied cells	**Home > Get Data > Paste**, or Ctrl+V on the start page	Copy a block in Excel, including the header row, and paste.
A database	**Data > Import > SQL Import**	Type a connection address (for example **sqlite:///C:/data/plant.db**) and a query.
A file that keeps growing	**Home > Get Data > Live Data**	Watches a file written by a gauge or data collector and re-runs your charts when new rows arrive.
Typed readings	**SPC > Shop Floor > Data Entry**	One reading at a time, judged as it is entered (Chapter 14).
Examples	**Home > Get Data > Examples**	The datasets used in this guide.

The shape that works best

RS-1, like every statistics program, wants **one row per item measured and one column per thing recorded about it**. This is sometimes called tidy or long data.

Part	Machine	Operator	Shift	Diameter_mm	WallThick_mm
1	CNC-02	Dave	Night	10.0397	2.4668
2	CNC-03	Carol	Day	9.9889	2.4776
3	CNC-03	Bob	Day	10.0518	2.5299

- **Put readings in time order.** Control charts read the rows top to bottom as the order things happened. If the file is sorted by part number instead, sort it by date and time first (right-click a column > Sort).
- **Record the context.** A column for machine, operator, shift, lot and date costs nothing to collect and is what lets you find causes later (Part IV).
- **One kind of value per column.** Keep numbers numeric: write units in the column name (**Diameter_mm**), not in the cells ("10.02 mm").

- **Leave a missing reading empty.** Don't type 0 or "n/a"; a zero is a real measurement and will distort every result.

> **Watch out**
> A spreadsheet laid out for people (a header block with the part number, merged cells, a total row at the bottom, readings for each day across the page) has to be reshaped first. **Data > Organize > Stack** turns columns of readings into one long column with a label column beside it.

Spec limits: tell the program once

Many results need the **specification limits** from the drawing: the lower spec limit (LSL), the upper spec limit (USL) and the target. Set them once per column and every capability study, chart and report uses them.

Where: Quality > Capability > Set Spec Limits

1. Choose the column.
2. Type the LSL, the USL and the target. Leave a box empty for a one-sided specification (surface finish has only a maximum).
3. Click OK. **Quality > Capability > View Spec Limits** lists every limit set.

> **In real life**
> On a drawing that reads **Ø10.00 ±0.50**, the LSL is 9.50, the USL is 10.50 and the target is 10.00. For **Ra 3.2 max**, leave the LSL empty and set the USL to 3.2.

Check the data before you trust it

Two minutes of checking prevents the most embarrassing mistakes (a decimal point in the wrong place, a column of text that should be numbers):

- Look down the columns panel: does each range make sense? A diameter column ranging from 9.6 to 104 has a typing error.
- **Explore > Understand > Profile Report** summarizes every column at once, with missing values and suspicious entries.
- **Explore > Patterns > Flag Outliers** marks readings far from the rest. Investigate them; don't delete them just because they are inconvenient.

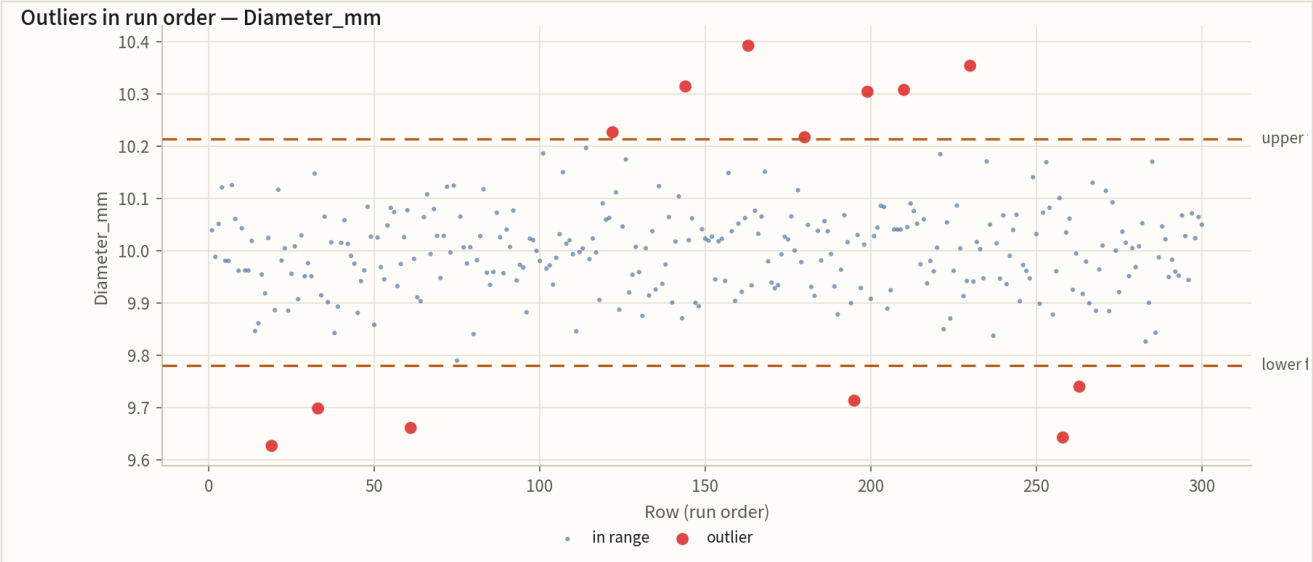


Figure 4.1 Flag Outliers on the aerospace diameters, in run order. Red points are outside the 1.5 x IQR fences.

Watch out

An outlier is a question, not a mistake. A diameter of 10.39 may be a typing error, or it may be the first part after the insert chipped. Check the paperwork before you remove anything, and never delete a reading only because it makes the capability number worse.

CHAPTER 5

Your first study in fifteen minutes

A complete, realistic piece of work on the example data: load it, set the limits, check stability, measure capability and hand over a report.

The situation: you machine a housing whose bore diameter is specified as 10.00 ±0.50 mm. The customer wants evidence that the process is capable before releasing the next order. You have 300 inspected parts.

Step 1: open the data

1. **Home > Get Data > Examples** and choose **Aerospace manufacturing**.
2. Glance at the columns panel: 300 rows, a diameter between 9.627 and 10.39 mm, three machines, four operators, two shifts.

Step 2: set the spec limits

1. **Quality > Capability > Set Spec Limits**, choose **Diameter_mm**.
2. LSL **9.5**, USL **10.5**, target **10.0**. Click OK.

Step 3: is the process stable?

Capability numbers only predict the future if the process is stable, so the first question is always the control chart.

Where: Home > Analyze > Control Chart

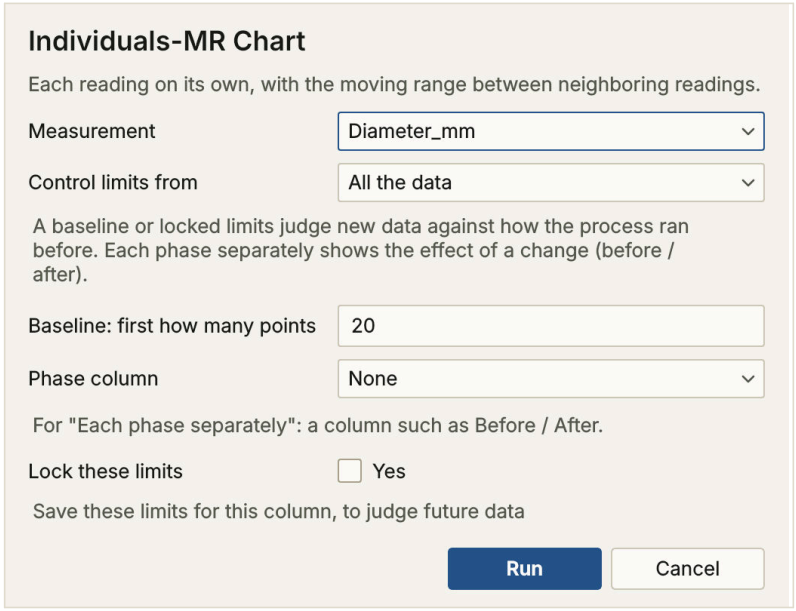


Figure 5.1 The I-MR chart's questions. The defaults are right for a first look.

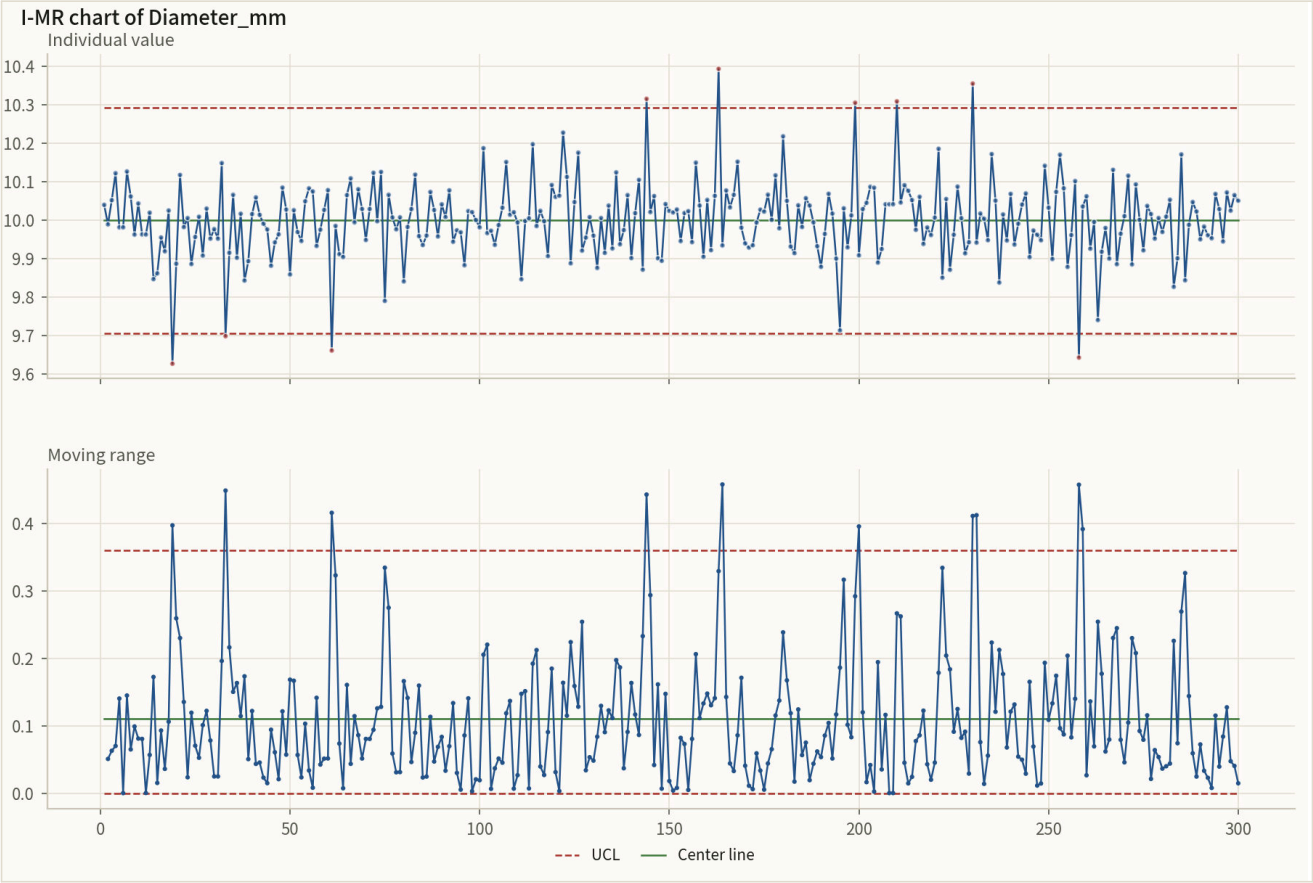


Figure 5.2 I-MR chart of the 300 bore diameters.

The points wander evenly around the center line at 9.9996 mm. The red dashed control limits at 9.706 and 10.293 show how far the process normally strays. Nine points fall outside them (observations 19, 33, 61, 144, 163, 199, 210, 230 and 258) and are marked red. Chapter 8 explains how to read this chart in detail; for now

the message is: **mostly stable, with occasional single excursions that deserve an explanation.**

Step 4: is it capable?

Where: Home > Analyze > Capability

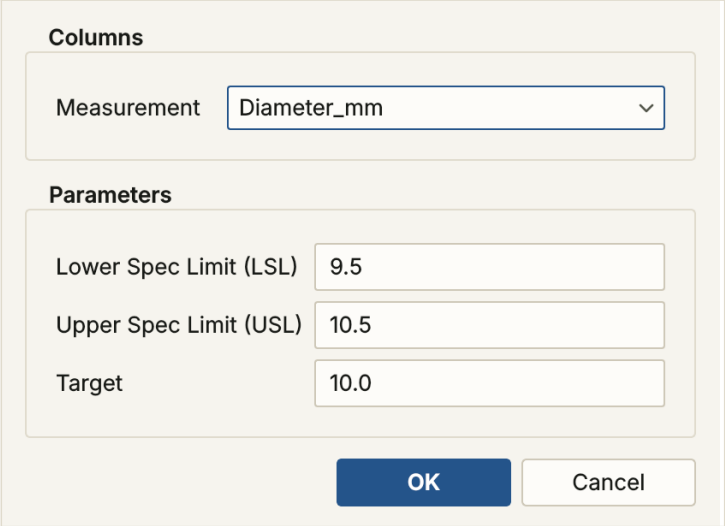
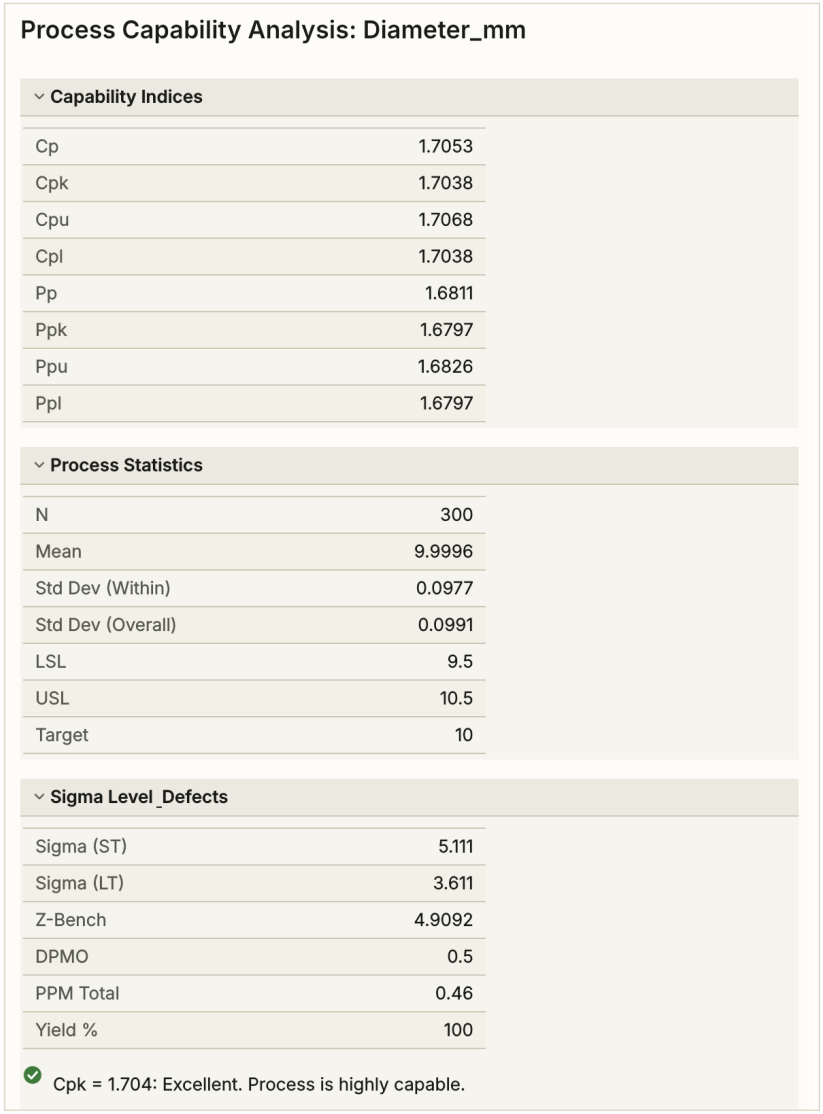


Figure 5.3 The capability study picks up the limits you set in step 2.



Process Capability Analysis: Diameter_mm

Capability Indices

Cp	1.7053
Cpk	1.7038
Cpu	1.7068
Cpl	1.7038
Pp	1.6811
Ppk	1.6797
Ppu	1.6826
Ppl	1.6797

Process Statistics

N	300
Mean	9.9996
Std Dev (Within)	0.0977
Std Dev (Overall)	0.0991
LSL	9.5
USL	10.5
Target	10

Sigma Level _Defects

Sigma (ST)	5.111
Sigma (LT)	3.611
Z-Bench	4.9092
DPMO	0.5
PPM Total	0.46
Yield %	100

Cpk = 1.704: Excellent. Process is highly capable.

Figure 5.4 The capability results.

Cpk = 1.70. The process spread fits inside the tolerance with a lot of room on both sides, and it is well centered (Cp 1.71 is almost the same as Cpk). The expected defect rate is under one part per million. Chapter 15 explains every number on this page.

Step 5: the report

1. **Reports > Generate > Templates** opens the report wizard.
2. Click **Cpk Validation**, then **Next**. The wizard already knows the limits you set, so it only asks for the ones it is missing.
3. Add the part number and customer, click **Next**, choose **PDF**, then **Build report**. The report opens when it is ready.

Chapter 23 walks through the wizard screen by screen.

In real life

This is the exact sequence a supplier quality engineer follows for a production part approval: show stability with a control chart, show capability with Cpk against the customer's minimum (usually 1.33, or 1.67 for a critical characteristic), explain any out-of-control points, and attach the report.

Try it yourself

Repeat steps 3 and 4 for **WallThick_mm** with limits 2.40 to 2.65. You should find a much lower Cpk. Which is the problem: the spread or the centering? (Compare Cp with Cpk, see Chapter 15.)

Run the I-MR chart on **Hardness_HRC**. It flags a trend and an alternating pattern rather than points outside the limits. Chapter 10 explains what those patterns usually mean.

PART II

Control charts

Is the process stable? How to tell a real change from everyday noise, which chart to use, how to read every rule, and what to do when a point turns red.

CHAPTER 6

Variation: the idea behind every chart

Every chart in this part answers one question: has something changed, or is this the ordinary scatter the process always has?

Two kinds of variation

Walter Shewhart, who invented the control chart in the 1920s, noticed that a process varies for two quite different reasons:

	Common-cause variation	Special-cause variation
What it is	The small, random scatter built into the process: machine vibration, tiny differences in material, temperature drifting a degree.	A specific event that is not part of the normal process: a worn tool, a new material lot, a wrong setting, an untrained operator.
How it looks	Points scatter evenly around the center, inside the limits, with no pattern.	A point far outside, a run on one side, a trend, a sudden jump.
Who fixes it	Management: only changing the process itself (a better machine, tighter material spec) reduces it.	The people on the line: find the event, remove it, prevent it coming back.
The mistake to avoid	Reacting to each point ("tweaking") makes the process worse.	Ignoring the signal lets scrap continue.

A process with only common-cause variation is called **stable** or **in statistical control**. It is predictable: next week it will behave as it did this week. That predictability is what makes capability numbers (Part III) meaningful.

Seeing the difference

The two run charts below come from a seeded teaching dataset. Both have the same amount of everyday noise. The second has one special cause: at reading 31 the process mean moved up by 2.5 units and stayed there.

Where: Explore > Time Series > Run Chart

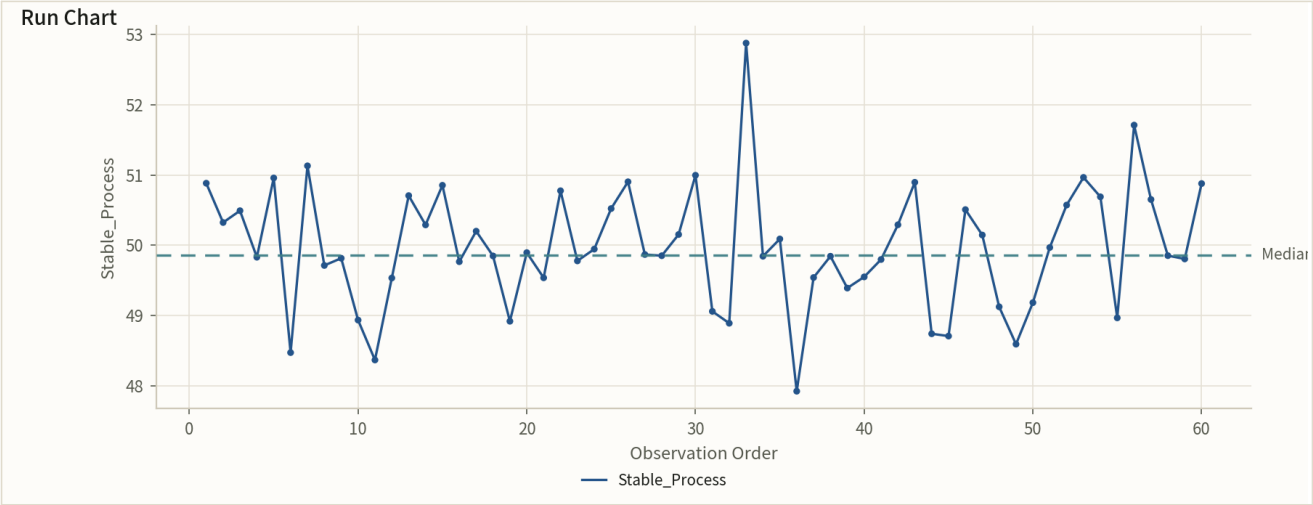


Figure 6.1 A stable process: 60 readings scattered evenly around the median (dashed line).

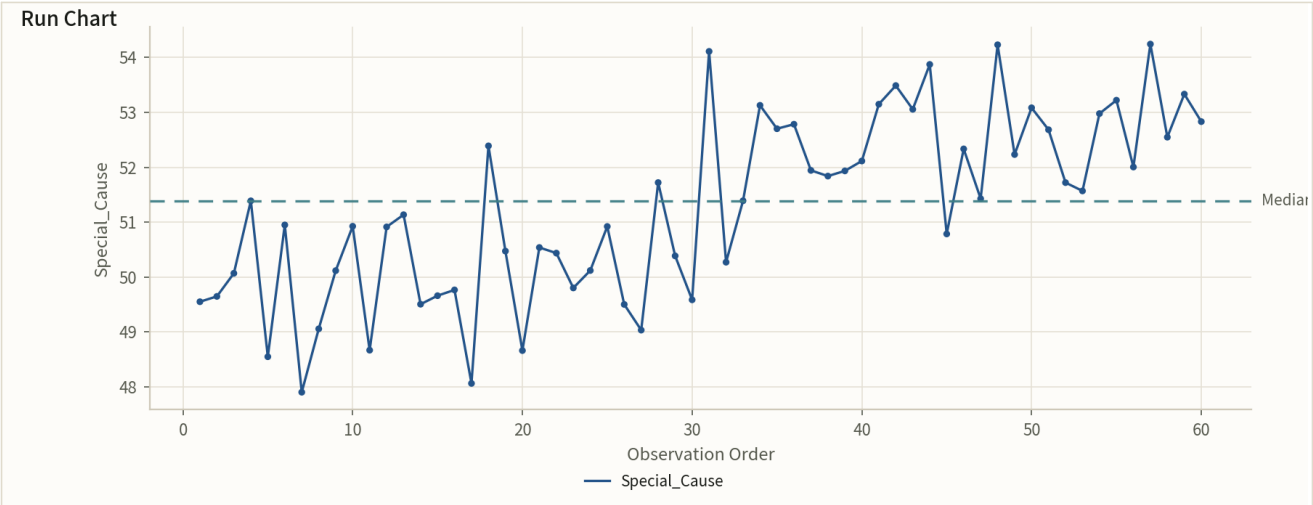


Figure 6.2 The same noise with a special cause: from reading 31 almost every point sits above the median.

Look at the second chart for a few seconds and the change is obvious. Looking at a table of the same 60 numbers, almost nobody would see it. **That is the whole point of charting data in time order.**

> **In real life**
> A filling line's weights look fine in the daily average, but the run chart shows the afternoon readings always sit higher. The cause turns out to be the product warming up after lunch and flowing faster. No average or histogram would have shown it, because both throw away the time order.

The shape of the scatter

A **histogram** stacks the readings into bars by size. It shows the center, the spread and the shape of the data, but not the time order. Most machined dimensions are close to the bell-shaped normal distribution; times, flatness, roundness and impurity levels usually are not.

Where: Visualize > Distribution > Histogram

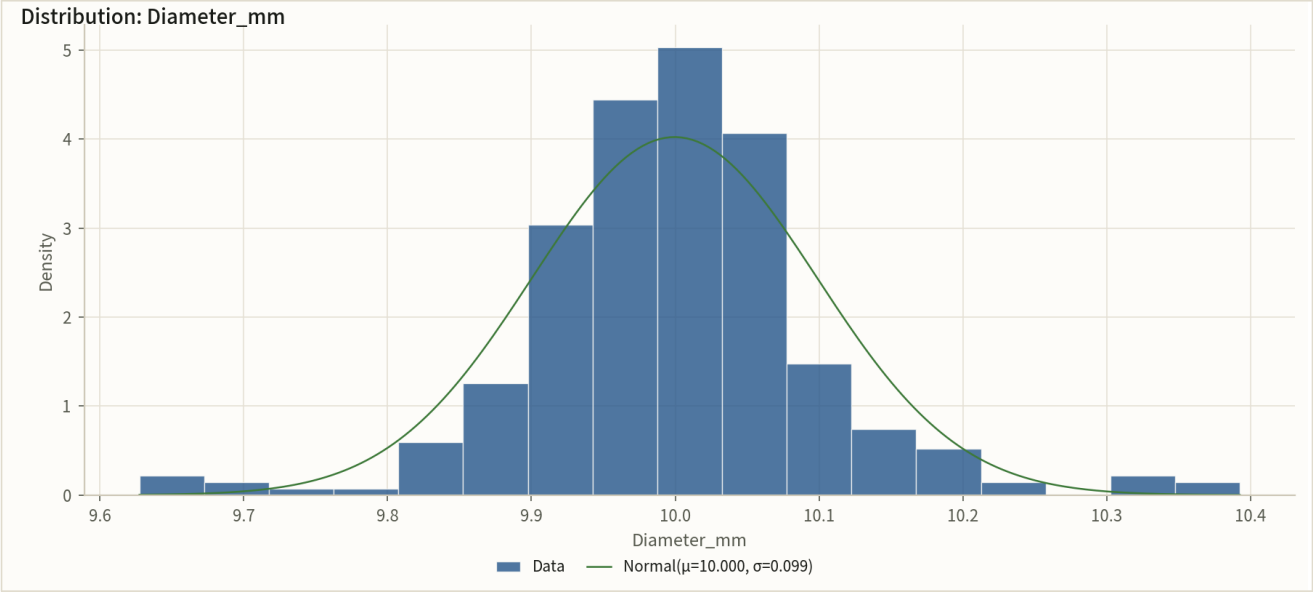


Figure 6.3 300 bore diameters: roughly bell-shaped and centered on 10.00 mm. The green curve is the normal distribution with the same mean and standard deviation.

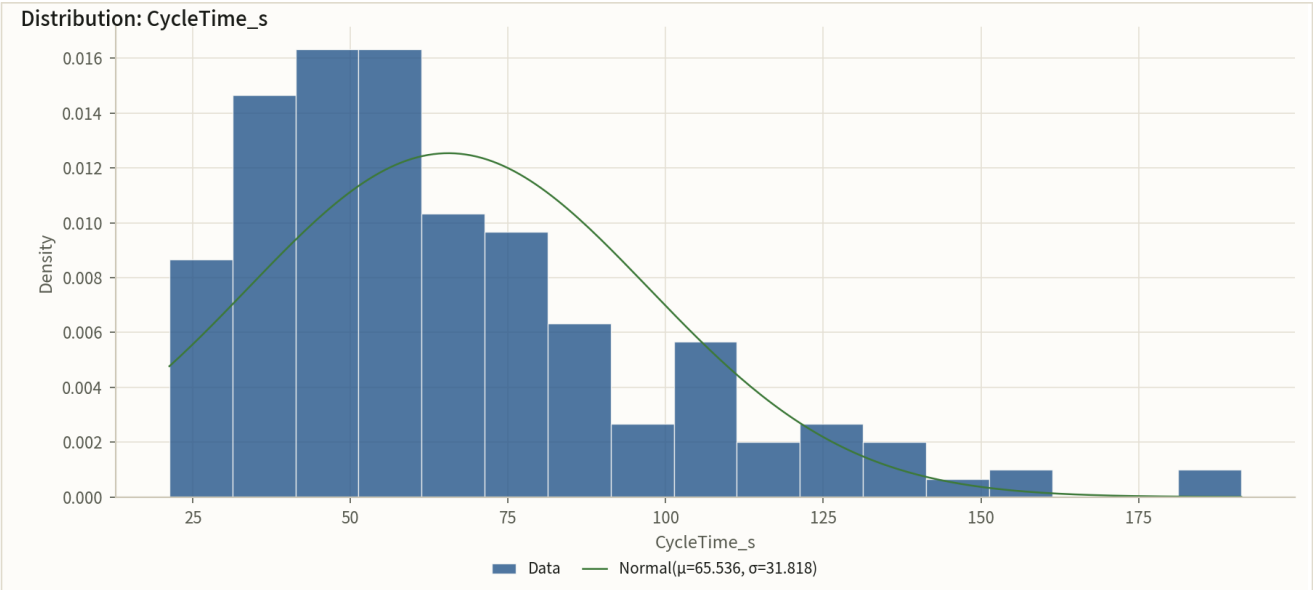


Figure 6.4 300 cycle times: skewed to the right. Most cycles take 35 to 60 s, a few take far longer. Times are rarely normal.

Watch out

A histogram can look perfectly normal while the process shifted halfway through: the two halves blend into one bell. Always look at the control chart as well as the histogram.

Three numbers you will meet everywhere

Number	What it tells you	In the example
Mean (average)	Where the process is centered.	Diameter 9.9996 mm
Standard deviation (sigma, σ)	The typical distance of a reading from the mean. About two thirds of a normal process lies within 1 σ of the mean, 95% within 2 σ and 99.73% within 3 σ.	0.099 mm
Range / **moving range**	The difference between the largest and smallest reading in a group, or between two neighboring readings. A quick measure of short-term spread.	Average moving range 0.110 mm

Tip

Select any block of cells in the data sheet: the bar under the sheet shows the count, sum, mean, min, max and standard deviation of the selection instantly.

CHAPTER 7

Choosing the right control chart

Two questions pick the chart: what kind of data do you have, and how is it collected?

The decision table

Your data	How it is collected	Chart	Chapter
A measurement (mm, kg, °C, s)	One reading at a time (each part, each batch, each shift)	**I-MR** (individuals and moving range)	8
A measurement	Small groups of 2 to 8 taken together (5 parts every hour)	**Xbar-R**	9
A measurement	Groups of 9 or more	**Xbar-S**	9
A measurement, small shift matters	One at a time, in order	**EWMA** or **CUSUM**	13
A measurement, many part numbers	Short runs on one process	**Short Run** (DNOM or Z-MR)	13
A measurement, simple go / adjust rules	At setup and during the run	**Pre-Control**	13
Defective items (pass / fail)	Samples of varying size	**p chart**	12
Defective items	Samples of the same size	**np chart**	12
Defects (scratches, voids)	Same area or amount inspected each time	**c chart**	12
Defects	Area or amount varies	**u chart**	12
Proportions from very large samples	Thousands per sample	**Laney p'**	12
Rare events (injuries, spills)	Count between events	**G chart**	12

> **Tip**
> Not sure? **Home > Analyze > Which Analysis?** asks a few questions about your data and points you to the right command.

Rational subgroups

When you sample in groups, choose the group so that **everything inside it was made under the same conditions**, a few consecutive parts from one machine at one time. Then the spread inside a group is pure common-cause noise, and any difference between groups is a candidate signal.

> **Watch out**
> Mixing machines or cavities inside one subgroup is the classic error. Five parts taken one from each of five spindles make the within-group range large, the limits wide, and the chart blind to a single spindle going wrong. Chart each spindle separately, or take the five parts from one.

CHAPTER 8

The I-MR chart: one reading at a time

The most used chart in industry. Each reading is plotted on its own, with a second chart of the change from one reading to the next.

Where: SPC > Control Charts > I-MR (also Home > Analyze > Control Chart)

Running it

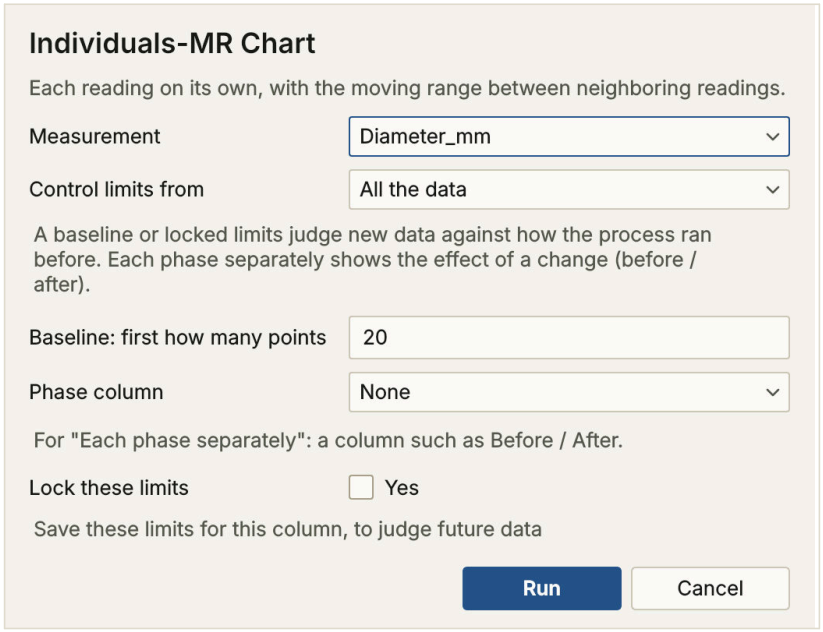


Figure 8.1 The I-MR chart's questions.

Question	What to choose
Measurement	The column to chart. Rows must be in time order.
Control limits from	All the data for a first look at history. A baseline (the first points) to judge new readings against a known good period. Saved (locked) limits to use limits saved earlier. Each phase separately to compare before and after a change (Chapter 11).
Baseline: first how many points	Only used with a baseline: how many of the first readings define normal. 20 to 30 is the usual minimum.
Phase column	Only for Each phase separately: a column naming the phases.
Lock these limits	Saves the limits for this column, to judge future data against them.

Reading the chart

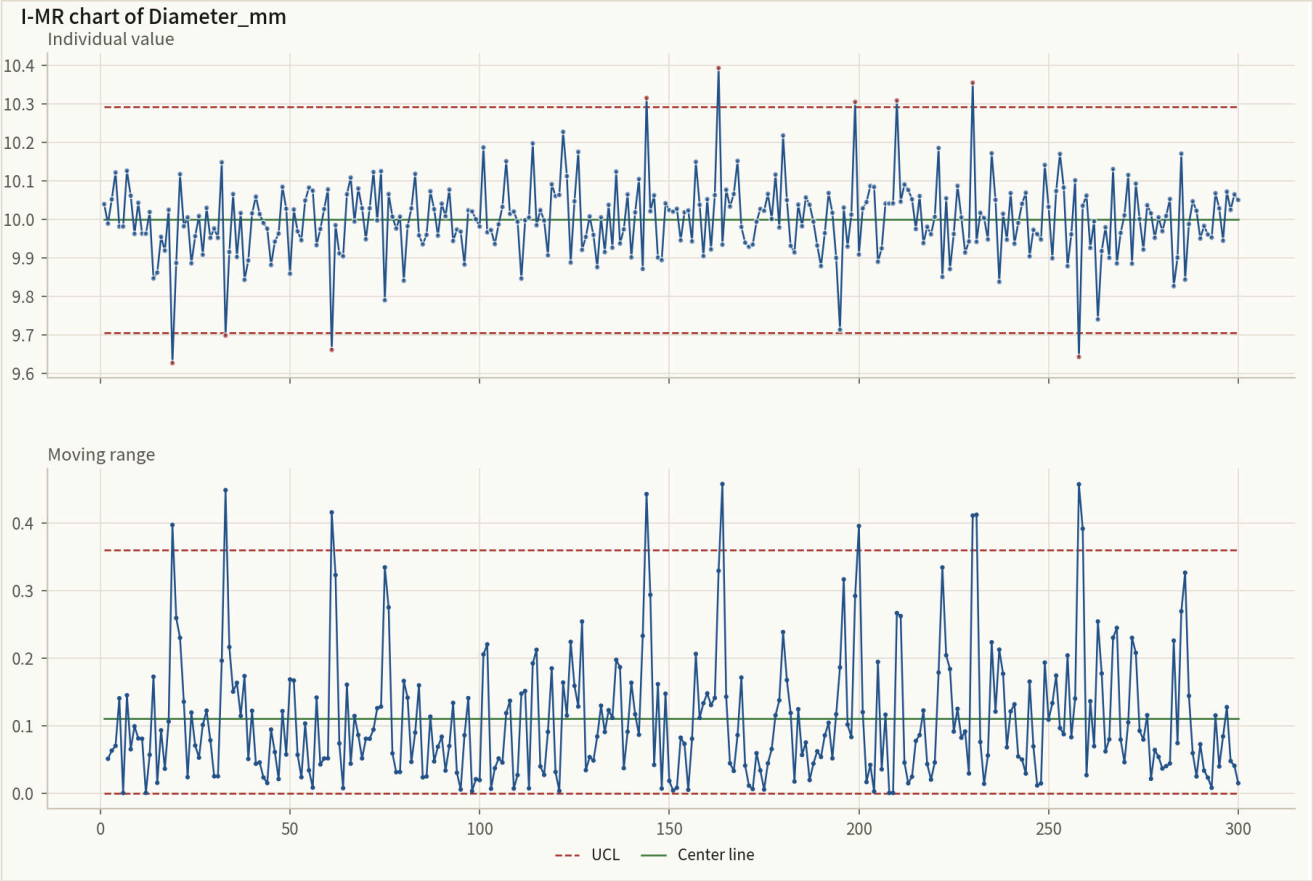


Figure 8.2 I-MR chart of 300 bore diameters from the aerospace example.

The top chart: individual values

- Each dot is one reading, in the order it was made.
- The **green center line** is the average: 9.9996 mm.
- The **red dashed lines** are the **control limits**, 3 sigma either side of the center: UCL 10.293 and LCL 9.706. They are calculated from the data, not from the drawing. They show where the process does run, not where it should run.
- **Red points** broke one of the alarm rules (Chapter 10). Here nine readings are beyond the limits.

The bottom chart: moving range

Each point is the absolute difference between a reading and the one before it. The moving range chart watches **short-term consistency**. A spike means two neighboring parts were unusually different, often the sign of a sudden event: a tool change, a material change, a gauge mis-read. Points above the MR UCL (0.360 here) deserve the same attention as points outside the top chart's limits.

The numbers



Figure 8.3 Under the chart: the limits and the list of signals.

Sigma here is estimated from the average moving range (MR-bar / 1.128), not from the overall standard deviation. That is deliberate: the moving range only sees short-term noise, so a shift in the data cannot inflate its own limits and hide itself.

Stable is not the same as good

A control chart says whether the process is **consistent**. It says nothing about whether the parts meet the drawing. A process can be perfectly stable and make every part out of tolerance, or wildly unstable and, for now, make every part good. Stability is judged with control limits; acceptability with spec limits and capability (Chapter 15).

	Control limits	Spec limits
Come from	The data (the voice of the process)	The drawing (the voice of the customer)
Answer	Has the process changed?	Is this part acceptable?
Act on a point outside by	Finding what changed	Containing the part
Change when	The process is deliberately changed and re-baselined	Engineering changes the drawing

> **Watch out**
> Never draw spec limits on a control chart and treat them as control limits. The two answer different questions, and mixing them leads either to ignoring real changes (specs wider than the process) or to constant adjustment (specs tighter than the process).

What to do about a red point

1. **Contain**: if the reading is also out of spec, segregate the part and anything made since the last good check.
2. **Look for the special cause**: what was different at that time? Check the tool log, material lot, operator, maintenance, the gauge. The columns for machine, operator and shift in your data are the first place to look.
3. **Fix and prevent**: remove the cause and change something so it cannot recur (a tool-life counter, a material check at receiving).
4. **Record it**: **SPC > Limits and Causes > Record a Cause** notes what happened and what was done against that point (Chapter 11). The next person to see the chart, and the auditor, will know the point was dealt with.

> **In real life**
> In the aerospace example, observations 19, 33 and 61 are low and 144, 163, 199, 210, 230 and 258 are high. On a real line you would pull the machine, operator and time for each of those rows. If most of the high ones came from one machine after a tool change, you would have found your special cause in five minutes.

A second example: patterns instead of points

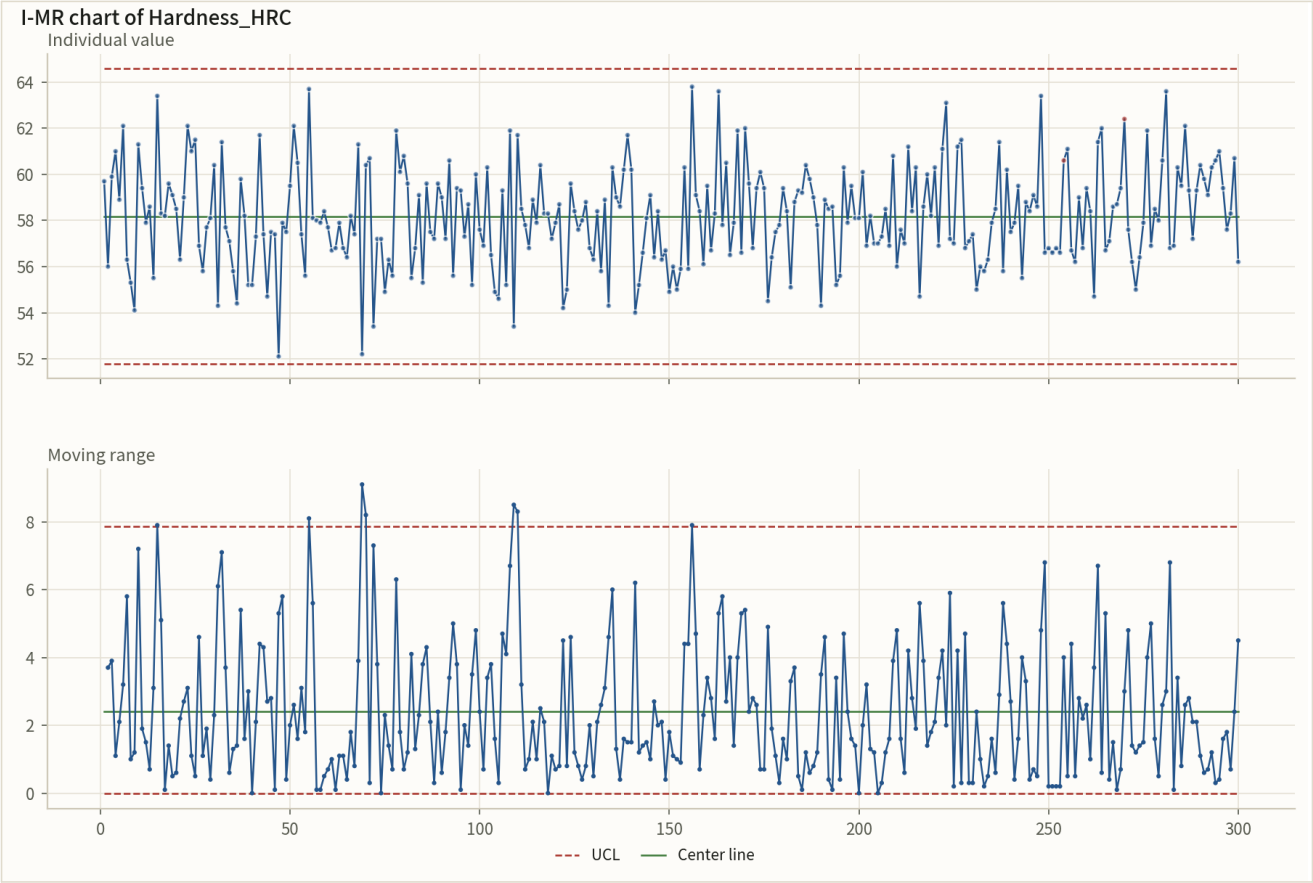


Figure 8.4 I-MR chart of Hardness_HRC. No point is outside the limits, but two pattern rules fire.

Here the program reports **Rule 3 (6 trending)** at observation 270 and **Rule 4 (14 alternating)** at observation 254. Nothing is out of the limits, yet the process is not behaving randomly. A trend in hardness often means furnace temperature drifting; a strict up-down alternation often means two sources taking turns, such as two furnaces or two fixtures feeding one inspection station. Chapter 10 shows each pattern on its own.

> **Try it yourself**
> Open the aerospace example and chart **Surface_Ra_um**. Is it stable? Then chart it again with Control limits from set to A baseline of 50 points. Do the signals change?

CHAPTER 9

Xbar-R and Xbar-S: charts for subgroups

When you measure a few parts together every hour, chart the group's average and its spread. Averages react faster to a shift than single readings do.

Where: SPC > Control Charts > Xbar-R / S

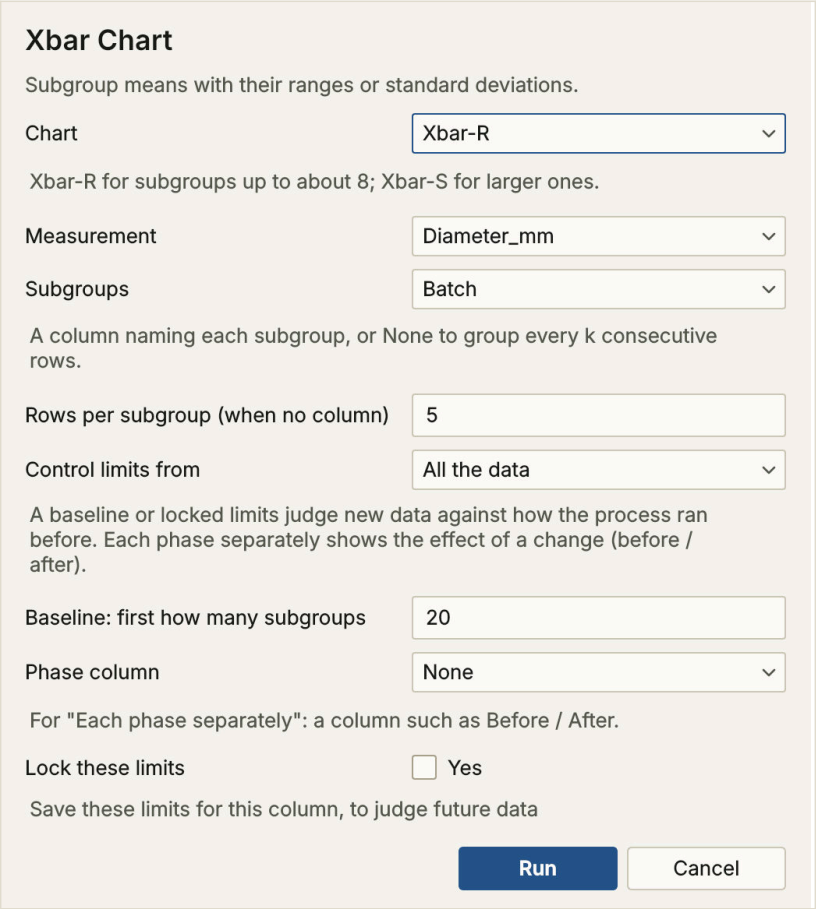


Figure 9.1 The Xbar chart's questions.

1. Choose **Xbar-R** for subgroups of up to about 8, **Xbar-S** for larger ones.
2. Choose the **Measurement**.
3. For **Subgroups**, choose the column that names each group (a batch, a sample time, a lot). If there is none, choose None and set **Rows per subgroup**: every 5 consecutive rows then form one group.
4. Choose where the limits come from, as for I-MR, and click **Run**.

Reading the Xbar-R chart

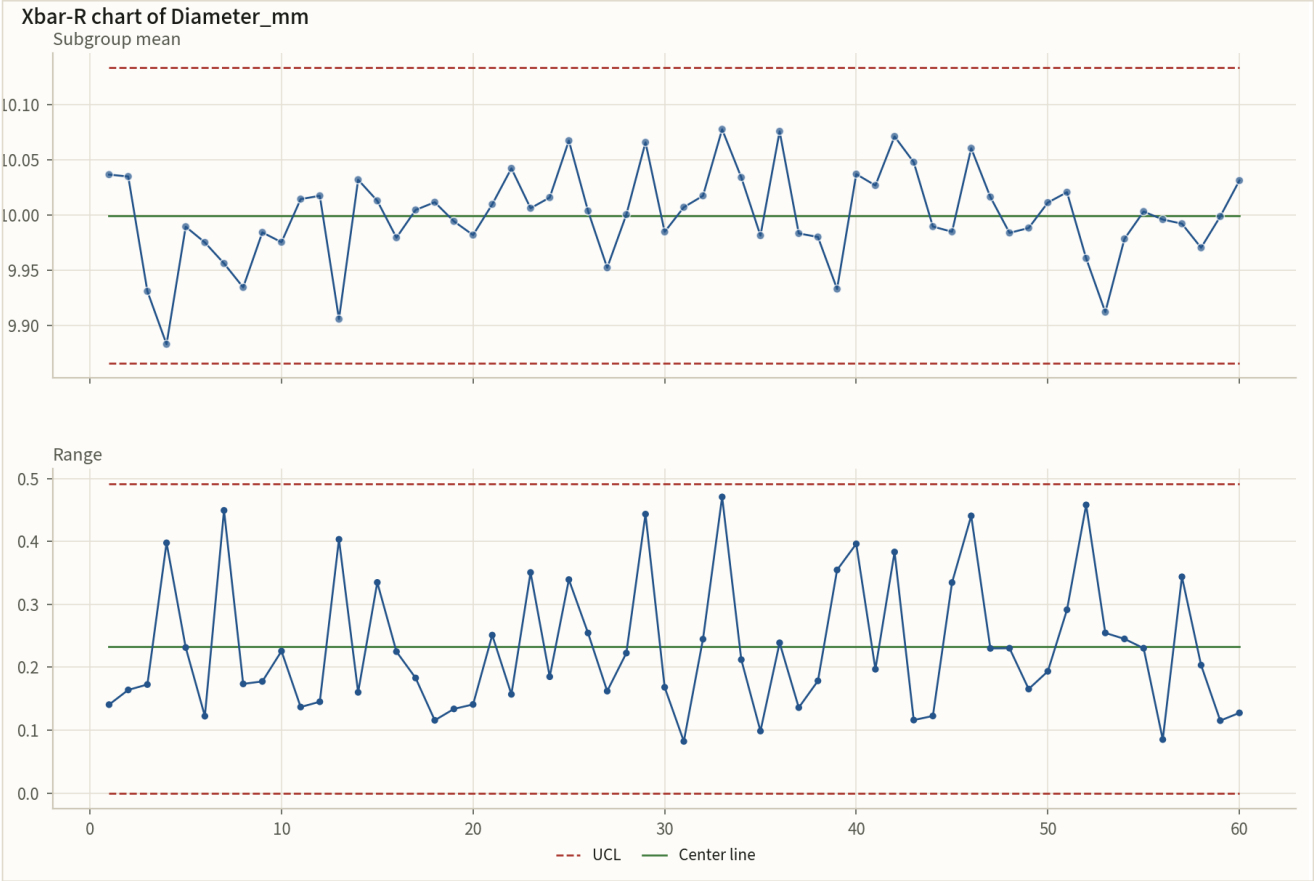


Figure 9.2 Xbar-R chart of the diameters, 60 batches of 5 parts.

- **Top: subgroup means.** Each point is the average of the 5 parts in one batch. The limits (9.866 to 10.134) are much narrower than the I-MR limits (9.706 to 10.293) because averages vary less than single parts, by a factor of the square root of the subgroup size.
- **Bottom: subgroup ranges.** Each point is the largest minus the smallest part in the batch: the spread within the batch. The average range is 0.232.

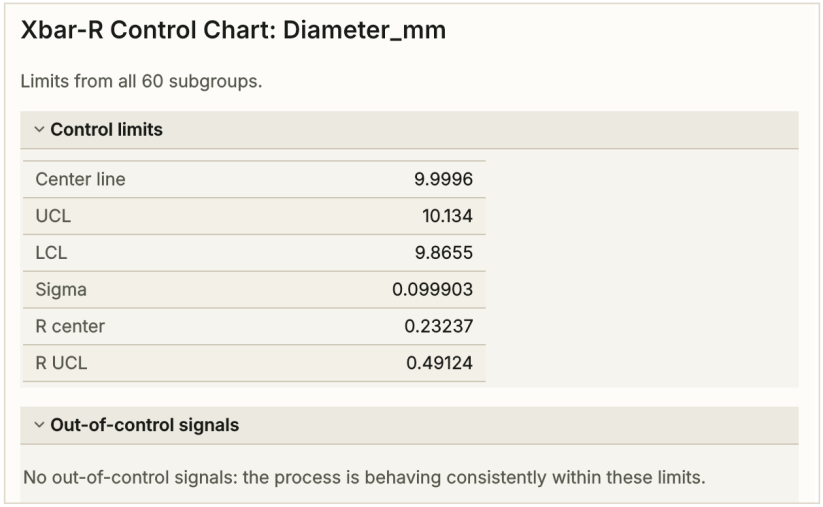


Figure 9.3 The limits, and the verdict: no out-of-control signals.

Read the bottom chart first. If the ranges are out of control, the within-group spread is unstable, and the limits on the top chart (which are built from the average range) cannot be trusted. Only when the ranges are stable does the means chart mean something.

Why the two charts disagree here

The I-MR chart of the same 300 parts flagged nine single parts; the Xbar-R chart flags nothing. Both are right. The nine parts were isolated extremes, and averaging each with four normal parts dilutes them. The Xbar chart is built to find **shifts in the process center**, which affect every part in a batch, and there are none. If isolated bad parts are what you need to catch, chart individuals.

Xbar-S

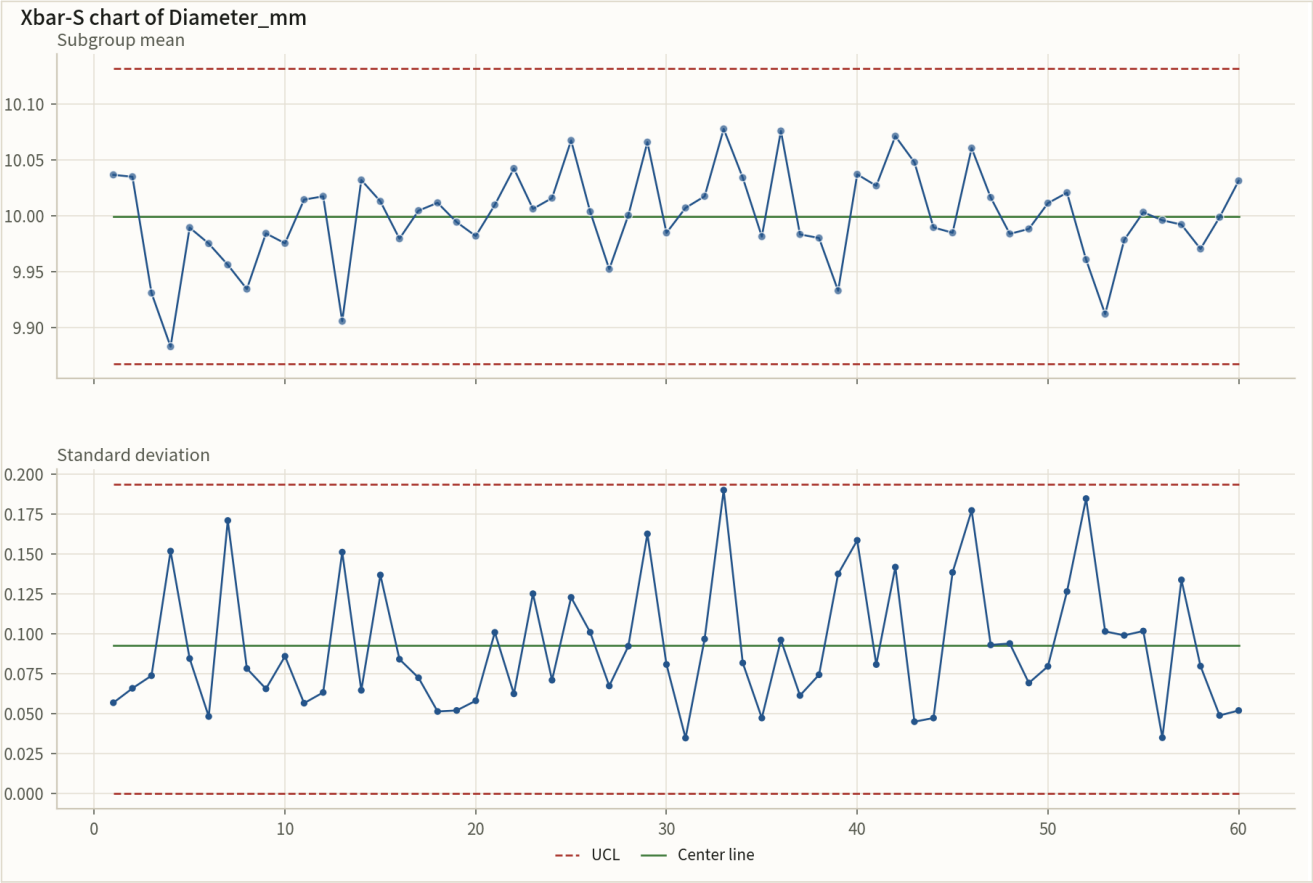


Figure 9.4 Xbar-S chart of the same data: the bottom chart shows each batch's standard deviation instead of its range.

With larger subgroups the range wastes information (it uses only the two most extreme parts), so the standard deviation of each subgroup is plotted instead. Read it exactly like Xbar-R.

> **In real life**
> A molding shop measures 5 parts from one cavity every 30 minutes and plots Xbar-R. When a heater band fails, the mold cools gradually and the means drift down for a few hours before any single part is out of spec. The means chart shows the drift by the second hour; the ranges stay flat, which tells the technician the problem is the process level, not its consistency.

CHAPTER 10

The eight alarm rules

Rule 1, a point beyond the limits, catches big changes. Rules 2 to 8 catch smaller changes and unnatural patterns long before a point leaves the limits.

Choosing which rules apply

Where: SPC > Chart Rules > Alarm Rules

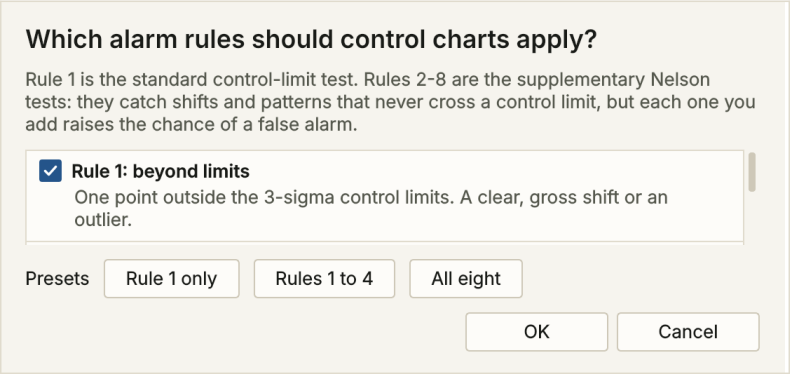


Figure 10.1 Choose the rules once; every control chart then uses them.

RS-1 starts with **Rules 1 to 4**, the common choice for production monitoring. The presets switch quickly between Rule 1 only, Rules 1 to 4 and All eight. Every extra rule finds smaller problems sooner, at the cost of a few more false alarms. A false alarm on any one stable chart is rare with one rule (about one point in 370) and becomes several times more likely with all eight.

The zones

The rules talk about **zones**. Between the center line and each control limit the chart is split into three equal bands, each one sigma wide:

Zone	Distance from the center	Share of a stable normal process
C	Within 1 sigma	about 68%
B	Between 1 and 2 sigma	about 27%
A	Between 2 and 3 sigma	about 4.3%
Beyond the limits	More than 3 sigma	about 0.27%

Each chart below comes from a teaching dataset: 25 calm readings that define the limits (Control limits from: A baseline of 25 points), then 15 more containing one pattern. All eight rules were switched on. Read the pictures as you would on a real chart: what would make a process do that?

Rule 1: one point beyond the limits

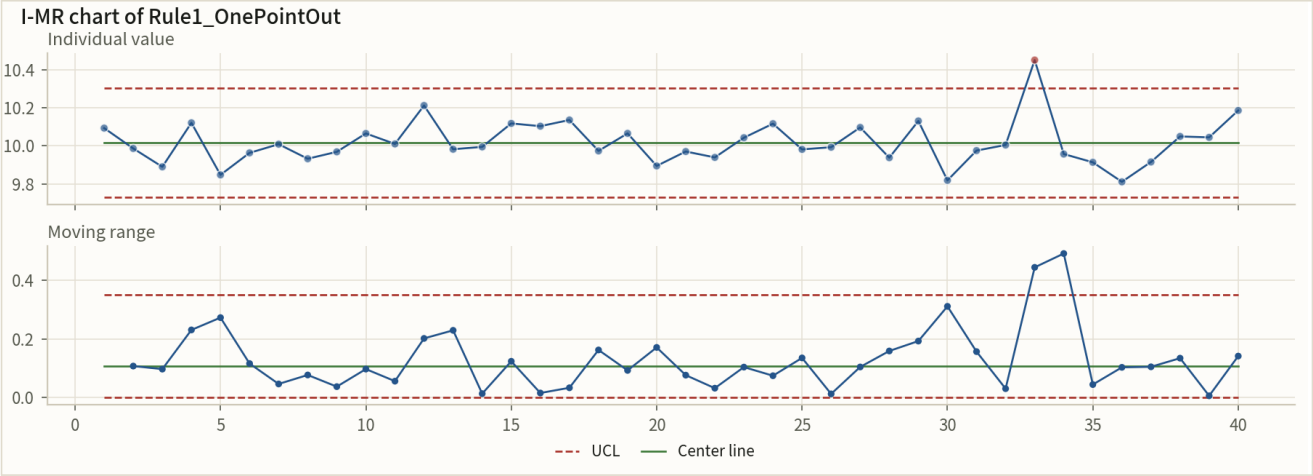


Figure 10.2 A single reading more than 3 sigma from the center: observation 33 jumps well above the UCL. Everything around it is normal.

What it usually means

A sudden, large, one-off event: a broken or chipped tool, a part loaded wrongly in the fixture, a measuring error, a typing error. Check the part and the gauge reading first.

Rule 2: nine points in a row on one side

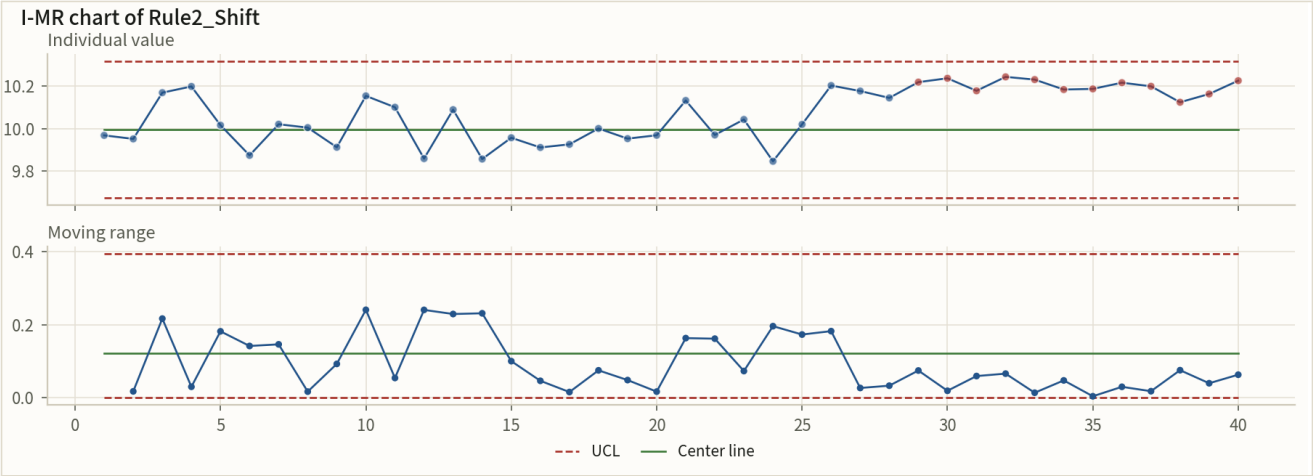


Figure 10.3 From observation 26 every reading sits above the center line. Rule 2 fires from observation 33 (once nine in a row have built up); Rules 5, 6 and 8 fire too, because the points crowd the upper zones.

What it usually means

The process level has moved and stayed: a new material lot, a setting changed, an offset entered, a different gauge, a new operator's technique. Look for what changed where the run **began** (observation 26), not where the alarm fired.

Rule 3: six points steadily rising or falling

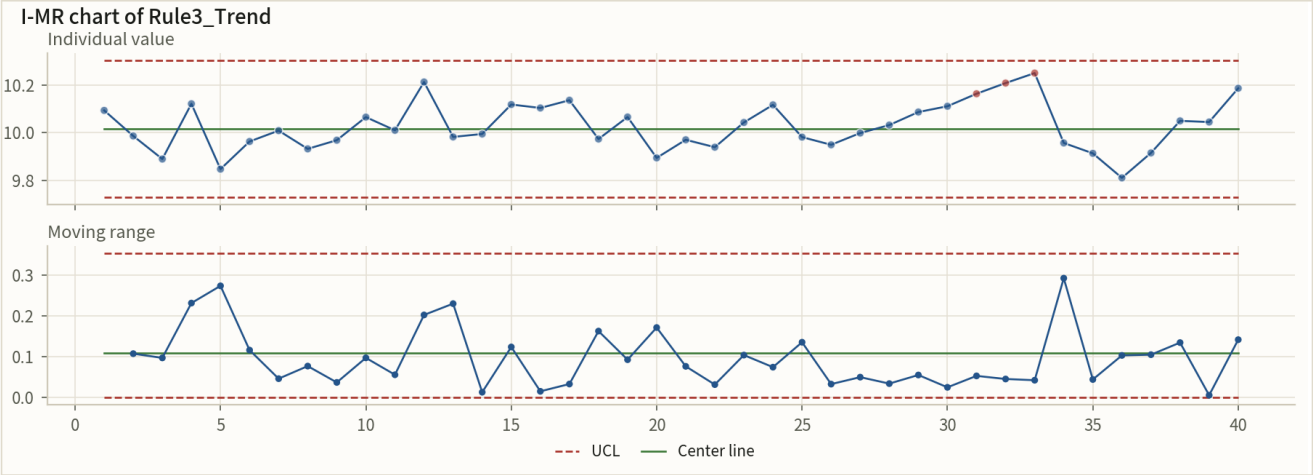


Figure 10.4 Observations 26 to 33 climb one after another; Rule 3 fires at 31 to 33, as soon as six increases in a row have been seen.

What it usually means

Something is wearing or warming: tool wear, a heating machine, a drying adhesive, a filter loading up. A trend predicts where the process is heading; act before it gets there.

Rule 4: fourteen points alternating up and down

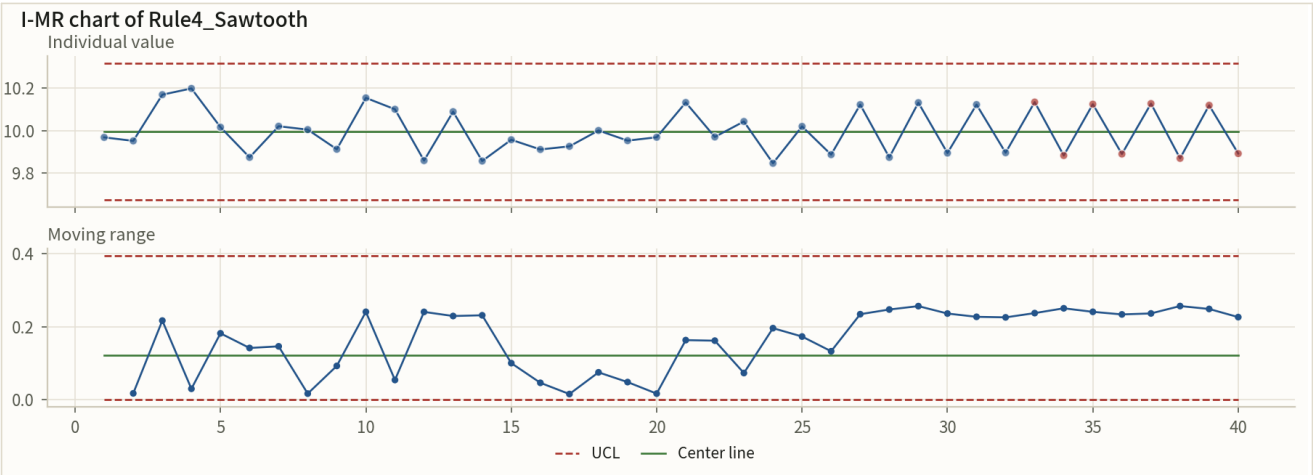


Figure 10.5 From observation 26 the readings zigzag high, low, high, low. Rule 4 fires from observation 33 onward.

What it usually means

Two sources taking turns: two machines, two fixtures, two spindles or two operators feeding the same chart. Also over-adjustment: an operator correcting every reading in the opposite direction.

Rule 5: two out of three points beyond 2 sigma

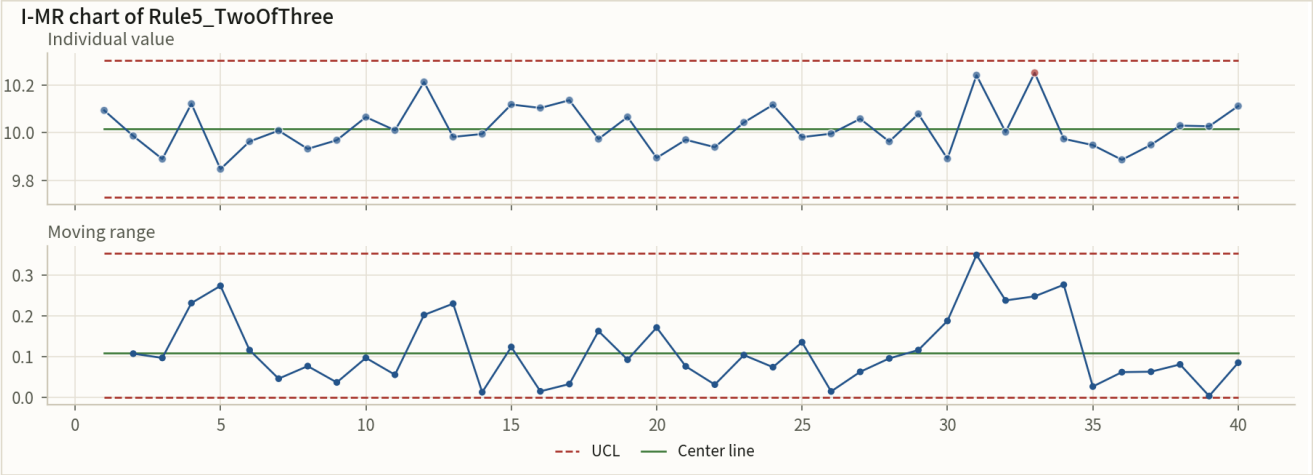


Figure 10.6 Observations 31 and 33 are both high, near the limit but not over it. Rule 5 fires at 33. Rule 1 alone would have stayed silent.

What it usually means

An early warning of a shift or of growing spread. Treat it as you would a near miss: look before a part goes out of the limits.

Rule 6: four out of five points beyond 1 sigma

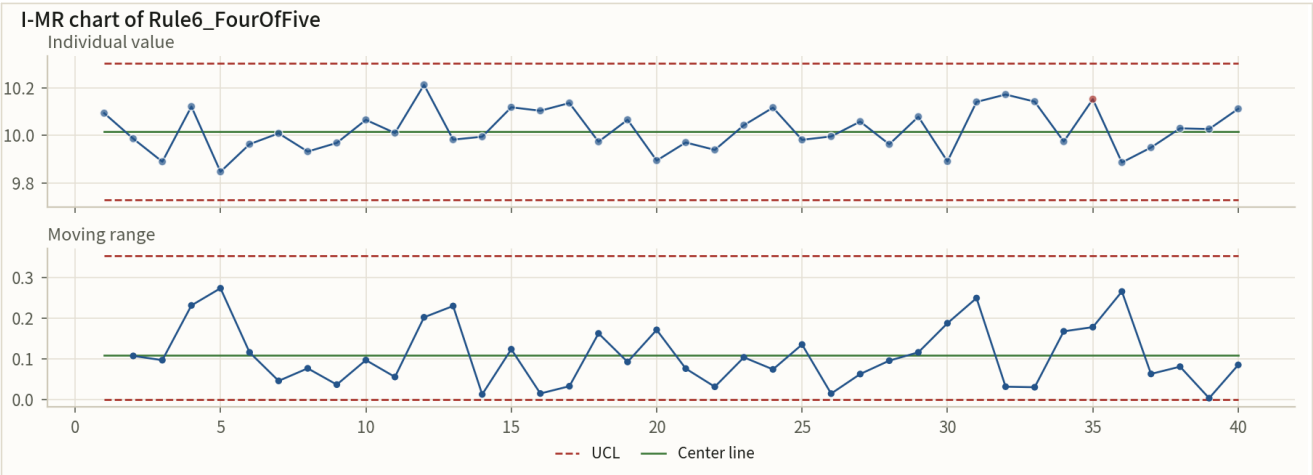


Figure 10.7 Observations 31, 32, 33 and 35 sit in the upper zones; Rule 6 fires at 35.

What it usually means

A moderate shift of the process level, too small for Rule 1 to see.

Rule 7: fifteen points hugging the center

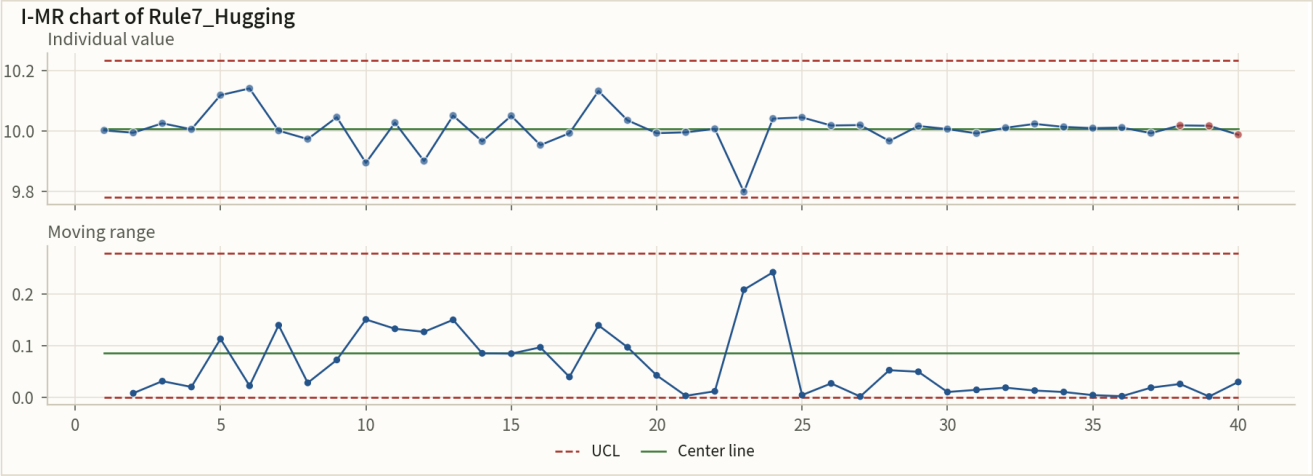


Figure 10.8 From observation 26 the readings barely move, all within 1 sigma. Rule 7 fires at 38 to 40, once fifteen have been seen.

What it usually means

Too good to be true. Usually the limits are wrong (calculated from mixed data), the gauge resolution is too coarse, readings are being rounded or written down from memory, or the process genuinely improved and the limits need recalculating.

Rule 8: eight points avoiding the center

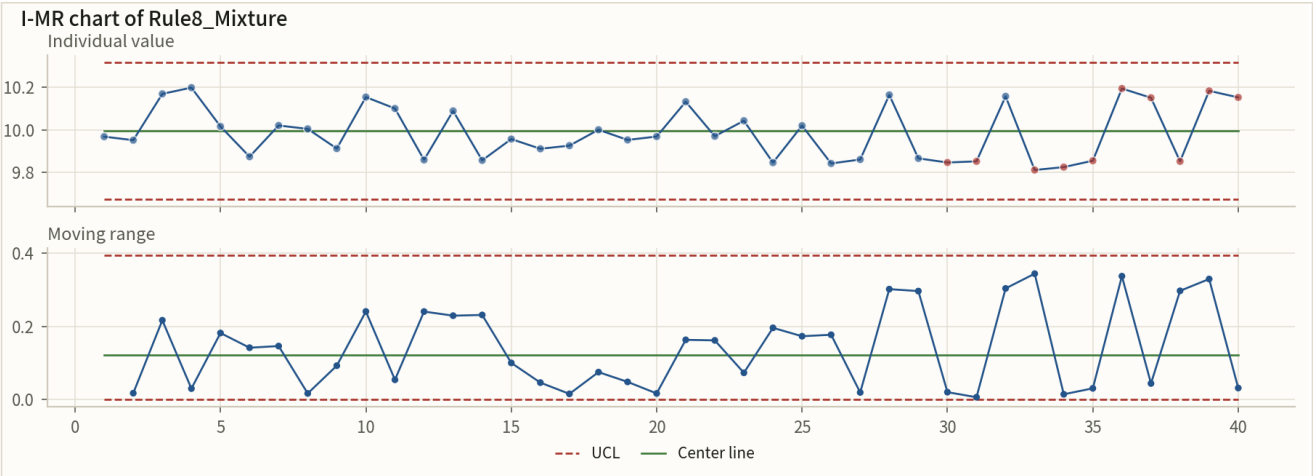


Figure 10.9 From observation 26 the readings jump between high and low with none near the center. Rule 8 fires from observation 33; Rule 6 flags the points far out.

What it usually means

A mixture of two different processes: two machines with different settings, two material suppliers. Stratify the data by machine or lot and chart each separately.

A summary to pin by the machine

Rule	Pattern	First suspects
1	One point beyond 3 sigma	Breakage, setup error, measurement error
2	9 in a row on one side	New lot, changed setting, new operator or gauge
3	6 in a row rising or falling	Wear, warming, drying, loading
4	14 alternating	Two alternating sources, over-adjustment
5	2 of 3 beyond 2 sigma	Early shift or more spread
6	4 of 5 beyond 1 sigma	Moderate shift
7	15 within 1 sigma	Wrong limits, rounding, poor gauge resolution
8	8 avoiding the center	Mixture of two processes

Watch out

The rules assume the readings are independent and roughly normal. Strongly skewed data (times, flatness) and readings that naturally follow each other (a slowly cooling oven measured every second) trigger the pattern rules more often. Chart such data with Rule 1 only, chart averages, or use EWMA (Chapter 13).

CHAPTER 11

Limits that last: baselines, phases, locking and causes

Limits recalculated from all the data every time move with the process and can hide a change. Real SPC fixes the limits from a good period and judges new data against them.

The four ways to set limits

Choice	Use it when	What happens
All the data	You are looking back at history for the first time.	Limits from every reading. Good for exploring, poor for monitoring.
A baseline (the first points)	You know the first N readings were a good, normal period.	Limits from those N readings, applied to all of them. Later points are judged against how the process ran then.
Saved (locked) limits	Day-to-day monitoring.	Limits saved earlier for this column are used; new readings never move them.
Each phase separately	Comparing before and after a deliberate change.	A column such as Before / After splits the chart; each phase gets its own limits.

Before and after an improvement

The teaching dataset below is a bore diameter measured 40 times before and 40 times after a fixture was replaced. The column **Phase** says which is which.

1. **SPC > Control Charts > I-MR**.
2. Measurement **Bore_mm**; Control limits from **Each phase separately**; Phase column **Phase**.

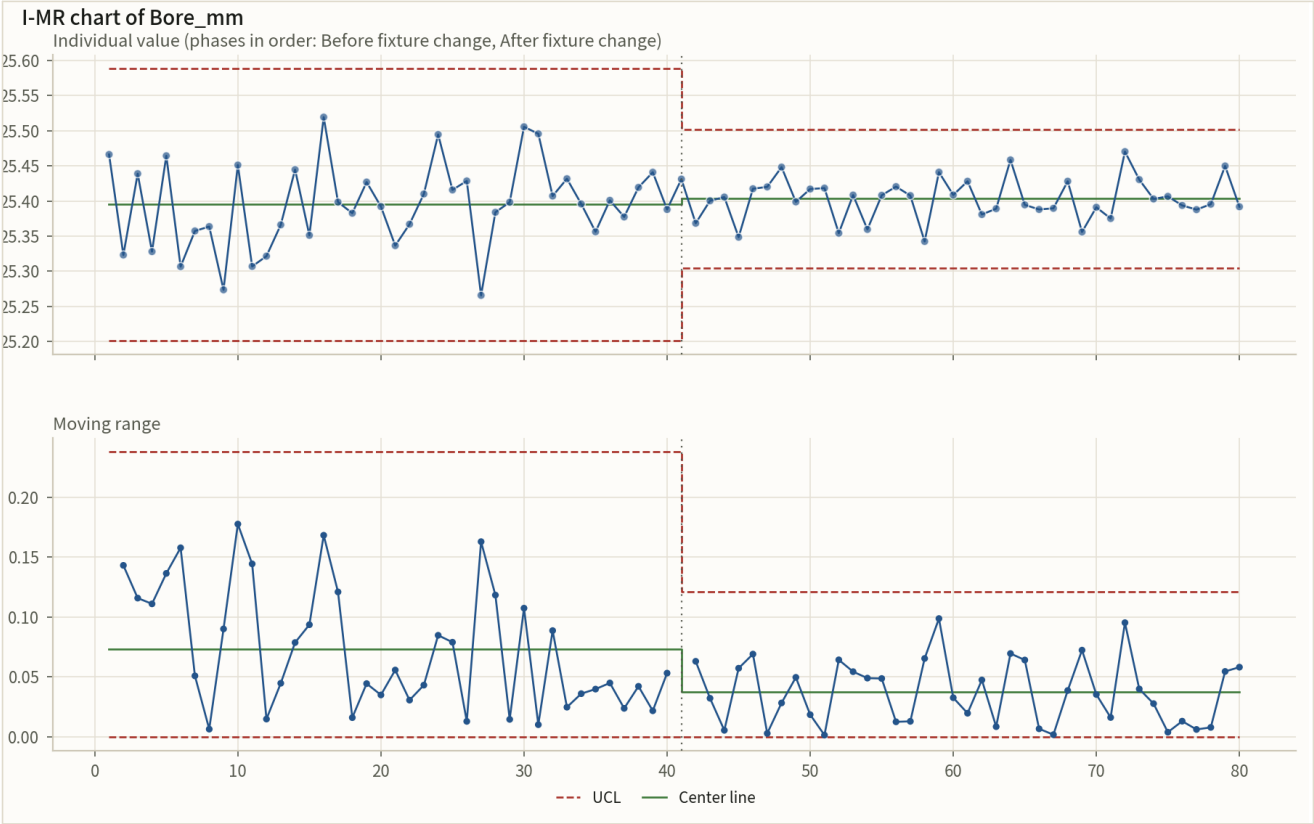


Figure 11.1 Each phase gets its own center line and limits. The dotted line marks the change.

The center barely moved (25.395 to 25.403 mm) but the limits closed in from ±0.194 to ±0.099 mm: the new fixture cut the process spread roughly in half. The moving range chart shows the same thing. This picture, with the table of limits for each phase that comes with it, is the evidence an improvement project (8D, DMAIC Control phase) needs.

Fix the limits and lock them

1. Run the chart with **A baseline** and the number of good readings, and tick **Lock these limits**.
2. From now on, choose **Saved (locked) limits** when you chart that column. New readings are judged against the saved limits and cannot drag them along.
3. **SPC > Limits and Causes > Saved Limits** lists what is saved for each column; **Unlock Limits** removes them when the process is deliberately changed.

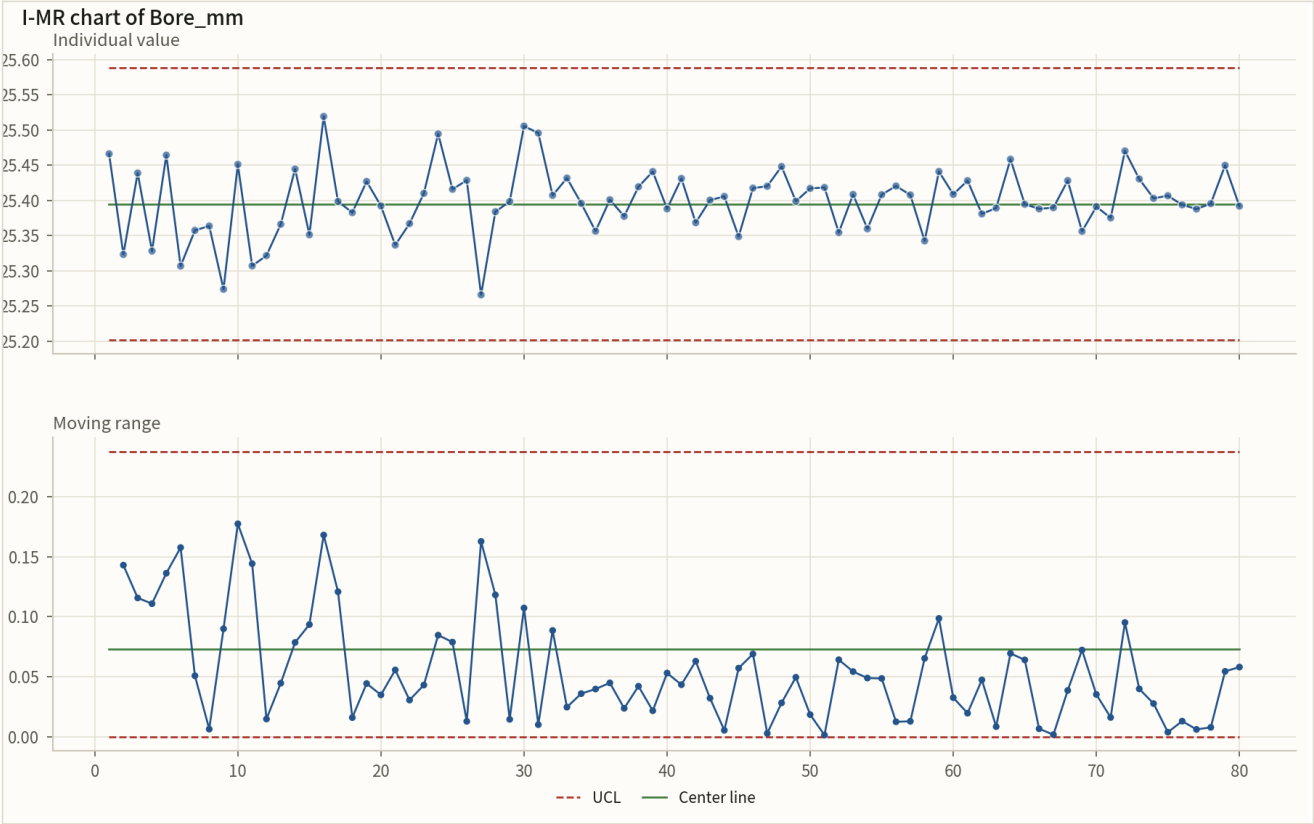


Figure 11.2 The first 40 readings (the old fixture) as the baseline, applied to all 80. The new fixture's readings sit well inside the old limits: the change helped.

Saved Control Limits

Charts run with "Saved (locked) limits" use these instead of recalculating from the data.

Column	Chart	Center line	Sigma	Subgroup size	From data	
Bore_mm	I-MR	25.395	0.064522	1	phase_demo	202

Figure 11.3 Saved Limits shows the locked limits for each column.

> **Watch out**
> Recalculate locked limits only for a reason you can name: a deliberate process change, a new machine, an improvement confirmed by a phase chart. Never recalculate just to make the red points go away.

Recording causes

Where: SPC > Limits and Causes > Record a Cause

When a point is investigated, record what was found. The form asks for the column, the point numbers as shown on the chart (for example **19** or **144, 163**), the cause and the corrective action. The note is saved with your name and the time, written to the audit trail, and listed in the cause log.

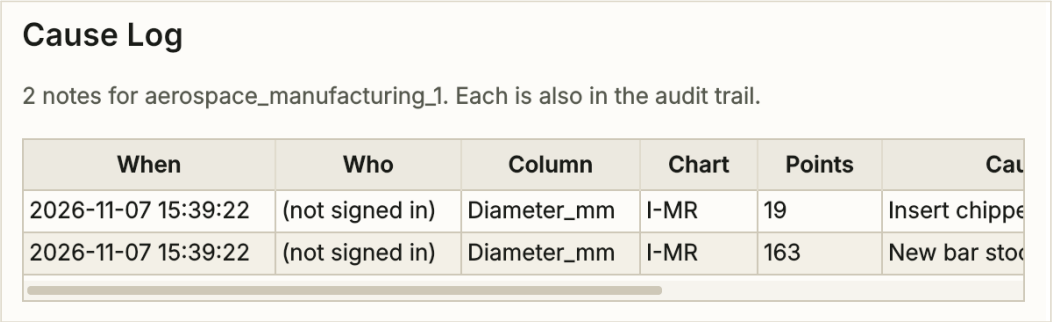


Figure 11.4 SPC > Limits and Causes > Cause Log lists every note for the open data.

Charts then say how many flagged points still have **no recorded cause**, which turns the chart into a to-do list.

> **In real life**
> Auditors for ISO 9001, IATF 16949 and AS9100 look for exactly this: evidence that out-of-control conditions were noticed, investigated and acted on. A chart full of red points with no notes is a finding; the same chart with a cause log is a strength.

CHAPTER 12

Attribute charts: counting defects and defectives

When the result is pass or fail, or a count of flaws, chart the counts.

Defectives or defects?

	Defectives	Defects
What is counted	Items that fail (a bad part is bad, however many flaws it has)	Flaws (one part can have three scratches)
Same sample size each time	**np chart**: number defective	**c chart**: number of defects
Sample size varies	**p chart**: proportion defective	**u chart**: defects per unit

Where: SPC > Control Charts > p, np, c, u

The example **Inspection lots** has 40 lots with the sample size, the number of defective pieces and the number of defects found.

p chart

1. Choose **p Chart**.
2. Choose the column of defectives (**Defectives**).
3. Choose the sample size column (**Sample_Size**).

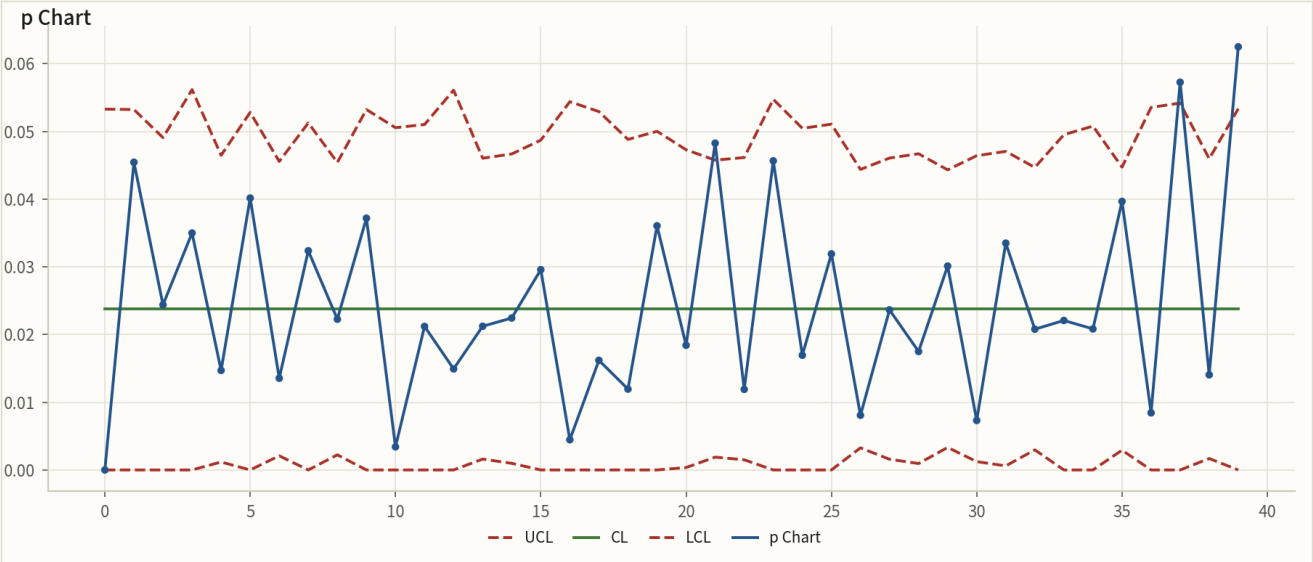


Figure 12.1 p chart: the proportion defective in each lot. The limits step in and out because each lot's sample size is different.

Larger samples give tighter limits, because a proportion from a big sample is more certain. That is why the limits are jagged. A lot is only judged against its own limits.

np chart

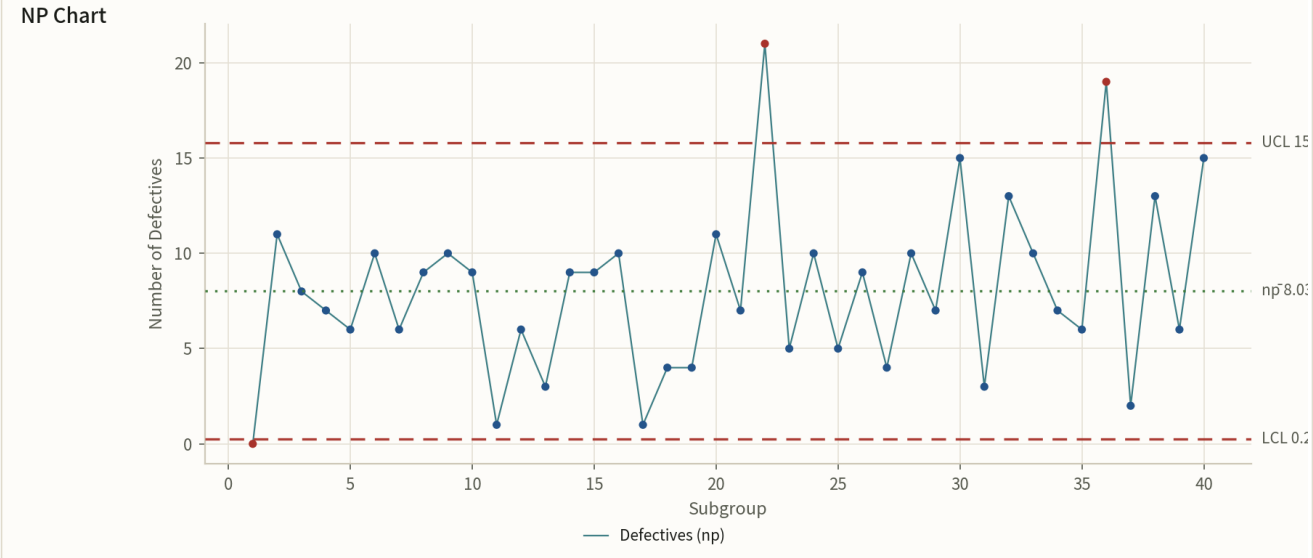


Figure 12.2 np chart with a fixed sample of 50: the count of defective items per sample. Three samples are above the UCL of 15.8.

u and c charts

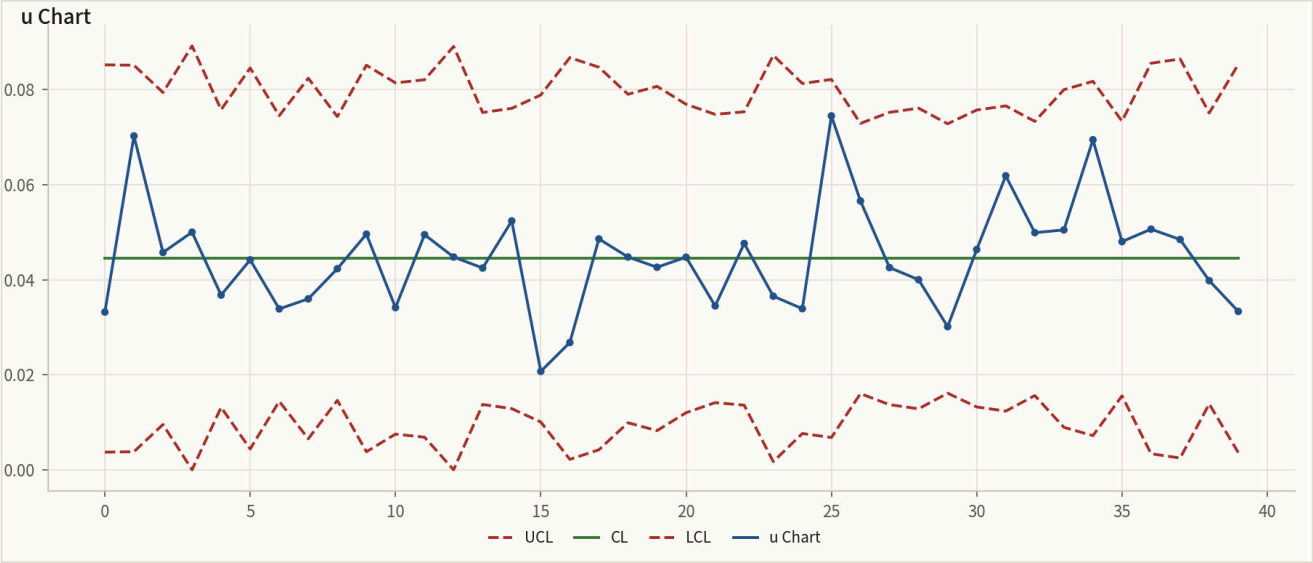


Figure 12.3 u chart: defects per unit inspected.

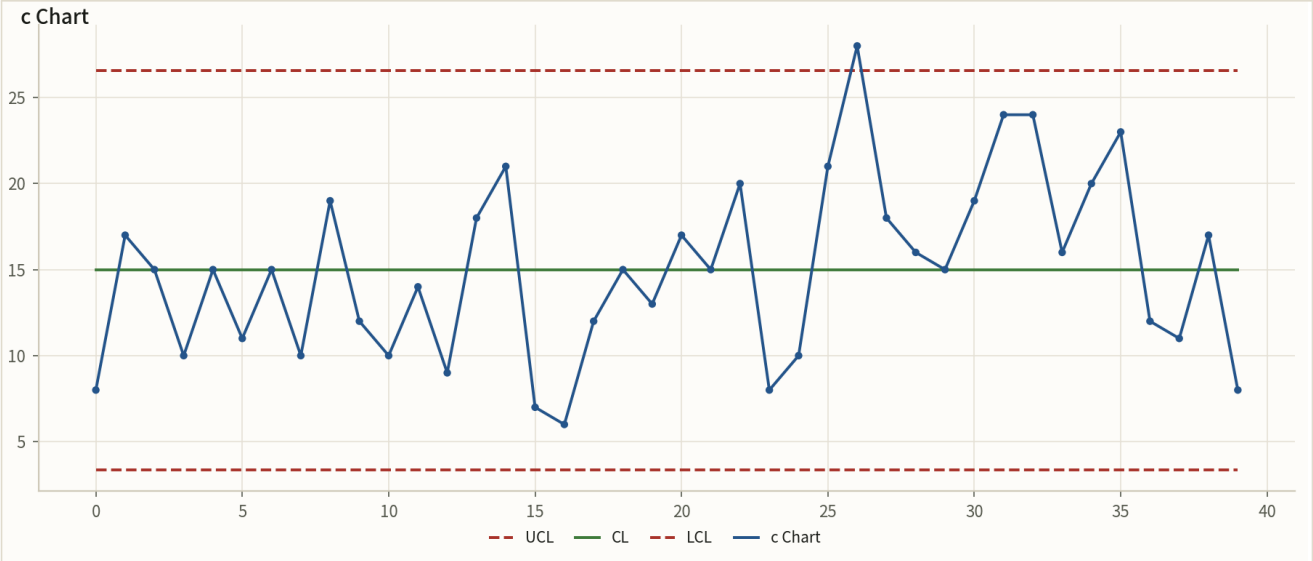


Figure 12.4 c chart: defects per lot, for a fixed amount inspected.

> **In real life**
> A paint line counts surface defects on each panel. The c chart shows a jump every Monday morning: dust from the weekend shutdown settles in the booth. Running the booth fans for an hour before the first panel removes the Monday signal.

Laney p': when there are thousands per sample

Where: SPC > Control Charts > Laney p'

With very large samples (thousands of items per day) an ordinary p chart's limits become so narrow that almost every point is out of control, because real day-to-day differences are bigger than the binomial model expects. Laney's p' chart measures that extra day-to-day variation and widens the limits by it.

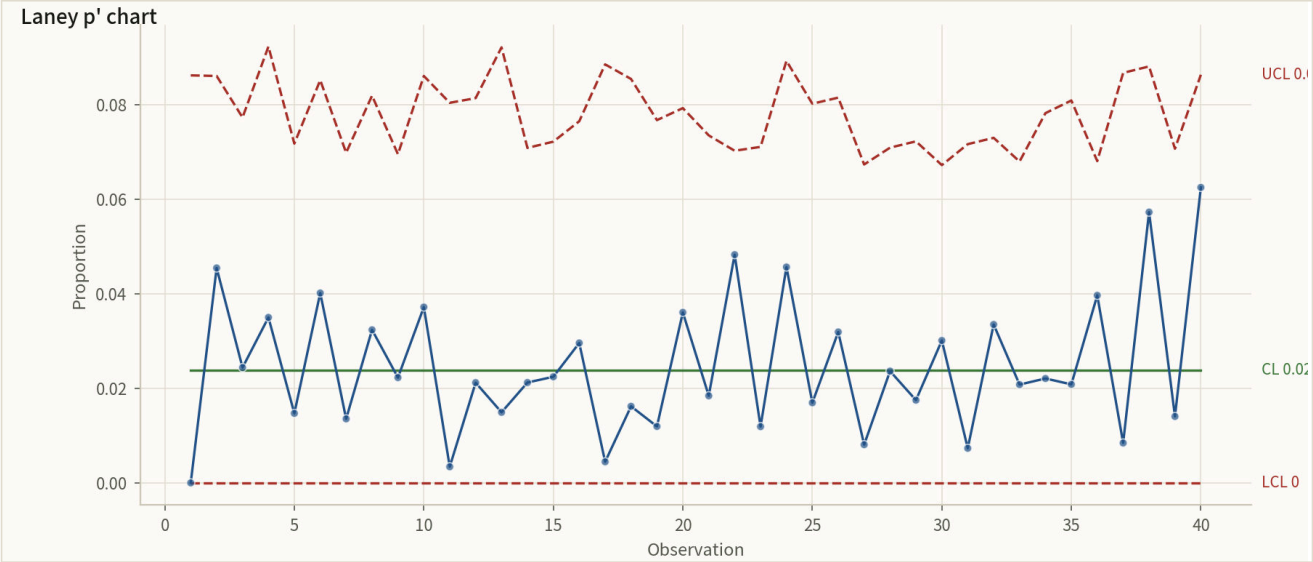


Figure 12.5 Laney p' chart. Sigma Z = 2.1: days differ twice as much as the simple model expects, so the limits are about twice as wide. No point is out.

G chart: rare events

Where: SPC > Control Charts > G Chart

When events are rare (injuries, spills, field failures), counting them per week gives mostly zeros. Instead, chart **how many days (or units) passed between events**.

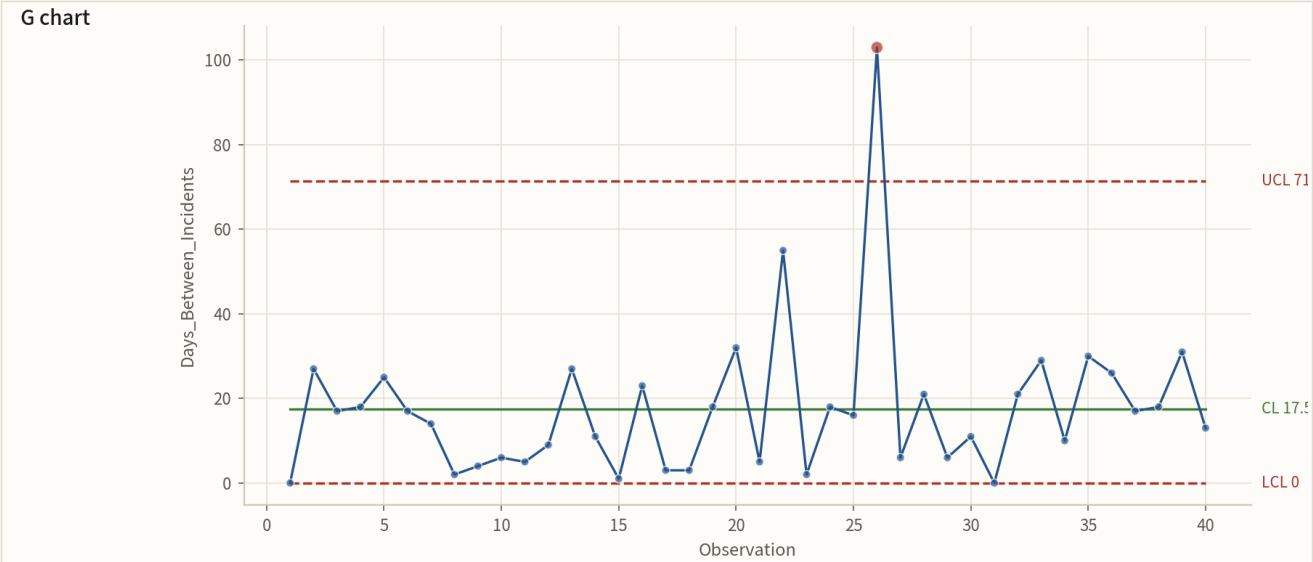


Figure 12.6 G chart of days between incidents. Mean 17.5 days; one interval of about 100 days is above the upper limit of 71.5.

Read a G chart upside down from the others: **high points are good news** (events became rarer) and a run of low points means they are getting more frequent.

CHAPTER 13

Special-purpose charts

Charts for small shifts, short runs and simple go/stop rules.

EWMA and CUSUM: catching small shifts

Where: SPC > Control Charts > EWMA / CUSUM

An I-MR chart needs a shift of about 2 to 3 sigma to flag a point quickly. Many costly problems are smaller: a filler that starts giving away 1 g per pack, a temperature 1 degree high. EWMA and CUSUM charts **remember** recent readings, so a small shift that lasts builds up until it is caught.

The teaching dataset is a filling machine with a target of 500 g. After reading 40 a worn valve seat makes it fill 1.2 g heavy, a little over one sigma. First, the I-MR chart:

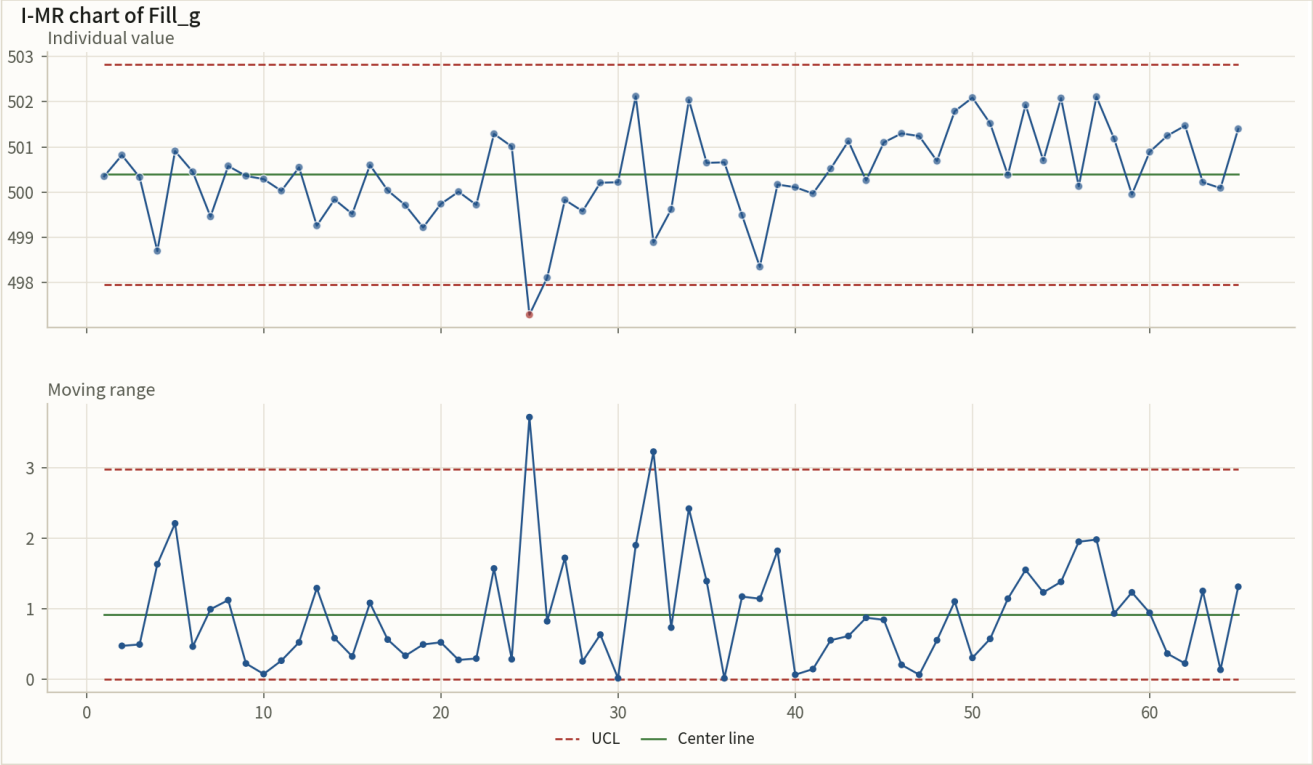


Figure 13.1 I-MR chart of the fill weights. It flags one random low reading (25) and misses the shift completely.

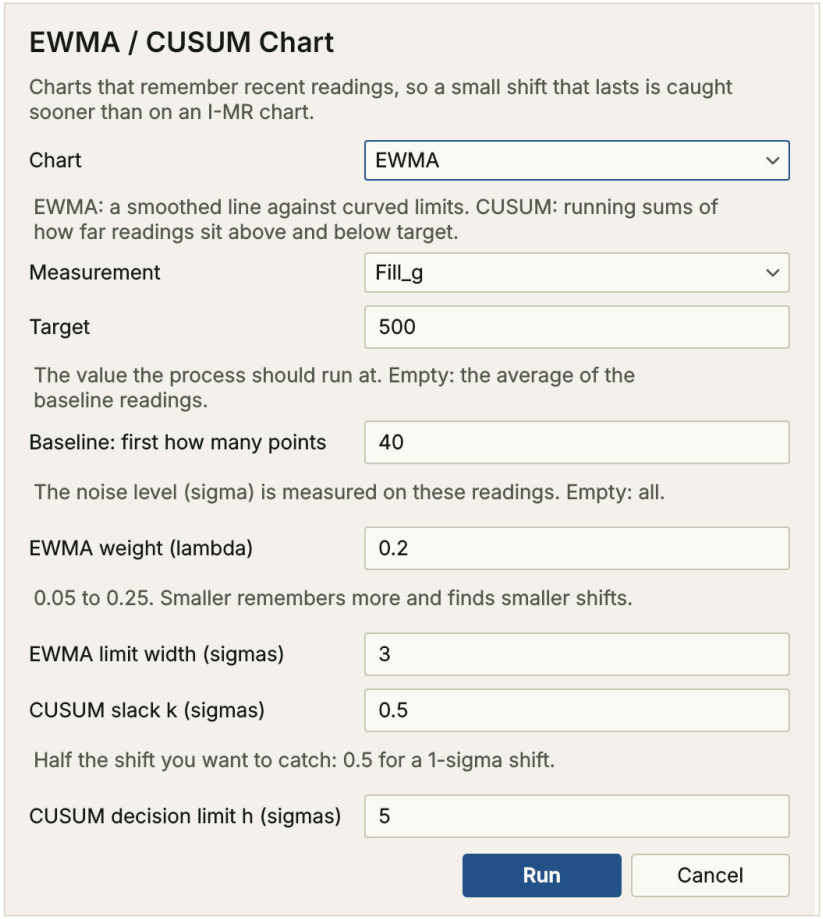


Figure 13.2 The EWMA / CUSUM questions. Give the target, and say which readings were normal (here the first 40).

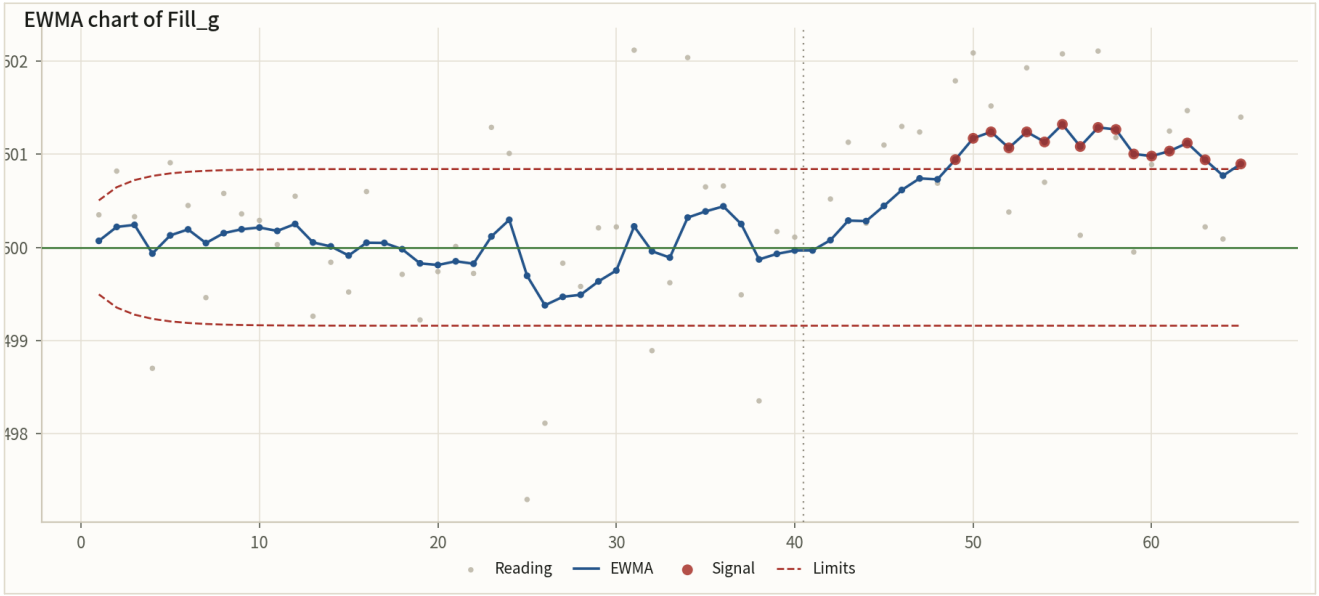


Figure 13.3 EWMA chart. The blue line is the smoothed average, the gray dots the raw readings. The EWMA crosses its upper limit at reading 49 and stays above.

Each EWMA point is 20% of the new reading plus 80% of the previous EWMA (weight 0.2). Random ups and downs cancel; a lasting shift accumulates. The limits start narrow and widen to a steady width after about ten readings. Because the line is smoothed, it reacts **after** the change: the result tells you to look a little before the first signal.

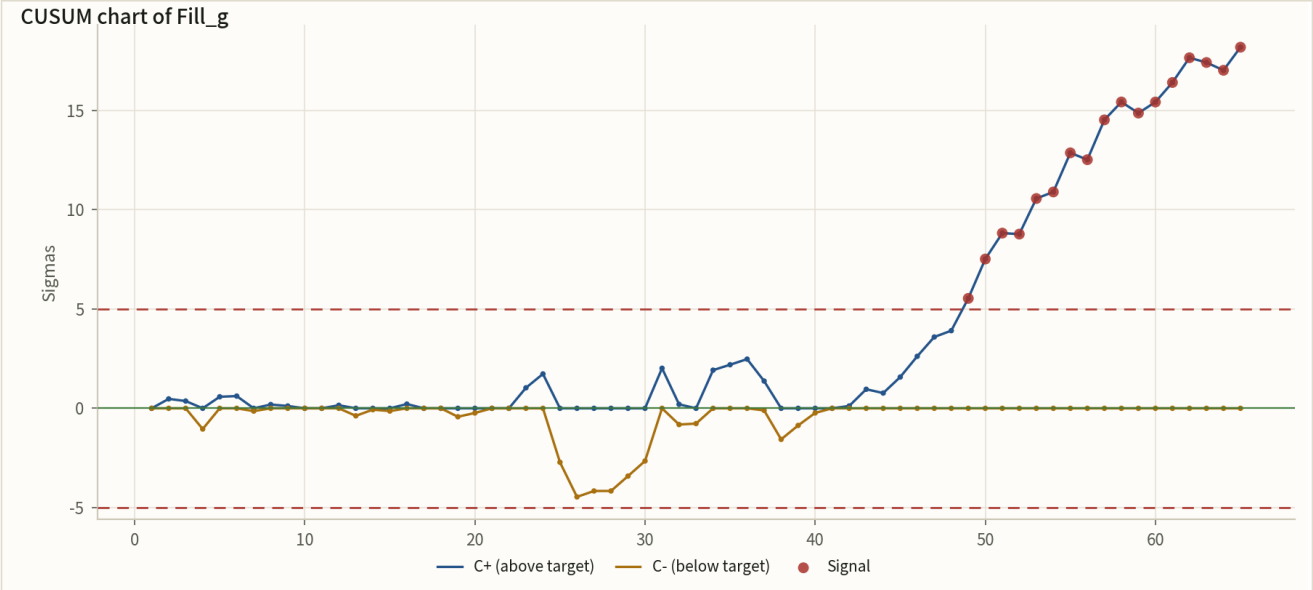


Figure 13.4 CUSUM chart. C+ (blue) adds up how far readings sit above target; it climbs from reading 42 and crosses the decision limit of 5 sigma at reading 49.

The CUSUM keeps two running sums in sigma units. **C+** grows when readings are above target by more than the slack k (half a sigma) and falls back to zero otherwise; **C-** does the same below target and is drawn downward. When either passes **h** (5 sigma) a sustained shift is signaled. The result names where the sum started climbing, the most likely start of the change: here reading 42, one reading after the real change at 41.

Setting	Usual value	Effect
EWMA weight (lambda)	0.2	Smaller (0.05 to 0.1) finds smaller shifts but reacts more slowly to big ones.
EWMA limit width	3 sigma	2.7 with a small weight gives the same false-alarm rate as an I-MR chart.
CUSUM slack k	0.5 sigma	Half the shift you want to catch.
CUSUM decision limit h	5 sigma	Lower catches sooner with more false alarms.

In real life

Give-away on a filler costs money on every pack. A 1 g overfill on 500 g packs at 60 packs a minute is over 28 kg of product a shift. The I-MR chart above would never have flagged it; the CUSUM found it nine packs after the valve seat wore.

Short-run charts: many part numbers, one process

Where: SPC > Control Charts > Short Run

A job shop may run a part number for 12 pieces, switch to another, and come back next week. No part number has enough readings for its own chart. The short-run chart plots every part on one chart as its **deviation from its own nominal** (DNOM), so the process (the machine, the setup method) is charted rather than the part.

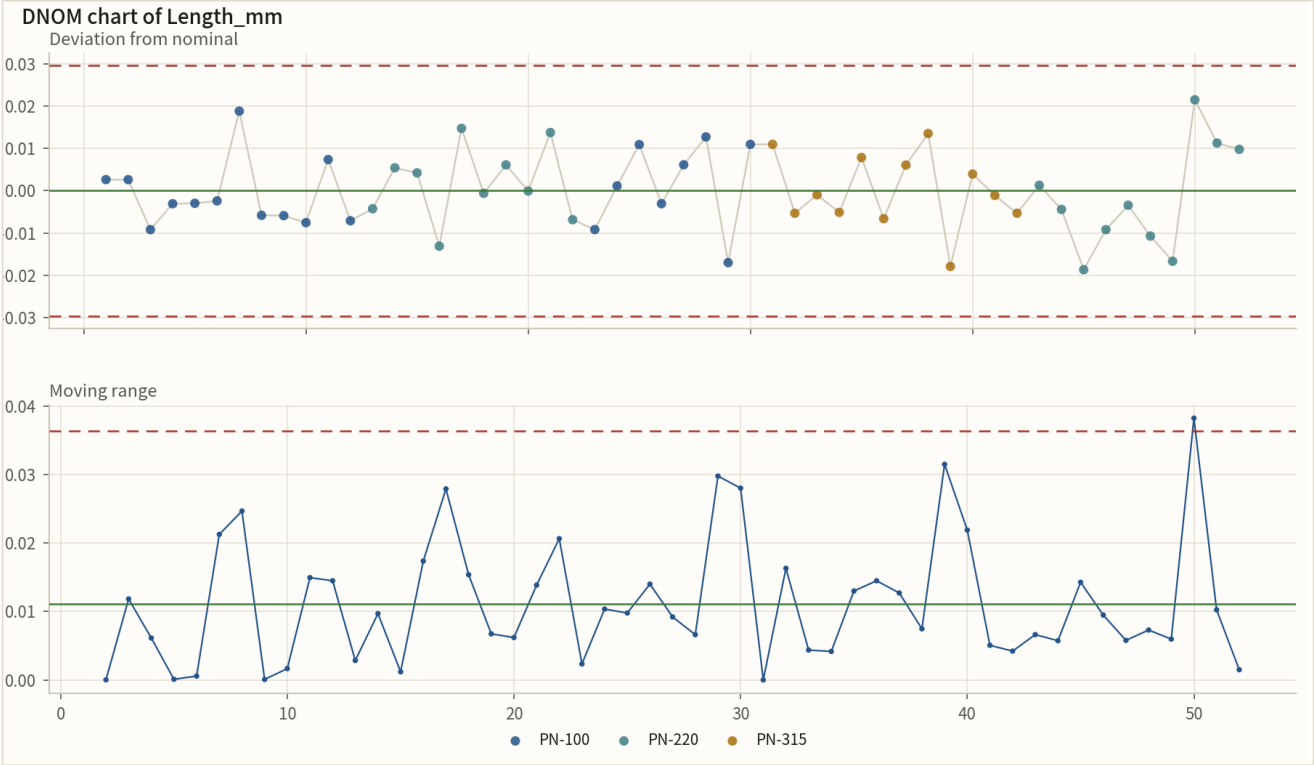


Figure 13.5 DNOM chart of three part numbers turned on one lathe, colored by part, in production order. All 52 readings are inside the limits: the lathe is stable across part numbers.

Use **DNOM** when the parts vary by about the same amount, **Z-MR** when they do not (each part's deviation is divided by its own sigma). The nominal can be each part's average or the spec target set for the measurement.

Pre-control: green, yellow, red

Where: SPC > Control Charts > Pre-Control

Pre-control needs no statistics on the floor. The tolerance is split into zones: the middle half is **green**, the outer quarters **yellow**, outside the specs **red**.

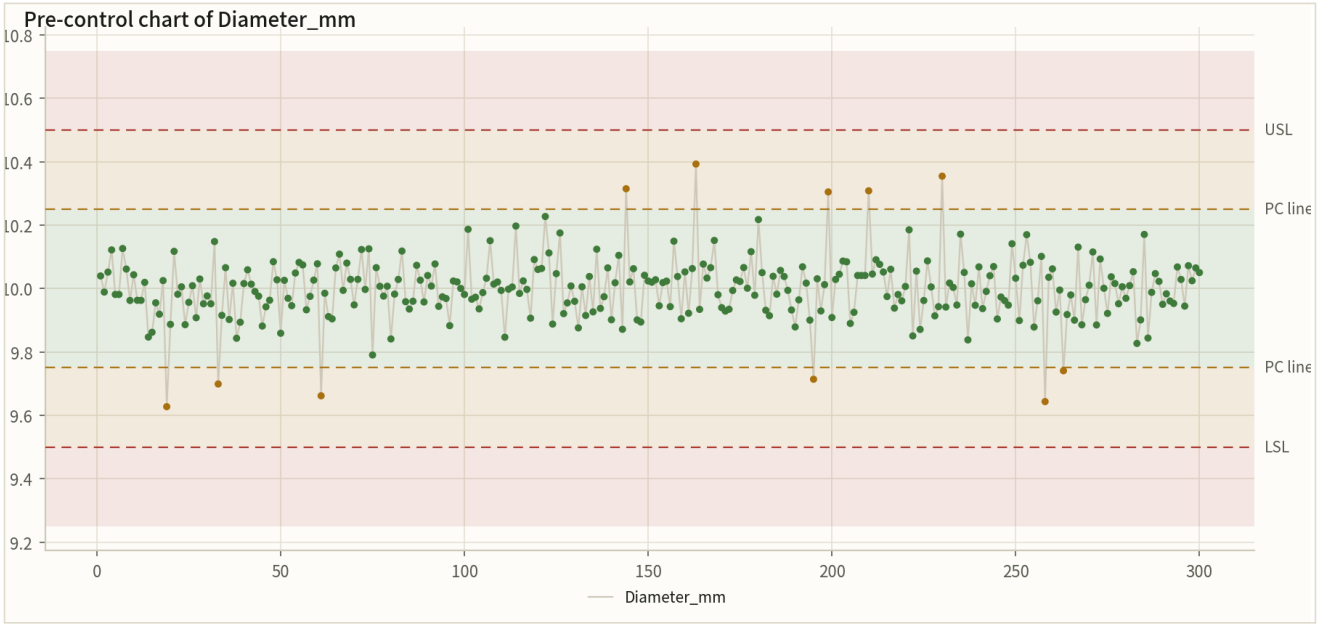


Figure 13.6 Pre-control chart of the diameters with specs 9.5 to 10.5: 289 green, 11 yellow, none red, no stops.

- **To start (qualify):** five greens in a row.
- **While running:** one red, or two yellows in a row, means stop, adjust, and qualify again.
- Otherwise keep running.

> **Watch out**
> Pre-control judges parts against the tolerance, not the process against itself. It is a good setup-approval tool and simple for operators, but it will not show trends or shifts inside the green zone. Use it alongside, not instead of, a control chart for important characteristics.

CHAPTER 14

SPC on the shop floor

Entering readings as parts are made, seeing every characteristic at a glance, and getting warned the moment something goes wrong.

Data entry

Where: SPC > Shop Floor > Data Entry

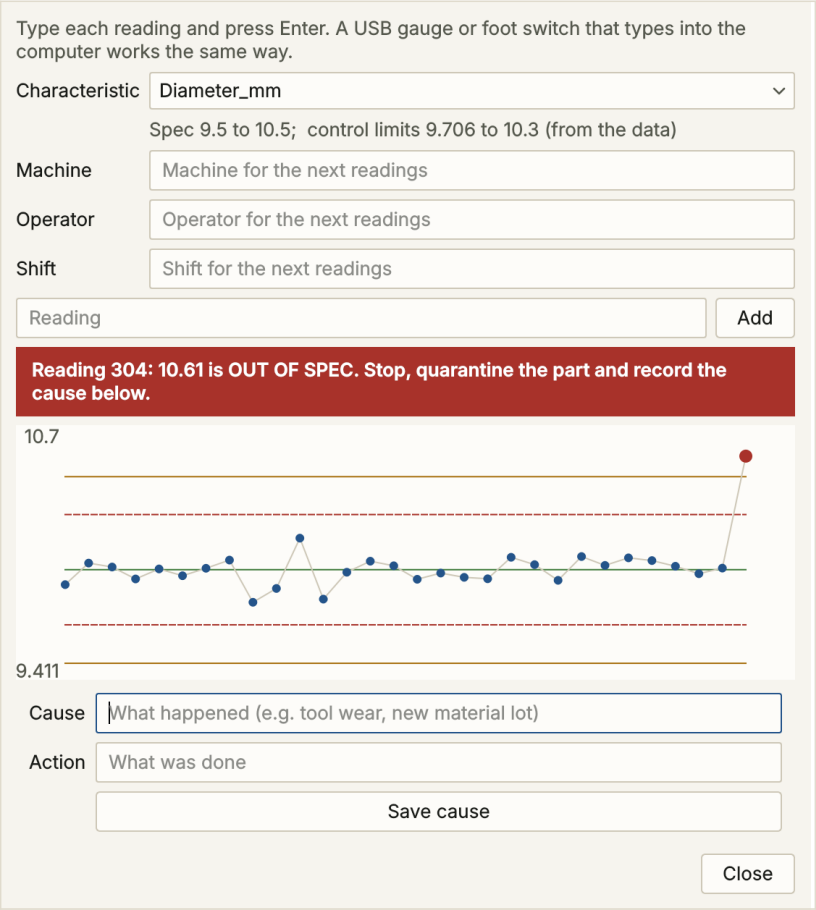


Figure 14.1 Data entry: the fourth reading, 10.61, is out of spec, so the window says so and opens the cause boxes.

1. Choose the **Characteristic**. The window shows its spec and control limits.
2. Optionally fill in **Machine**, **Operator** and **Shift**: they are written to every reading that follows, so the data carries its context.
3. Type each reading and press **Enter** (or click **Add**). A USB gauge or foot switch that types into the computer works the same way.

4. Each reading is added to the data sheet and judged at once: in spec and in control (green), a rule broken (amber), or **out of spec** (red, with the instruction to stop and quarantine the part).
5. When a reading is flagged, type the **Cause** and the **Action** taken and click **Save cause**. It goes into the cause log against that reading.

In real life

An operator measures every tenth bore with a digital gauge plugged into the USB port. Each press of the gauge's button enters a reading; the run chart in the window updates as she works. When a reading turns red she stops the machine, sets the part aside and types insert chipped / replaced insert. The quality engineer sees the note in the cause log the next morning.

The SPC dashboard

Where: SPC > Shop Floor > Dashboard

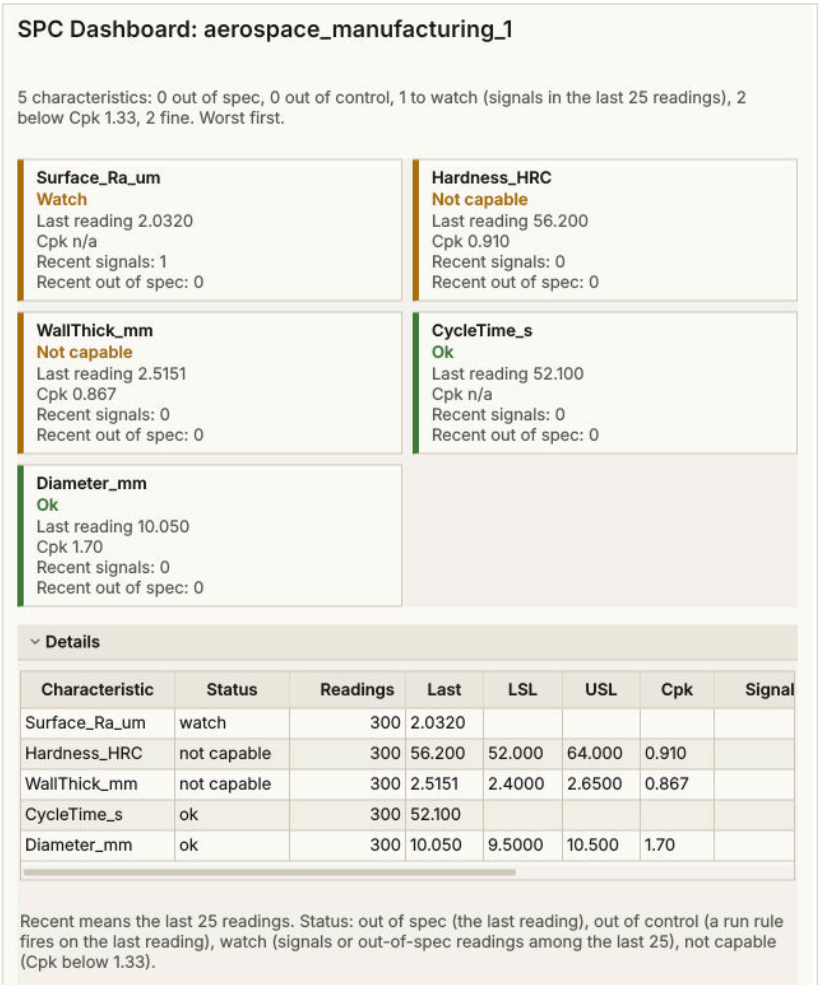


Figure 14.2 The dashboard for the aerospace data: one card per characteristic, worst first.

Each card shows the last reading, the Cpk, the number of rule signals and out-of-spec readings among the **last 25**, and a status:

Status	Meaning	Do
Out of spec	The last reading is outside the tolerance.	Stop and contain.
Out of control	A rule fires on the last reading.	Find what changed.
Watch	Signals or out-of-spec readings among the last 25.	Check the chart; confirm the cause was dealt with.
Not capable	Cpk below 1.33.	Plan an improvement (Part III).
Ok	None of the above.	Nothing.

Here **Surface_Ra_um** is on watch (one recent signal), **Hardness_HRC** and **WallThick_mm** are stable but not capable (Cpk 0.91 and 0.87), and **Diameter_mm** is fine with Cpk 1.70.

SPC alarms

Where: SPC > Chart Rules > SPC Alarms

Checks every characteristic with spec or control limits against the alarm rules in one go and lists what needs attention. With **Live Data** (Chapter 4) watching a gauge's output file, the checks run again whenever new rows arrive.

> **Tip**
> Put the dashboard on a screen at the cell. A supervisor walking past sees the worst characteristic first, and an amber or red card is a conversation starter before it is a scrap report.

PART III

Capability and measurement

Can the process make good parts, and can the gauge tell good from bad? Cp, Cpk, Pp and Ppk explained with pictures, non-normal data, and the three measurement-system studies.

CHAPTER 15

Process capability: Cp, Cpk, Pp, Ppk

Capability compares the width of the process with the width of the tolerance, and asks whether the process is centered in it.

The idea in one picture

Imagine parking a car in a garage. The garage door is the **tolerance** (USL minus LSL). The car is the **process spread**, six standard deviations wide, which holds 99.73% of the parts of a normal process. Two questions matter:

- **Is the car narrower than the door?** That is **Cp**: tolerance width divided by process width. Cp = 1 means a perfect fit with no room to spare; Cp = 2 means the door is twice as wide as the car.
- **Is the car in the middle of the door?** That is what **Cpk** adds: it measures the room on the tighter side only. A wide door is no use if you drive in against one side.

Index	Formula	Reads as
Cp	(USL - LSL) / (6 sigma)	Potential: how good it could be if centered
Cpk	min(USL - mean, mean - LSL) / (3 sigma)	Actual: how good it is now
Cpu, Cpl	One side each of Cpk	Which side is tighter
Pp, Ppk	Same, using the overall standard deviation	Long-term performance including drift between subgroups

Within and overall: Cpk or Ppk?

Two standard deviations are worked out from the same data:

- **Within (short-term)** sigma comes from the moving ranges or subgroup ranges: the spread of parts made close together. **Cp and Cpk** use it. They describe what the process can do when nothing changes.
- **Overall (long-term)** sigma is the plain standard deviation of every reading. It includes shifts and drifts between subgroups. **Pp and Ppk** use it. They describe what the customer actually received.

For a stable process the two are almost equal, and so are Cpk and Ppk. **A big gap between Cpk and Ppk is a sign the process is not stable**: it moves around more over time than it does from part to part.

How big is big enough?

Cpk	Usual reading	Defects (centered, normal)
below 1.00	Not capable: defects are being made	more than 2,700 per million
1.00 to 1.33	Marginal: needs close control and 100% checks may be required	64 to 2,700 per million
1.33	The usual minimum for production	about 64 per million
1.67	The usual minimum for safety or critical characteristics and new processes (PPAP)	under 1 per million
2.00	Six Sigma level	about 0.002 per million

Your customer's quality manual is the authority. Automotive (IATF 16949) and aerospace (AS9145) customers typically require Ppk ≥ 1.67 for initial approval and Cpk ≥ 1.33 in ongoing production.

Four processes, four stories

The four teaching datasets below each have 125 parts against the same tolerance, 9.94 to 10.06 mm. Each was run with:

Where: Home > Analyze > Capability

1. Centered and capable

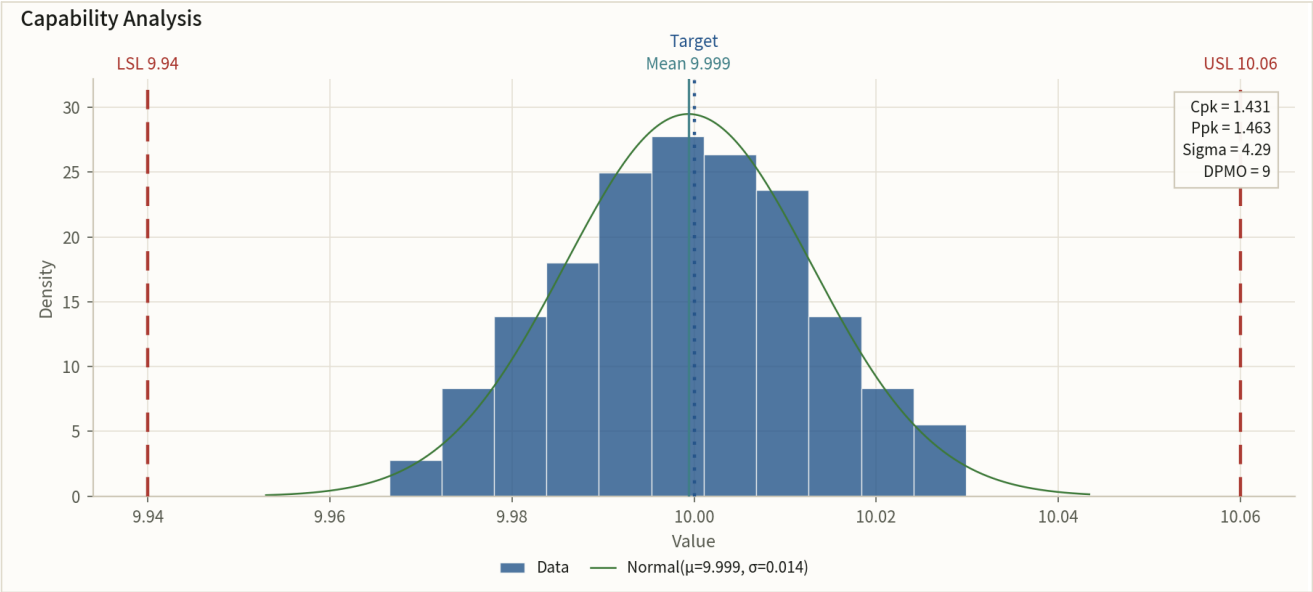


Figure 15.1 Centered on 9.9994, spread well inside the limits. Cp 1.4458, Cpk 1.4307.

The histogram sits in the middle and both tails stop short of the limits. Cp and Cpk are nearly equal, so the process is centered, and at **Cpk 1.4307** it clears the usual 1.33 minimum. The expected defect rate is about 9 per million. This is what a healthy production characteristic looks like.

2. Narrow but off-center

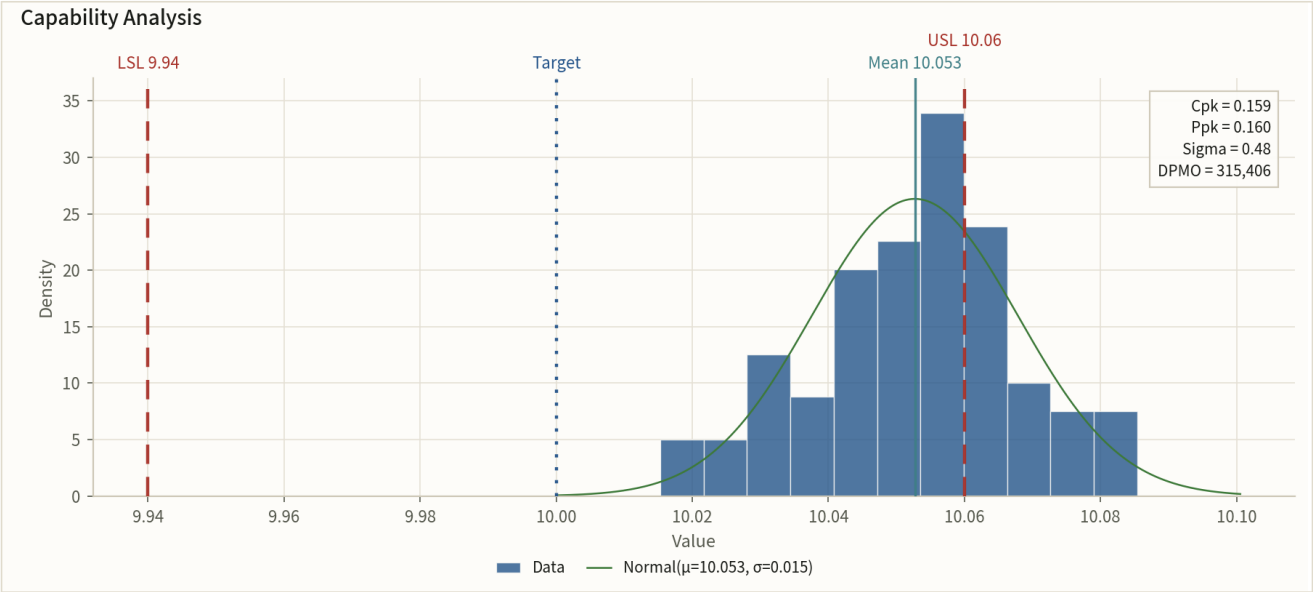


Figure 15.2 The same kind of spread, but the mean is at 10.0527, almost on the USL. Cp 1.3081, Cpk 0.1588.

Cp is 1.3081 but Cpk is 0.1588. The car fits the door but is parked against one side: about 32% of the parts are over the upper limit. The program says so in plain words: Cp >> Cpk: process is off-center. Adjust the mean toward the target to recover capability. This is the cheapest problem in quality to fix: an offset, not a new machine.

3. Centered but too wide

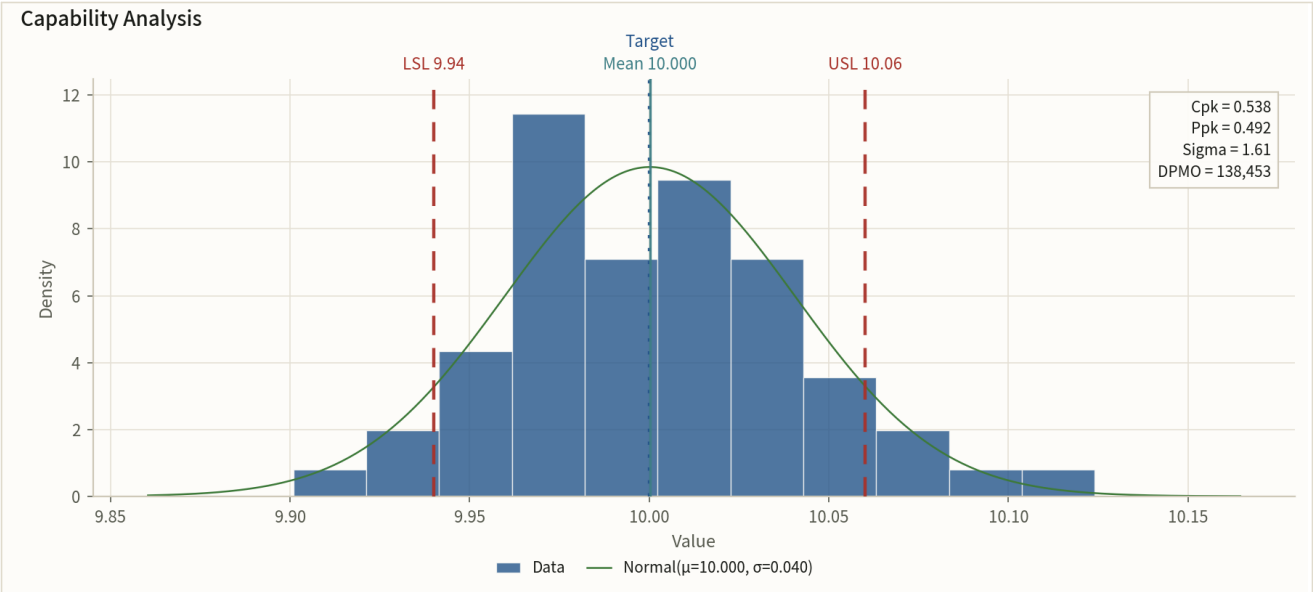


Figure 15.3 Centered on target, but the spread is far wider than the tolerance. Cp 0.5391, Cpk 0.5377.

Both indices are low and about equal: the process is centered, it is just too variable. About 14% of parts fall outside the limits. Adjusting the setting would not help. The spread itself must come down (tooling, fixturing,

material, method), or the tolerance must be reviewed with the designer, or every part must be inspected until it does.

4. Capable on paper, not stable

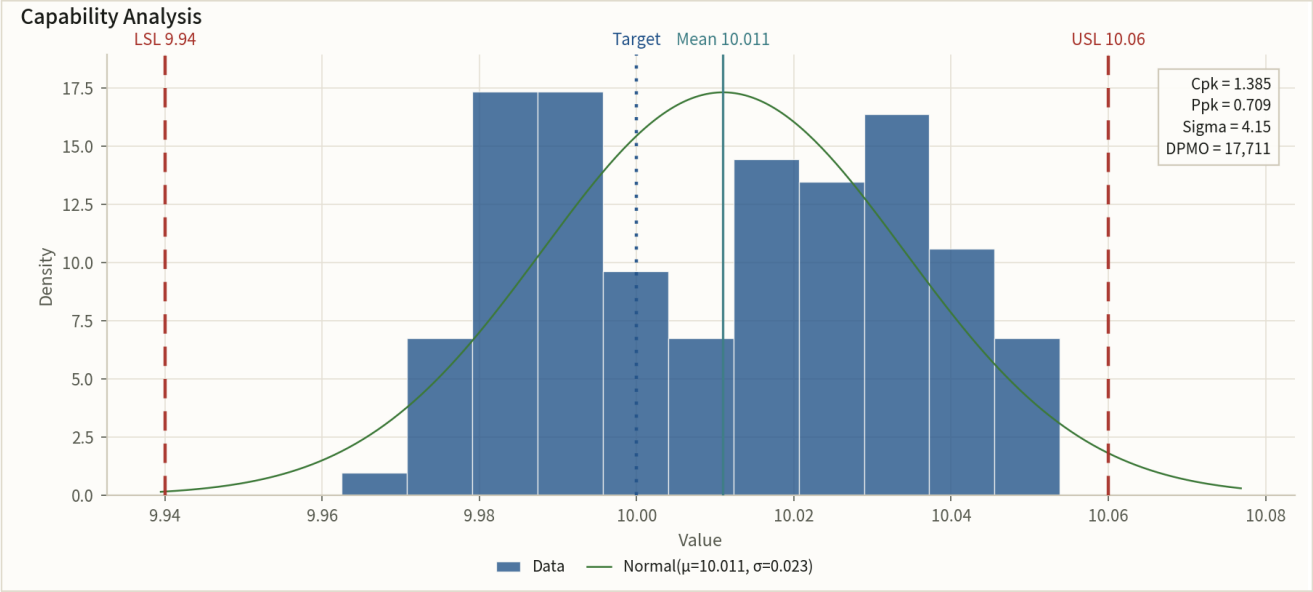
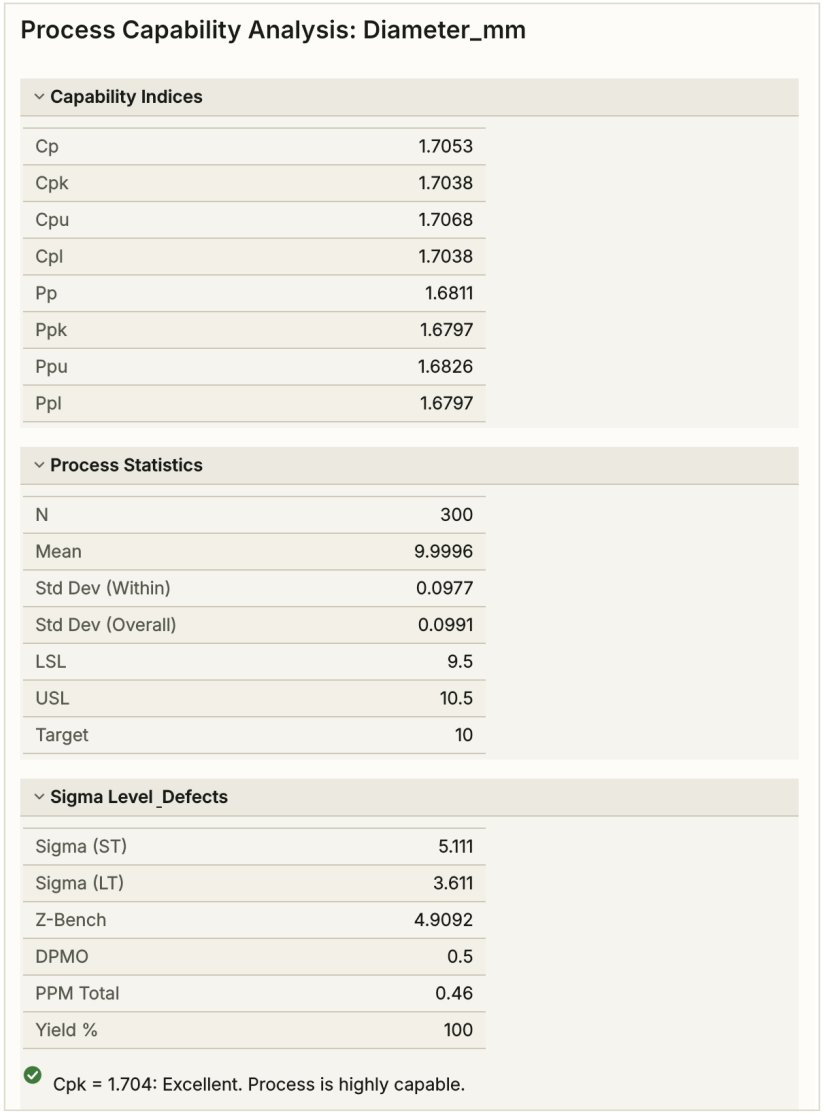


Figure 15.4 Cpk 1.3848 looks acceptable, but the process ran at two different levels. Pp 0.8684, Ppk 0.7092.

The **within** sigma (0.0118) only sees part-to-part noise, so Cpk says 1.3848. The **overall** sigma (0.023) sees the whole picture, and Ppk is only 0.7092. The histogram's two humps tell the story: the process ran near 9.99 for 60 parts and then jumped to 10.03. Chapter 16 shows the same data in the capability report, where the I chart makes the jump obvious.

> **Watch out**
> Never quote a Cpk without looking at the control chart. A capability index predicts the future only if the process is stable. For an unstable process, report Ppk (what was actually delivered) and fix the instability first.

Every number on the capability page



Process Capability Analysis: Diameter_mm

Capability Indices

Cp	1.7053
Cpk	1.7038
Cpu	1.7068
Cpl	1.7038
Pp	1.6811
Ppk	1.6797
Ppu	1.6826
Ppl	1.6797

Process Statistics

N	300
Mean	9.9996
Std Dev (Within)	0.0977
Std Dev (Overall)	0.0991
LSL	9.5
USL	10.5
Target	10

Sigma Level _Defects

Sigma (ST)	5.111
Sigma (LT)	3.611
Z-Bench	4.9092
DPMO	0.5
PPM Total	0.46
Yield %	100

Cpk = 1.704: Excellent. Process is highly capable.

Figure 15.5 The full capability results for the aerospace diameters (spec 9.5 to 10.5).

Number	Meaning
Cp, Cpk, Cpu, Cpl	Short-term capability (within sigma), as above.
Pp, Ppk, Ppu, Ppl	Long-term performance (overall sigma).
Std Dev (Within) / (Overall)	The two sigmas. Close together here (0.0977 and 0.0991): the process is consistent over time.
Sigma (ST) / (LT)	The sigma level: distance from the mean to the nearer limit in short-term sigmas, and the long-term level (short-term minus the conventional 1.5).
Z-Bench	The total fraction out of spec on both sides, expressed as a single sigma distance.
DPMO / PPM Total	Expected defective parts per million.
Yield %	Expected share of good parts.

Cpk 1.70 and under one defect per million: this characteristic is not where your quality problems are.

How many parts do you need?

Where: Quality > Sampling > Sample Size

A Cpk from 30 parts is a rough estimate; from 125 it is a good one. The sample size calculator's Capability Study (Cpk) option says how many parts are needed to show a Cpk target with confidence.

Sample Size: Capability Study

n	51
cpk_target	1.33
alpha	0.05
power	0.8
note	Boyles (1991) approximation for Cpk estimation precision

Figure 15.6 To demonstrate Cpk 1.33 with 95% confidence and 80% power: 51 parts.

In real life

For a PPAP, customers commonly ask for at least 100 consecutive parts (often 125 = 25 subgroups of 5) from a significant production run. A Cpk from a 30-piece first run is useful for a go / no-go decision on the tooling, but it is not the capability study.

CHAPTER 16

The capability report and non-normal data

Six charts on one page check every assumption behind the numbers, and a different method handles data that is not bell-shaped.

The capability report (sixpack)

Where: Quality > Capability > Capability Report

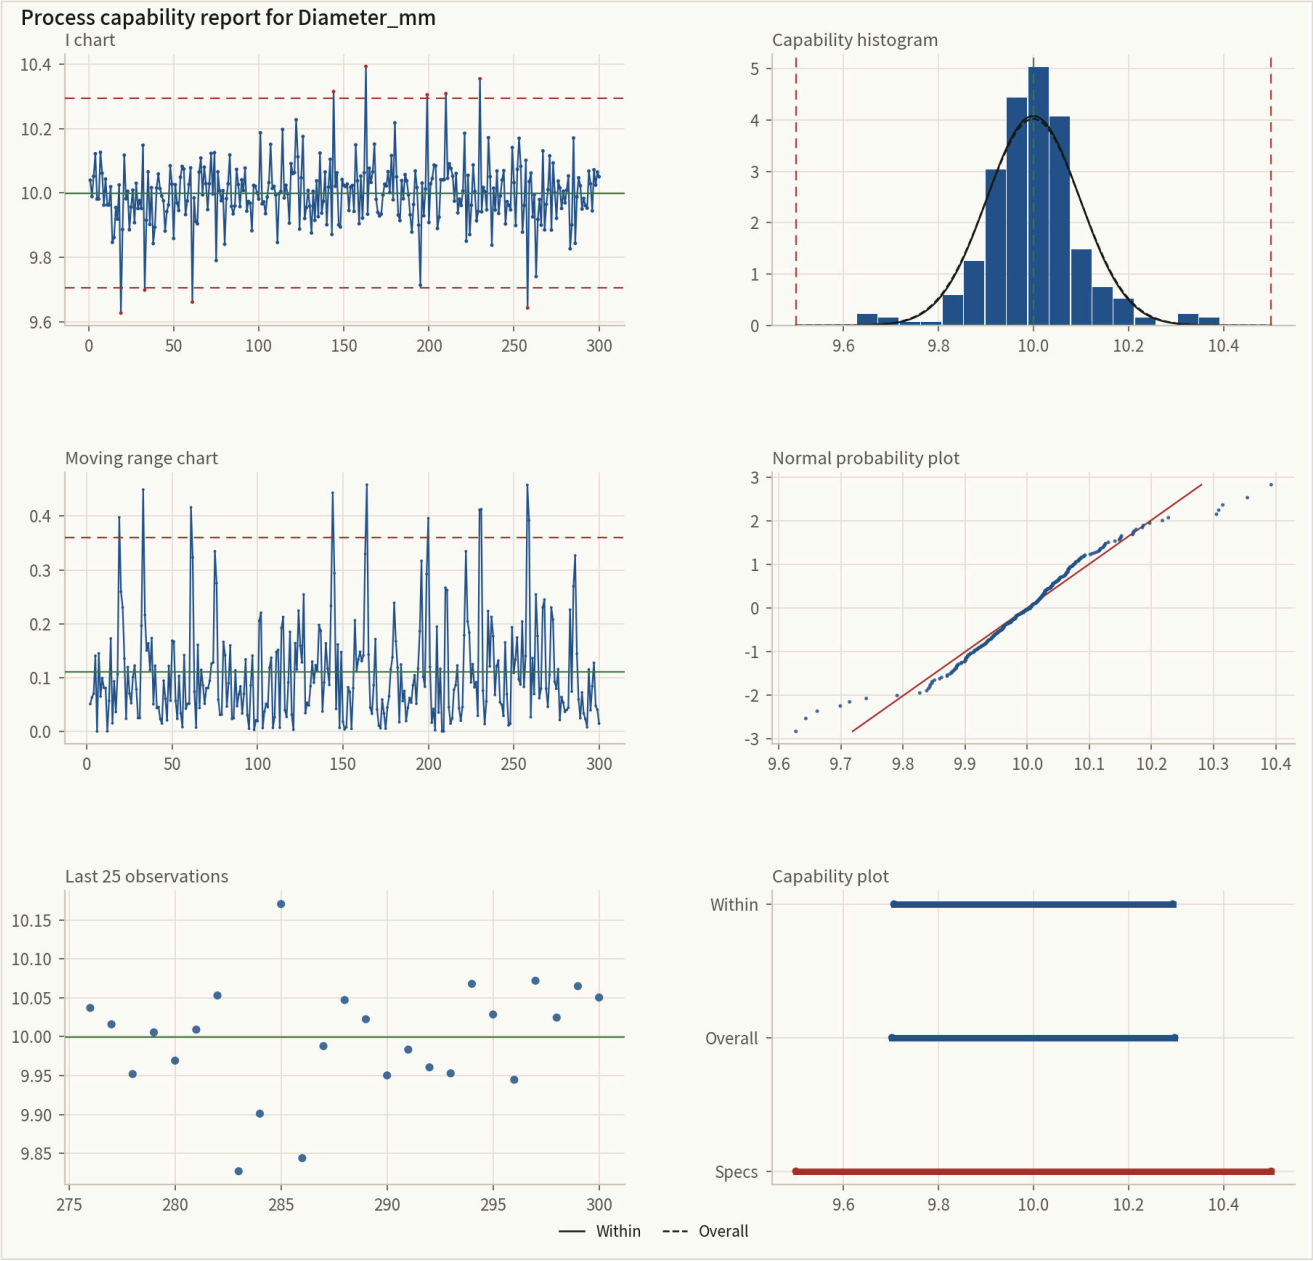


Figure 16.1 Capability report for the aerospace diameters.

Panel	Question it answers	Here
I chart	Is the process stable?	Mostly; 9 isolated points flagged
Capability histogram	Where does the data sit in the tolerance?	Centered, well inside 9.5 and 10.5
Moving range chart	Is short-term variation stable?	Yes, with spikes at the same points
Normal probability plot	Is the data normal enough for the calculation?	Close to the line, slight curve in the tails
Last 25 observations	What is the process doing right now?	Scattered around 10.00

Panel	Question it answers	Here
Capability plot	How do the within, overall and spec widths compare?	Both process bars are well inside the specs bar

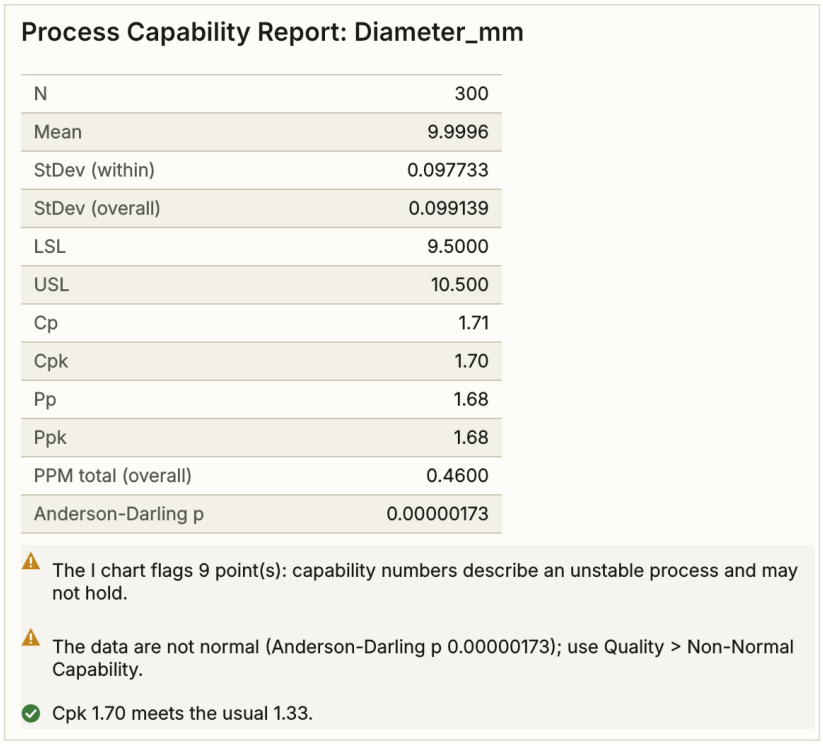


Figure 16.2 The numbers and notes under the report. The program warns about the 9 flagged points and the non-normal tails.

The notes are worth reading closely. Here the program says three things: the I chart flags 9 points, so the numbers describe a not-quite-stable process; the data is not normal by the Anderson-Darling test ($p < 0.0001$); and Cpk 1.70 meets 1.33. With 300 readings, tests of normality become very sensitive to small departures in the tails. The probability plot shows the departure is slight, so the normal-based Cpk is a reasonable guide here.

When the process is not stable

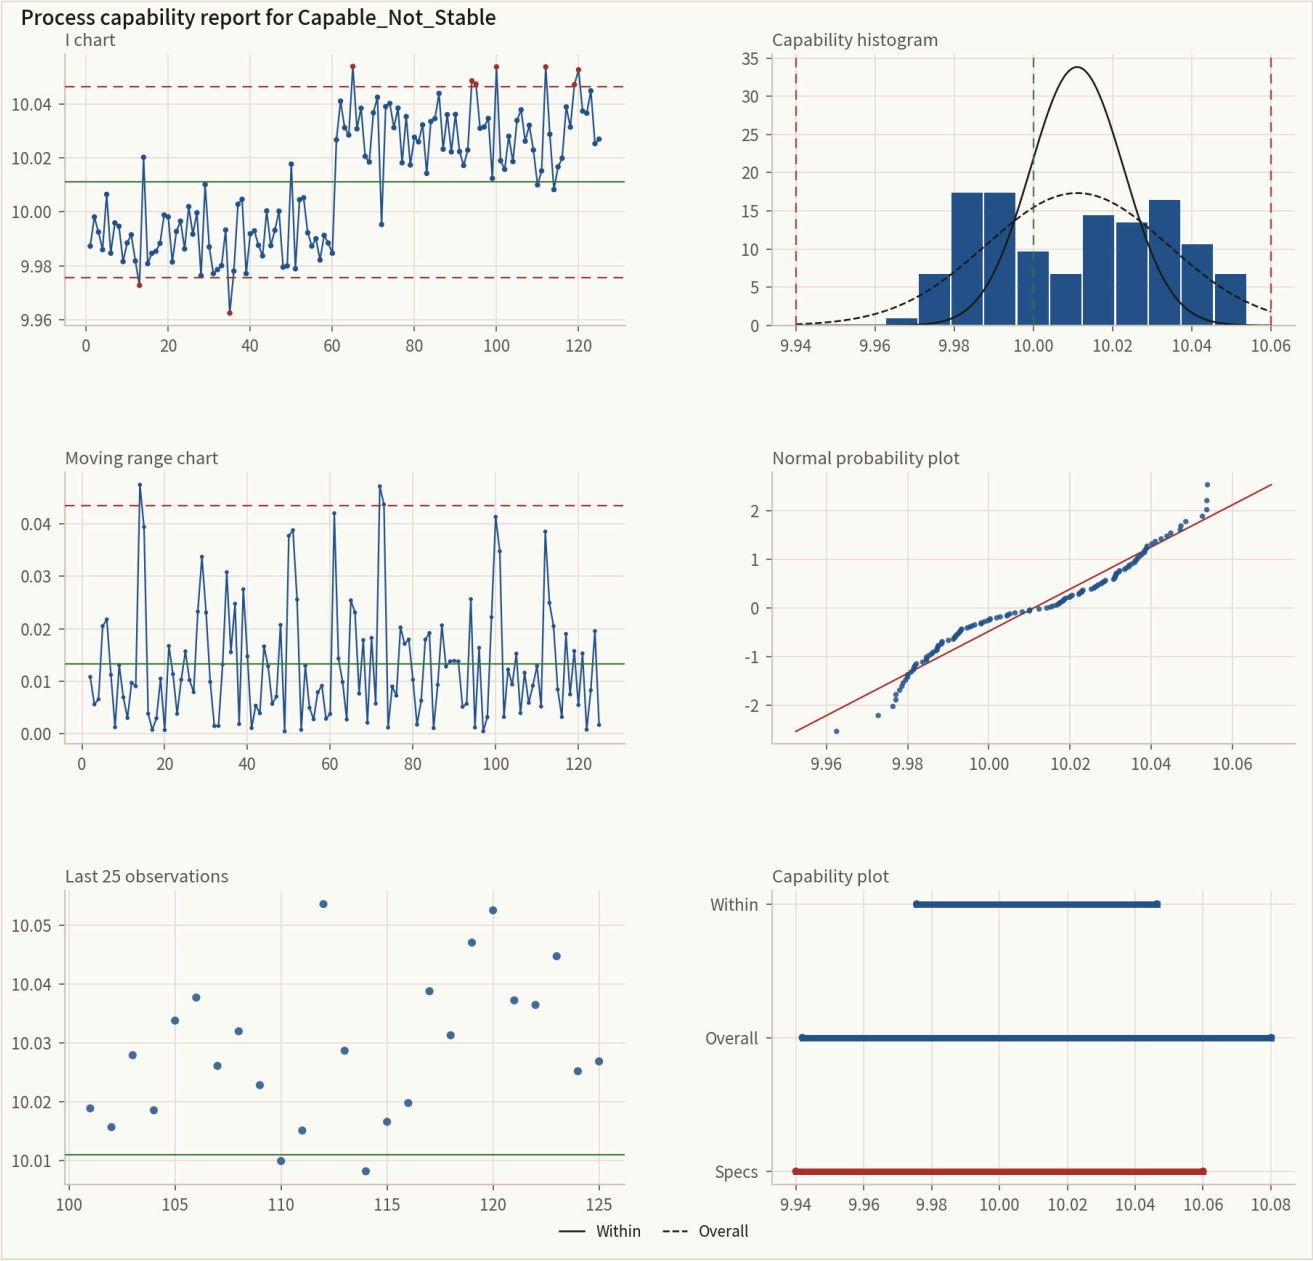


Figure 16.3 The same report for the two-level teaching data from Chapter 15. The I chart shows the jump at part 61; the histogram is two humps.

Every panel tells the same story from a different angle: the I chart shows two levels, the histogram two humps, the probability plot an S-shape, and the capability plot shows the overall bar far wider than the within bar. The note says the I chart flags points, so the capability numbers describe an unstable process and may not hold. Fix the jump first.

Normal or not? The probability plot

Where: Explore > Distributions > Probability Plot

A probability plot redraws the data so that a perfectly normal sample falls on a straight line. Bends show how it differs:

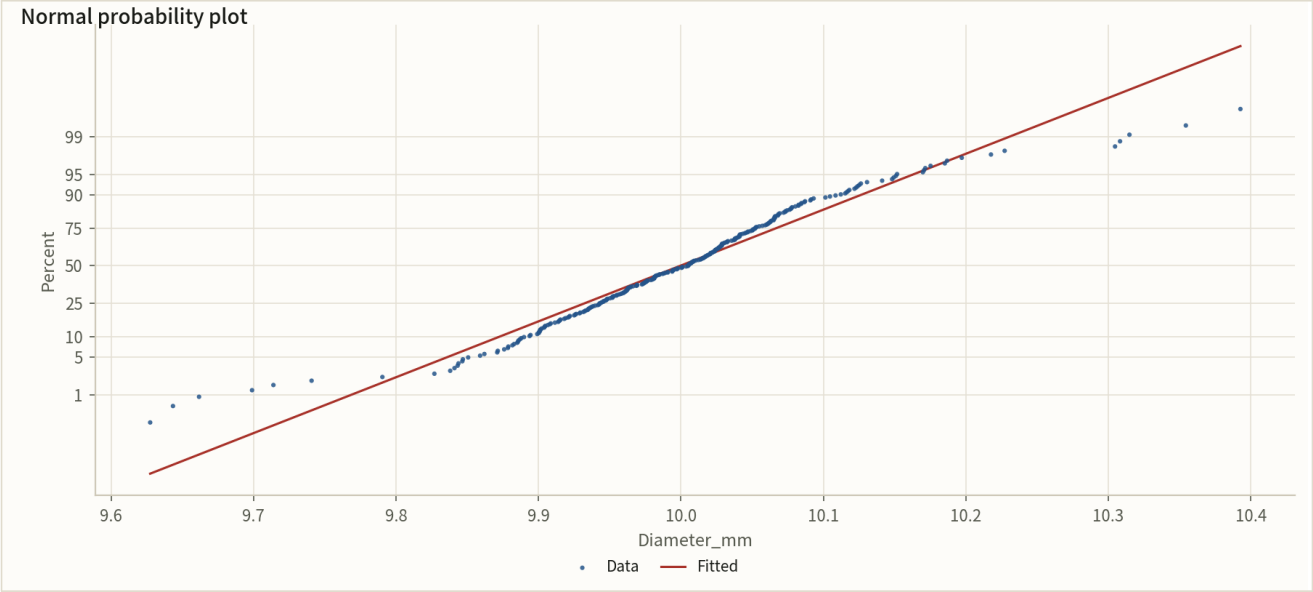


Figure 16.4 Diameters on a normal probability plot: nearly straight, with a few points leaving the line at the ends.

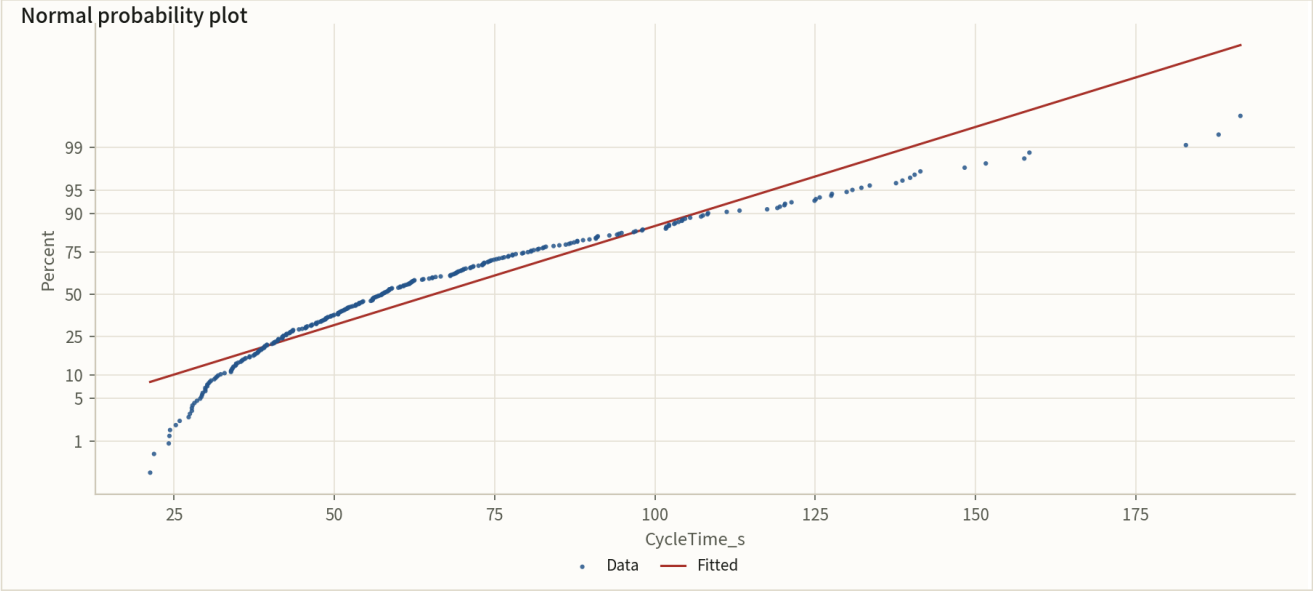


Figure 16.5 Cycle times: a strong curve. Short times bunch up, long times stretch out. This data is skewed, not normal.

Shape on the plot	Means
Straight line	Normal
Curve bending up (like a smile)	Skewed right: a long tail of large values

Shape on the plot	Means
S-shape	A mixture of two processes, or heavy tails
Steps or flat runs	Coarse gauge resolution or rounded readings
One or two points far off	Outliers

Non-normal capability

Where: Quality > Capability > Non-Normal Cpk

Many characteristics are naturally skewed: flatness, roundness, runout, surface finish, impurity, cycle time. Anything with a natural limit of zero tends to be. Forcing a normal curve onto them gets the tail, the part that matters, wrong. The non-normal study fits a better distribution and works out the out-of-spec fraction from it.

1. Choose the column.
2. Choose the distribution, or **auto** to let the program pick the best fit.
3. Read the fitted curve against the histogram before trusting the numbers.

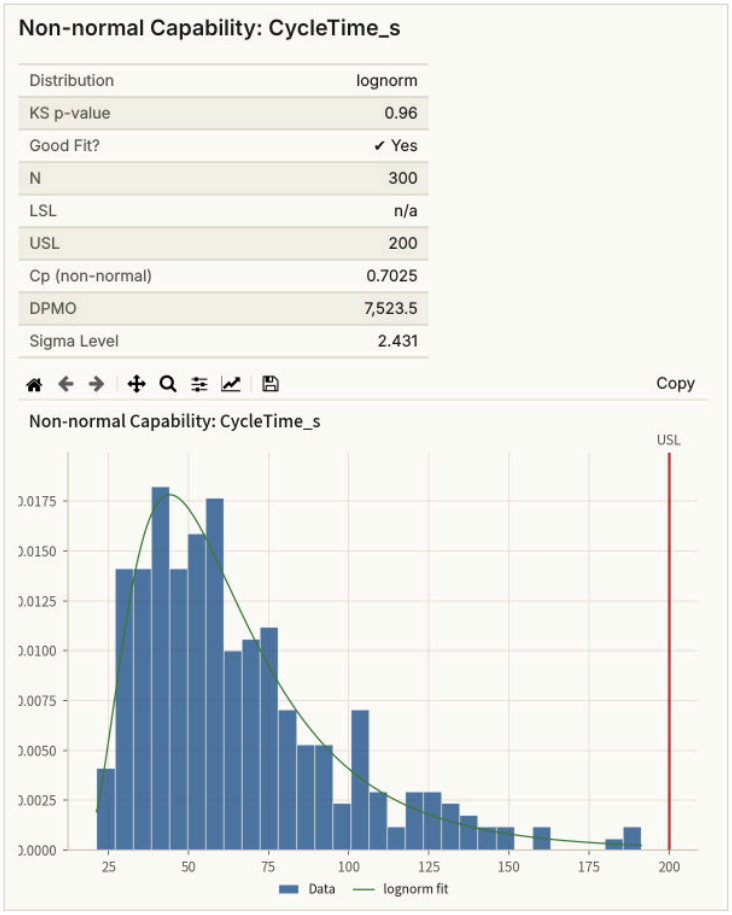


Figure 16.6 Cycle times with an upper limit of 200 s. A lognormal curve fits well (goodness-of-fit p = 0.96).

The lognormal curve follows the long right tail closely. From it, about 7,500 cycles per million (0.75%) would exceed 200 s, and the capability index is 0.70. A normal curve would have put the tail in the wrong place and given a misleading answer.

> **Tip**
> If the data is non-normal because the process is unstable (two humps, a drift), don't fit a fancy distribution: fix the instability. Non-normal methods are for processes that are stable but naturally skewed.

Other capability tools

Command	Use it for
Quality > Capability > Sigma Level	The sigma level and DPMO from spec limits.
Quality > Capability > Yield	Pass/fail data: first-pass yield and DPMO.
Quality > Attribute Capability > Binomial / **Poisson**	Capability of defectives or defect counts.
Explore > Distributions > Tolerance Interval	The range that contains a stated share of all future parts, with confidence.

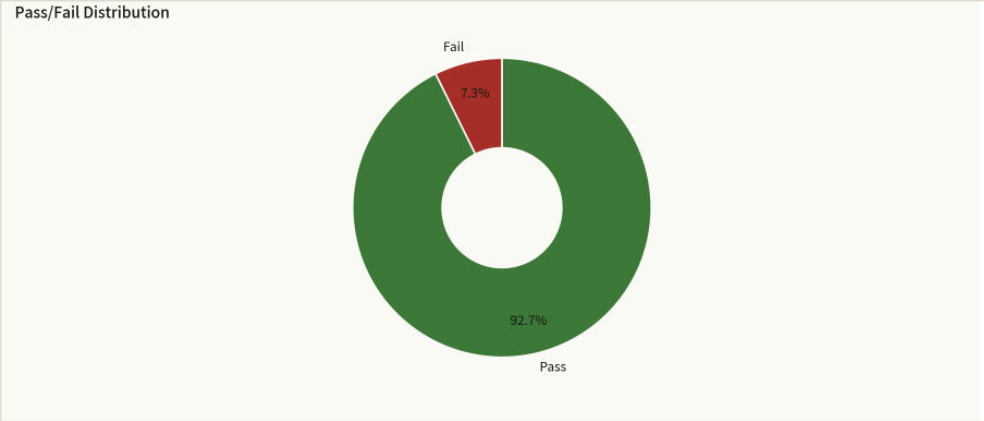


Figure 16.7 Yield of the aerospace Pass_Fail column: 278 of 300 passed, 92.7%, or 73,333 defective per million.

CHAPTER 17

Can you trust the gauge? Measurement systems

Every number in this guide comes through a gauge and a person. Before you blame the process, check the measurement.

The variation you see on a chart is the sum of two things: the parts really being different, and the measurement adding its own error. If the gauge's error is a large share of the tolerance, good parts get rejected, bad parts get shipped, and every capability number is pessimistic. **Measurement System Analysis (MSA)** puts a number on that error.

Study	Answers	Data
Gage R&R	How much do repeat readings scatter (repeatability) and how much do operators differ (reproducibility)?	About 10 parts x 3 operators x 2 or 3 readings each, blind and in random order
Attribute agreement	Do inspectors making pass/fail calls agree with each other and with the right answer?	30 to 50 parts, 2 or 3 inspectors, 2 or 3 rounds, with a known correct answer
Bias and linearity	Does the gauge read true, across its whole range?	Reference parts of known size across the range, each measured 10 or more times

Gage R&R

Where: Home > Measure > Gage R&R

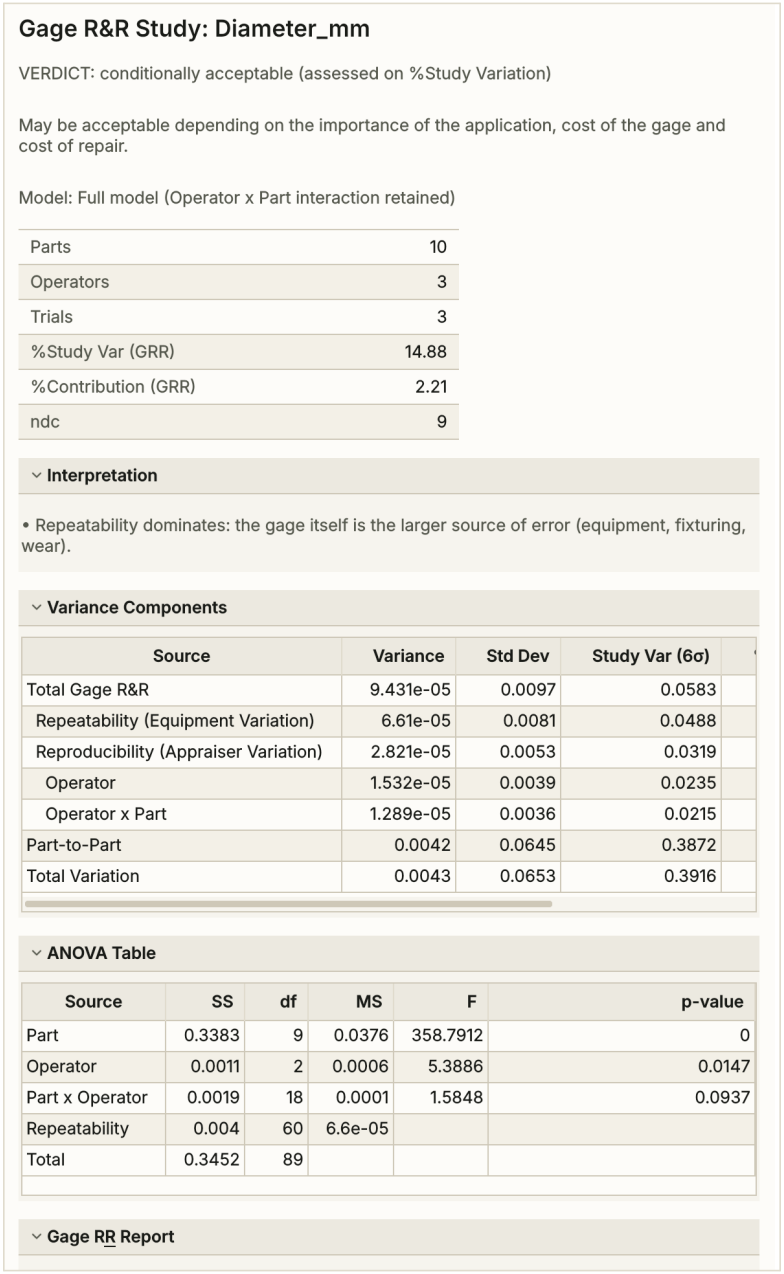
1. Choose the **part** column.
2. Choose the **operator** column.
3. Choose the **measurement**.
4. If you know the tolerance, enter it to see the gauge error as a share of the tolerance too.

The example **Gage R&R study** has 10 parts measured 3 times each by 3 operators, 90 readings.



Figure 17.1 The six-panel Gage R&R report.

Panel	What to look for
Components of variation	Bars for Gage R&R, repeatability, reproducibility and part-to-part. You want Gage R&R small and part-to-part tall. The dashed lines are the 10% and 30% guides.
R chart by operator	The range of each operator's repeat readings on each part. All points inside the limit means each operator is consistent.
Xbar chart by operator	Each operator's average per part. Here most points are outside the narrow limits, which is what you want: the gauge can tell the parts apart.
Measurement by part	Readings for each part with the part average. Parts should differ clearly; the dots for one part should be tight.
Measurement by operator	One box per operator. Boxes at the same level mean operators agree.
Part x operator interaction	One line per operator. Parallel, overlapping lines are good; crossing lines mean operators measure some parts differently.



Gage R&R Study: Diameter_mm

VERDICT: conditionally acceptable (assessed on %Study Variation)

May be acceptable depending on the importance of the application, cost of the gage and cost of repair.

Model: Full model (Operator x Part interaction retained)

Parts	10
Operators	3
Trials	3
%Study Var (GRR)	14.88
%Contribution (GRR)	2.21
ndc	9

Interpretation

• Repeatability dominates: the gage itself is the larger source of error (equipment, fixturing, wear).

Variance Components

Source	Variance	Std Dev	Study Var (6σ)
Total Gage R&R	9.431e-05	0.0097	0.0583
Repeatability (Equipment Variation)	6.61e-05	0.0081	0.0488
Reproducibility (Appraiser Variation)	2.821e-05	0.0053	0.0319
Operator	1.532e-05	0.0039	0.0235
Operator x Part	1.289e-05	0.0036	0.0215
Part-to-Part	0.0042	0.0645	0.3872
Total Variation	0.0043	0.0653	0.3916

ANOVA Table

Source	SS	df	MS	F	p-value
Part	0.3383	9	0.0376	358.7912	0
Operator	0.0011	2	0.0006	5.3886	0.0147
Part x Operator	0.0019	18	0.0001	1.5848	0.0937
Repeatability	0.004	60	6.6e-05		
Total	0.3452	89			

Gage RR Report

Figure 17.2 The verdict and the variance table.

Reading the verdict

%Study Var (gauge share of total variation)	Verdict (AIAG)
Under 10%	Acceptable
10% to 30%	Conditionally acceptable: depends on the importance of the characteristic, the cost of a better gauge and the cost of mistakes
Over 30%	Not acceptable: improve the measurement before using it

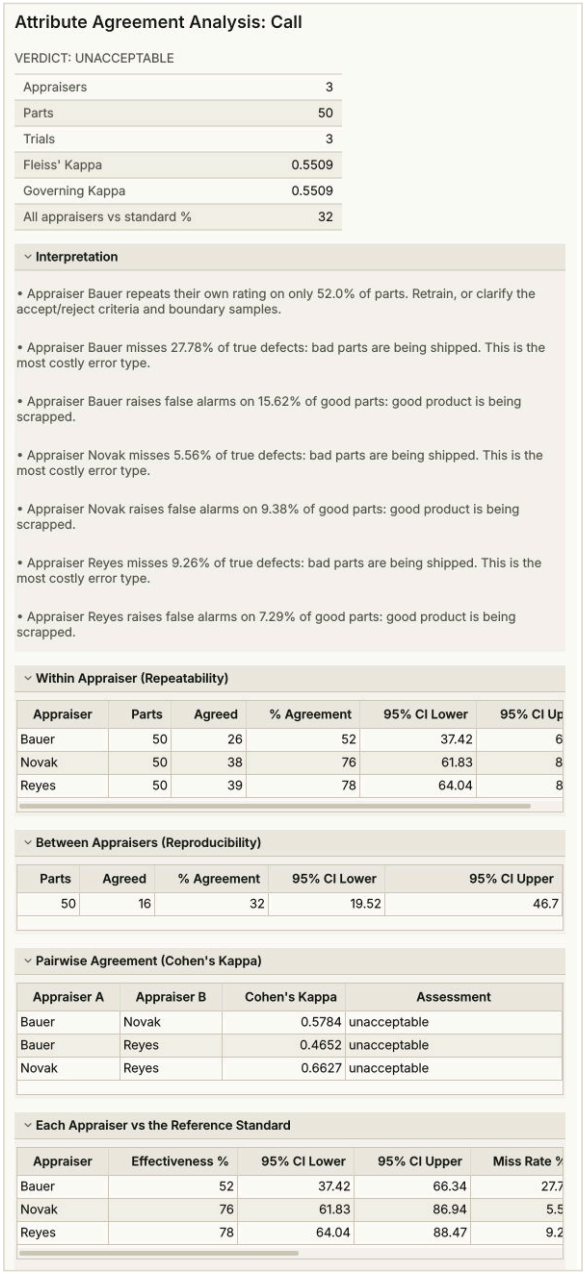
Here %Study Var is **14.9%** and the number of distinct categories (**ndc**) is **9**: the gauge can sort these parts into 9 groups, comfortably above the minimum of 5. The verdict is **conditionally acceptable**, and the note says repeatability dominates: the gauge itself (equipment, fixturing, wear) contributes more error than the differences between operators.

> **In real life**
> When repeatability dominates, look at the gauge and how the part sits in it: a worn anvil, a fixture that lets the part rock, a reading taken at a different place each time. When reproducibility dominates, look at the people: a clearer method sheet, training, a fixture that removes the operator's judgment.

Attribute agreement: pass/fail inspection

Where: Home > Measure > Attribute MSA

Visual inspection, go / no-go gauges and leak testers give a decision, not a number. The attribute agreement study checks whether those decisions can be trusted. The example has 50 parts judged 3 times by 3 inspectors, with the known correct answer for each part.



Attribute Agreement Analysis: Call

VERDICT: UNACCEPTABLE

Appraisers	3
Parts	50
Trials	3
Fleiss' Kappa	0.5509
Governing Kappa	0.5509
All appraisers vs standard %	32

Interpretation

- Appraiser Bauer repeats their own rating on only 52.0% of parts. Retrain, or clarify the accept/reject criteria and boundary samples.
- Appraiser Bauer misses 27.78% of true defects: bad parts are being shipped. This is the most costly error type.
- Appraiser Bauer raises false alarms on 15.62% of good parts: good product is being scrapped.
- Appraiser Novak misses 5.56% of true defects: bad parts are being shipped. This is the most costly error type.
- Appraiser Novak raises false alarms on 9.38% of good parts: good product is being scrapped.
- Appraiser Reyes misses 9.26% of true defects: bad parts are being shipped. This is the most costly error type.
- Appraiser Reyes raises false alarms on 7.29% of good parts: good product is being scrapped.

Within Appraiser (Repeatability)

Appraiser	Parts	Agreed	% Agreement	95% CI Lower	95% CI Up
Bauer	50	26	52	37.42	6
Novak	50	38	76	61.83	8
Reyes	50	39	78	64.04	8

Between Appraisers (Reproducibility)

Parts	Agreed	% Agreement	95% CI Lower	95% CI Upper
50	16	32	19.52	46.7

Pairwise Agreement (Cohen's Kappa)

Appraiser A	Appraiser B	Cohen's Kappa	Assessment
Bauer	Novak	0.5784	unacceptable
Bauer	Reyes	0.4652	unacceptable
Novak	Reyes	0.6627	unacceptable

Each Appraiser vs the Reference Standard

Appraiser	Effectiveness %	95% CI Lower	95% CI Upper	Miss Rate %
Bauer	52	37.42	66.34	27.7
Novak	76	61.83	86.94	5.5
Reyes	78	64.04	88.47	9.2

Figure 17.3 Attribute agreement results: **unacceptable**.

Fleiss' kappa measures agreement beyond what chance would give: 1 is perfect, 0 is no better than guessing. Above 0.9 is excellent, 0.7 to 0.9 acceptable, below 0.7 needs action. Here it is **0.55**, and all three inspectors agree with the correct answer on only **32%** of parts.

The notes name the problems inspector by inspector. Inspector **Bauer** repeats their own decision on only 52% of parts, misses 28% of the truly defective parts (bad parts shipped, the costliest error), and rejects 16% of good ones (good parts scrapped).

In real life

The usual fixes are a set of boundary samples (a board of photos or real parts showing just acceptable and just rejectable), better lighting, a magnifier, and a clear written criterion. Re-run the study afterwards to prove it worked.

Bias and linearity

Where: Home > Measure > Bias / Linearity

Bias is the difference between the average reading and the true value. **Linearity** asks whether the bias changes across the gauge's range. The example has five reference masters (10, 20, 30, 40, 50) each measured 12 times.

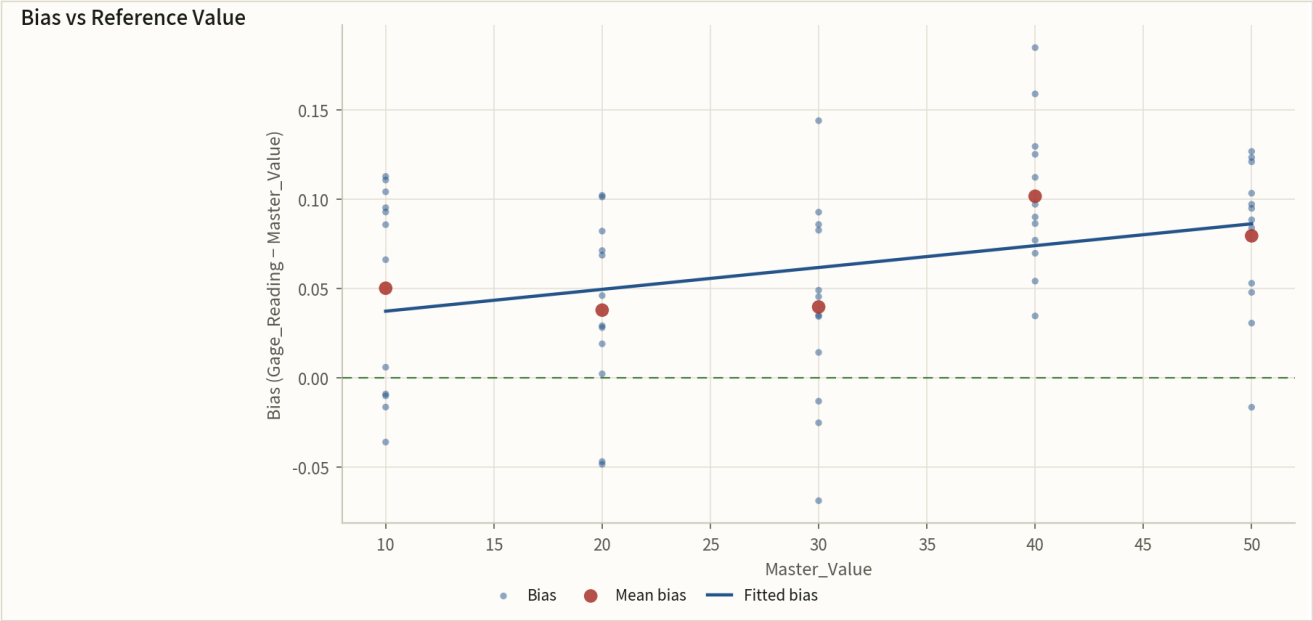
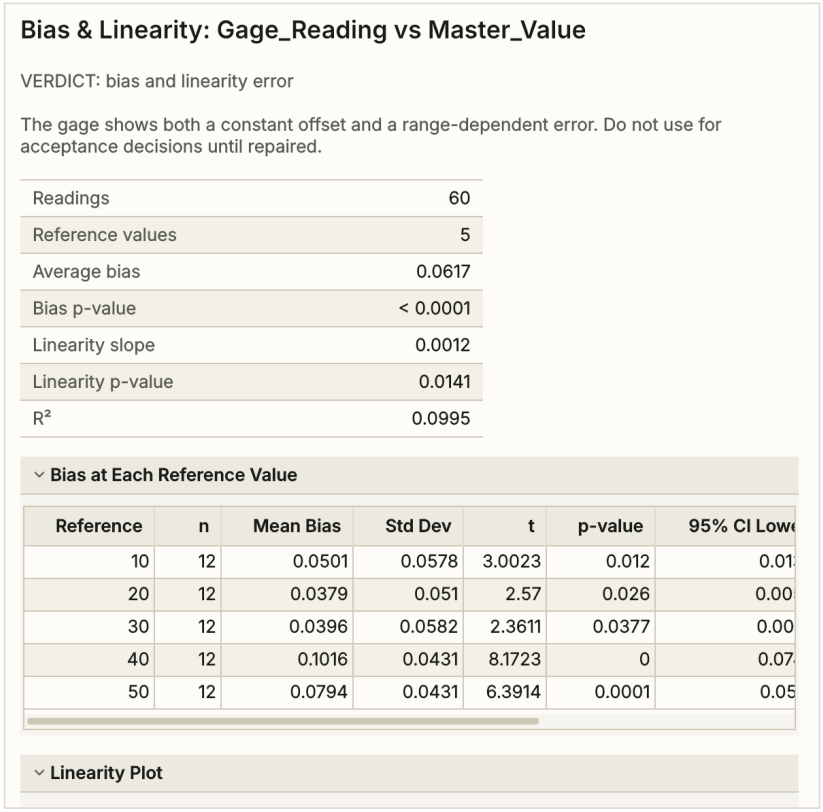


Figure 17.4 Bias at each reference value. Every mean bias (red) is above zero and the fitted line rises.



Bias & Linearity: Gage_Reading vs Master_Value

VERDICT: bias and linearity error

The gage shows both a constant offset and a range-dependent error. Do not use for acceptance decisions until repaired.

Readings	60
Reference values	5
Average bias	0.0617
Bias p-value	< 0.0001
Linearity slope	0.0012
Linearity p-value	0.0141
R^2	0.0995

Bias at Each Reference Value

Reference	n	Mean Bias	Std Dev	t	p-value	95% CI Lowe
10	12	0.0501	0.0578	3.0023	0.012	0.01
20	12	0.0379	0.051	2.57	0.026	0.00
30	12	0.0396	0.0582	2.3611	0.0377	0.00
40	12	0.1016	0.0431	8.1723	0	0.07
50	12	0.0794	0.0431	6.3914	0.0001	0.05

Linearity Plot

Figure 17.5 Average bias 0.062, significant (p < 0.0001); the slope is also significant (p = 0.014).

The gauge reads about 0.06 high on average, and the error grows slightly with size. The verdict: **do not use it for acceptance decisions until it is repaired or recalibrated**.

> **Watch out**
> A gauge with a good Gage R&R can still be wrong: R&R measures scatter, not accuracy. A gauge that reads consistently 0.06 high passes R&R and fails bias. You need both studies.

> **Try it yourself**
> Open the Gage R&R example and run the study again, entering a tolerance of 0.20 (the width of the part spec). Is the gauge acceptable as a share of the tolerance? Which figure would your customer care about?

PART IV

Finding causes

Something is wrong. Which problem first, where is it coming from, which inputs drive it, and what settings fix it.

CHAPTER 18

Which problem first? Pareto and the fishbone

Most of the trouble comes from a few causes. Find them before spending effort on the rest.

The Pareto chart

Where: Quality > Problem Solving > Pareto

A Pareto chart sorts categories (defect types, reject reasons, machines) from most to least frequent and adds a running total. The **80/20 rule** says a few categories usually account for most of the occurrences. Fix those first.

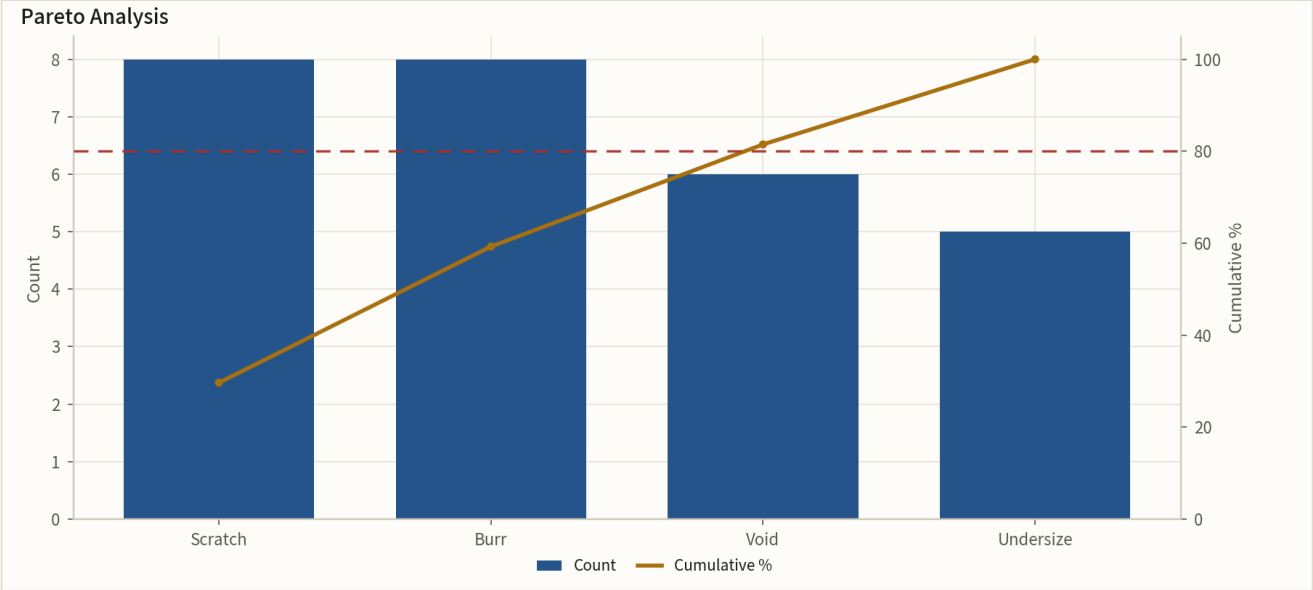


Figure 18.1 Pareto of defect types in the aerospace example. Bars are counts; the amber line is the cumulative percentage.

Scratches and burrs (8 each) together are 59% of all defects; adding voids takes the total past the dashed 80% line. The program's note says 2 of 4 categories account for about 80% of occurrences. Here the categories are fairly even, so the "vital few" are less dominant than in a typical shop, and the decision should also weigh cost.

> **Tip**
> Run the Pareto on **cost** as well as on count when you can. Ten scratches that are polished out in a minute matter less than two undersize parts that are scrapped after five hours of machining. Add a cost column and use it as the value.

In real life

A stamping plant's Pareto of 2,000 rejects shows three of 18 reasons make up 74%: burr, split and wrinkle, all traced to one worn die. One die repair removes three quarters of the rejects; the other 15 reasons are left until later.

The fishbone (cause and effect) diagram

Where: Quality > Problem Solving > Fishbone

The fishbone sorts possible causes into the classic six branches: **Man** (people), **Machine**, **Method**, **Material**, **Measurement** and **Environment**. RS-1 builds it from your columns: it places each column on the branch that fits its name, and if you choose an effect to explain it works out how much of that effect's variation goes with each column.

The teaching dataset **shaft turning** records 120 shafts with the spindle temperature, how many parts the current tool has already cut, the feed rate, and the finished diameter.

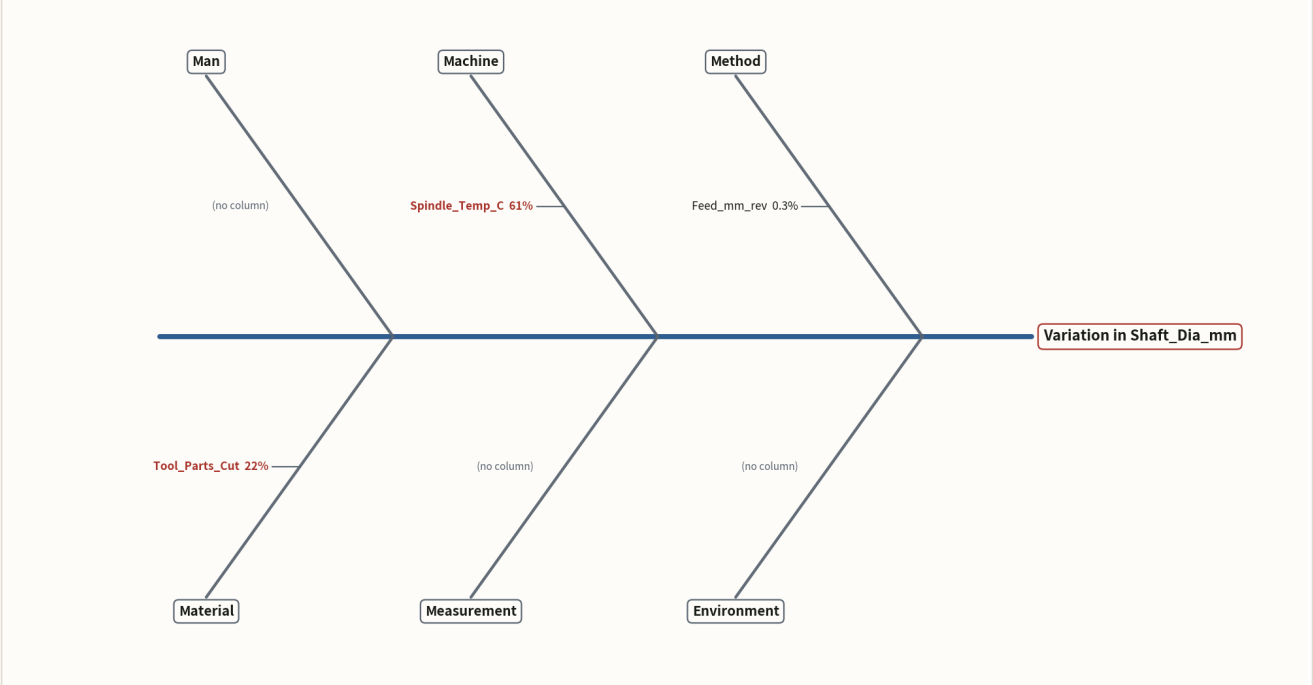


Figure 18.2 Fishbone for variation in shaft diameter. The percentages are each column's share of the diameter's variation.

Spindle temperature goes with **61.4%** of the variation and tool wear with **22.4%**; feed rate with almost none (0.3%). Branches with (no column) are causes your data does not record. They are reminders of what else to check, not evidence that they don't matter.

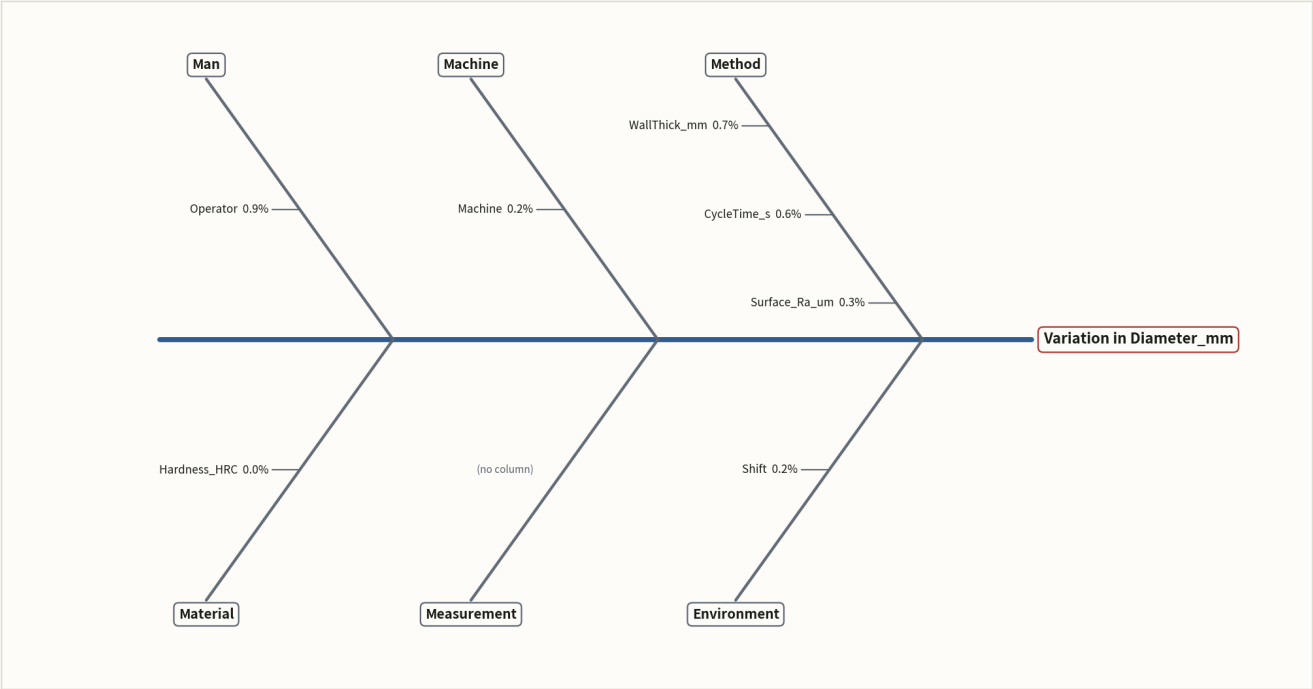


Figure 18.3 The same diagram for the aerospace diameters: no column explains more than 5% of the variation.

A fishbone with nothing above a few percent is useful news too: **the cause is not in this data**. The program says so (the main cause may not be recorded in this data). The next step is to collect what is missing: tool changes, material lots, temperatures.

> **Watch out**
> The percentages show **association, not proof**. Hot spindles and large diameters go together in this data, but so might the afternoon, a particular operator or a particular bar of material. Confirm a suspected cause by changing it deliberately (Chapter 22) or by checking the physics.

CHAPTER 19

Comparing groups: machines, shifts, suppliers

Are the machines different? Did night shift run differently? A picture first, then a test that says whether the difference is real.

Look first: box plots and mean plots

Where: Visualize > Distribution > Box Plot

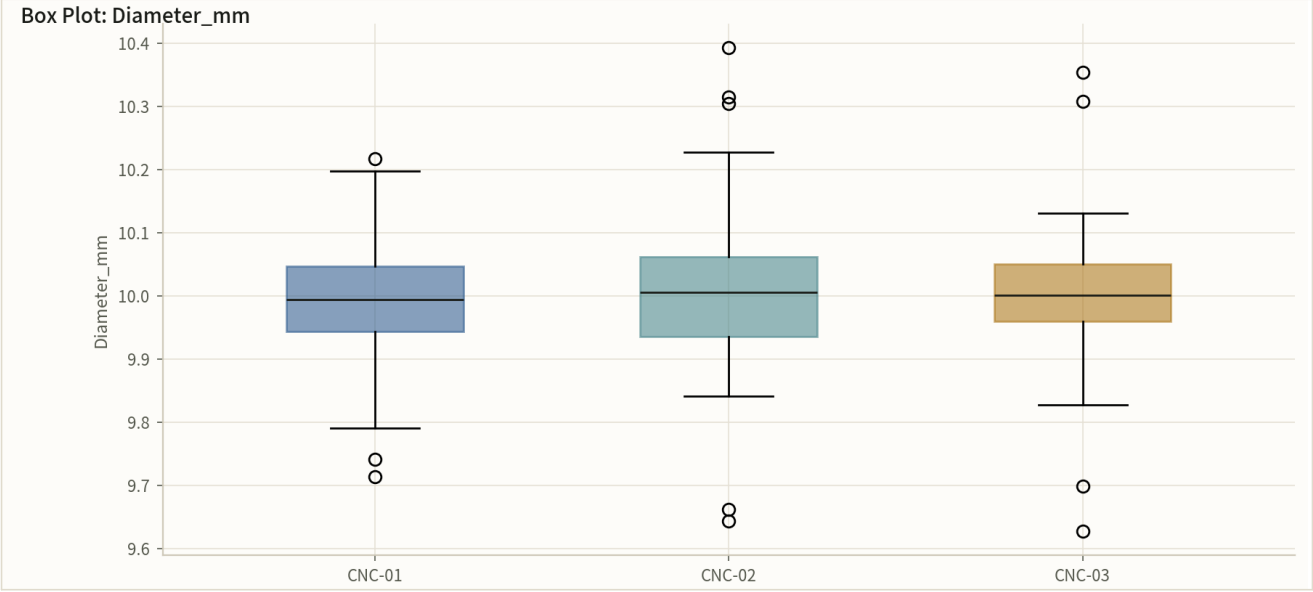


Figure 19.1 Diameter by machine. Each box holds the middle half of the parts; the line inside is the median; whiskers reach the bulk of the data; circles are unusual parts.

The three boxes sit at the same level and are about the same height: the machines look alike. CNC-02 has the most unusual parts (circles) on both ends, which is worth a look on its own.

Where: Visualize > Comparison > Mean Plot

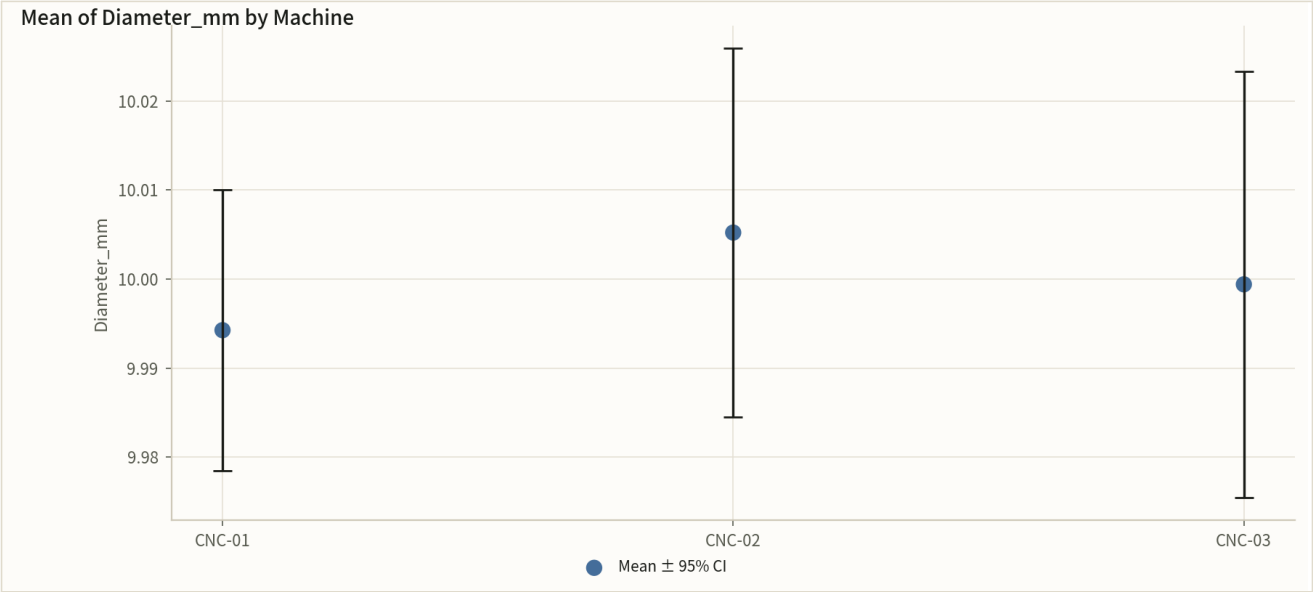


Figure 19.2 Mean diameter of each machine with its 95% confidence interval.

The intervals overlap heavily: any difference between the machines' averages is smaller than the uncertainty in them.

Then test: ANOVA

Where: Home > Analyze > ANOVA

One-way ANOVA asks whether the averages of several groups differ more than chance would explain. Choose the value column and the group column.

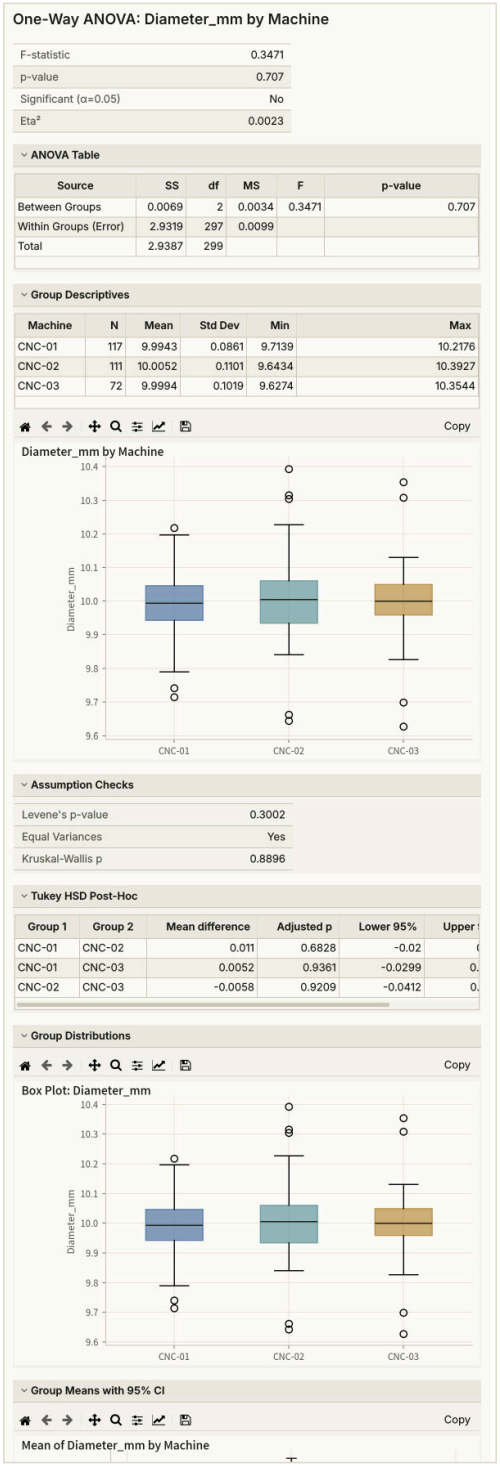


Figure 19.3 One-way ANOVA of diameter by machine (top of the results).

The **p-value** is 0.707. A p-value is the chance of seeing differences at least this big if the groups were really the same. Small p-values (below 0.05 by convention) mean the difference is probably real. Here it is large, so the program's verdict is **no significant difference between groups**.

The rest of the page checks the test's assumptions (equal spreads by Levene's test, a non-parametric Kruskal-Wallis test that does not assume normality) and, under Tukey HSD Post-Hoc, compares every pair of

groups so you can see **which** groups differ when the overall test says some do.

Two groups: the 2-sample t-test

Where: Home > Analyze > 2-Sample t

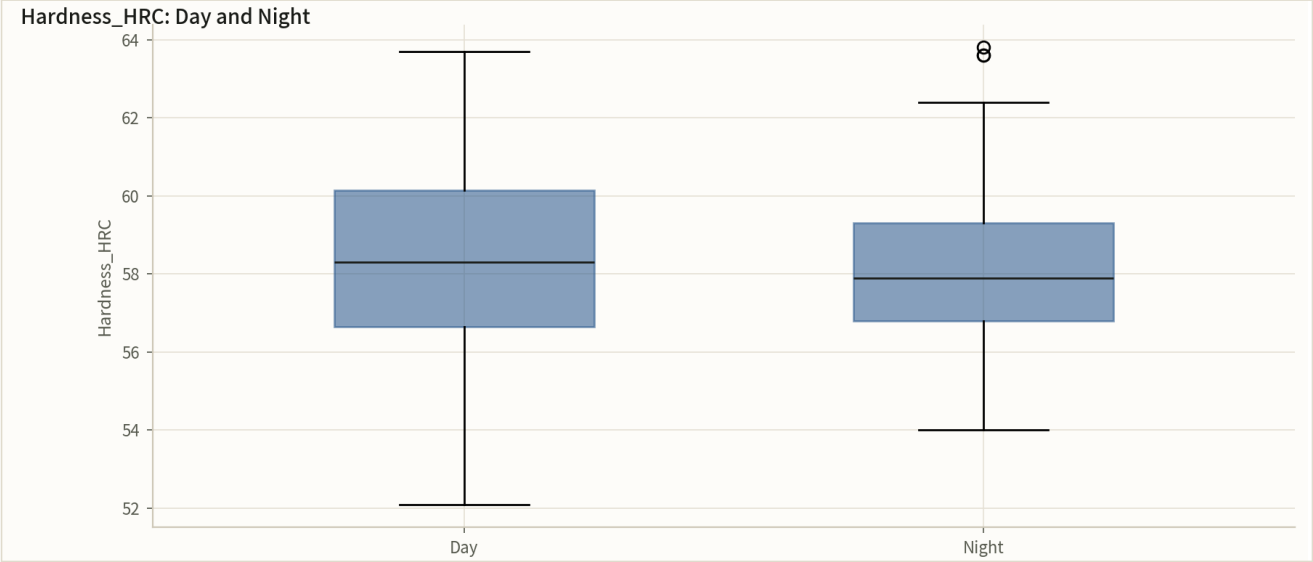


Figure 19.4 Hardness on day and night shift.

2-sample t: Day vs Night

Difference = mean of Day minus mean of Night. Null hypothesis: difference = 0. Alternative: difference ≠ 0.

Day

N	143
Mean	58.31
Standard deviation	2.4078

Night

N	157
Mean	58.063
Standard deviation	1.9787

Test

Difference	0.24743
95% CI for the difference	-0.2563 to 0.7512
t	0.9669
Degrees of freedom	275.5
p-value	0.3344
Method	Welch

No evidence that the two means differ.

Figure 19.5 The 2-sample t results.

Day shift averages 58.31 HRC and night shift a little less. The difference is 0.24743 HRC, and the 95% confidence interval for it (-0.2563 to 0.7512) includes zero. The p-value is 0.3344: **no evidence the shifts differ**. The program used Welch's version of the test, which does not assume the two groups have the same spread.

Significant is not the same as important

With enough data, a tiny difference becomes statistically significant. With too little, a large one does not. Always read the size of the difference and its confidence interval, and ask whether it matters against the tolerance:

Result	Meaning	Do
Significant and large	A real difference that matters	Act on it
Significant but tiny	Real, but too small to matter	Usually leave it
Not significant, interval narrow	The groups are about the same	Stop looking here
Not significant, interval wide	Not enough data to tell	Collect more

Other comparisons

Question	Command
Two groups, data not normal or with outliers	**Compare Groups** tab: Mann-Whitney
The same parts measured twice (before / after on each part)	**Compare Groups** tab: Paired t
Two proportions (scrap rate of two lines)	**Compare Groups** tab: 2 Proportions
Do two groups have the same spread?	**Compare Groups** tab: 2 Variances
Two factors at once (machine and shift)	**Compare Groups** tab: Two-Way ANOVA
Where does the variation come from, layer by layer?	**Quality > Problem Solving > Variability**

> **In real life**
> A supplier change is a classic two-group question. Measure 30 parts from each supplier, run a 2-sample t on the critical dimension and a 2 Variances test on its spread. A new supplier whose parts average the same but vary twice as much will cost you capability, and only the variance test shows it.

CHAPTER 20

Relationships: what drives the result?

When an input is a number (a temperature, a pressure, a feed rate), look at how the output moves with it.

The scatter plot

Where: Visualize > Relationships > Scatter

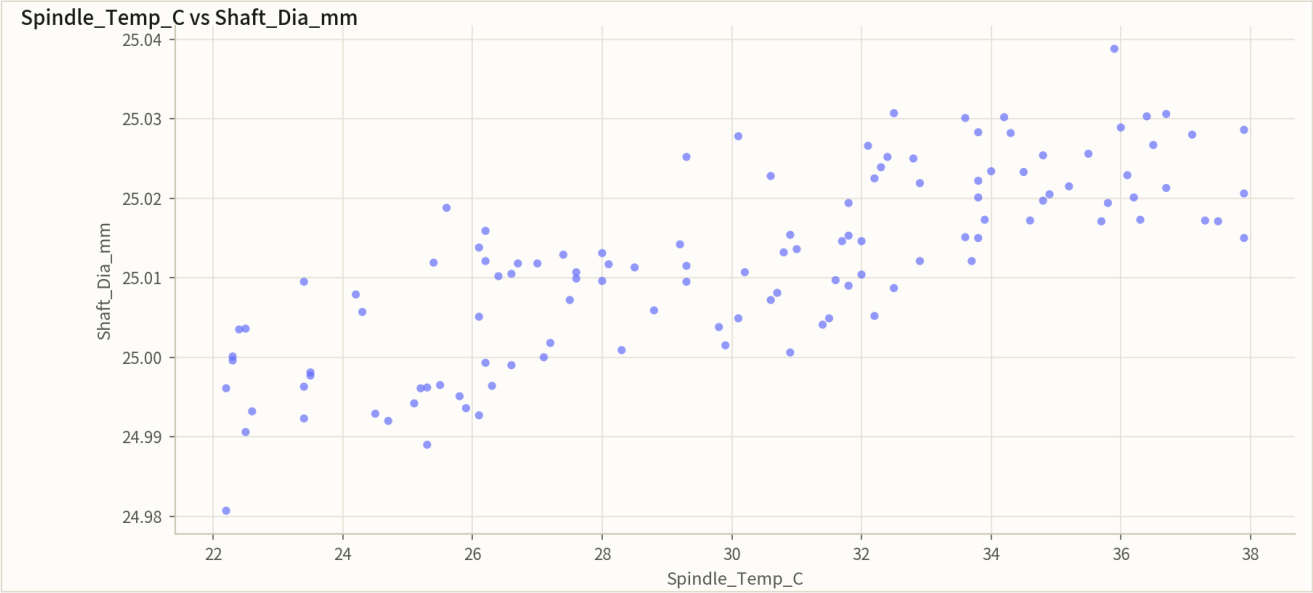


Figure 20.1 Shaft diameter against spindle temperature: a clear upward slope. Warmer spindle, bigger shaft.

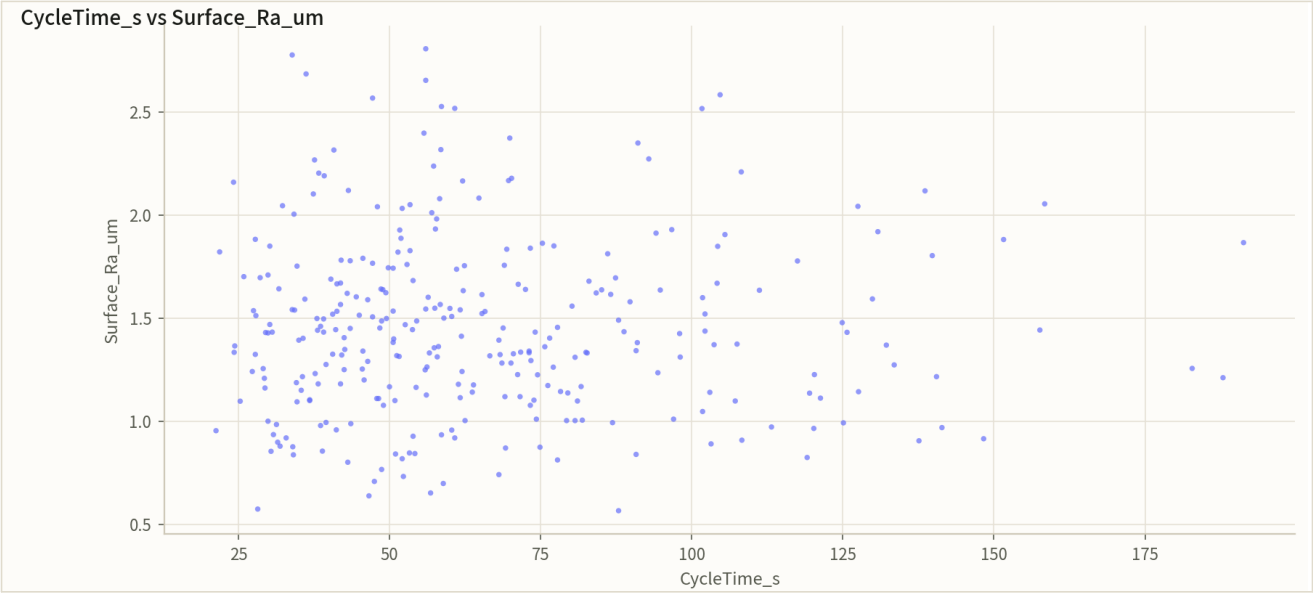


Figure 20.2 For contrast: cycle time against surface finish in the aerospace data. A shapeless cloud means no relationship.

The correlation matrix

Where: Explore > Relationships > Correlation

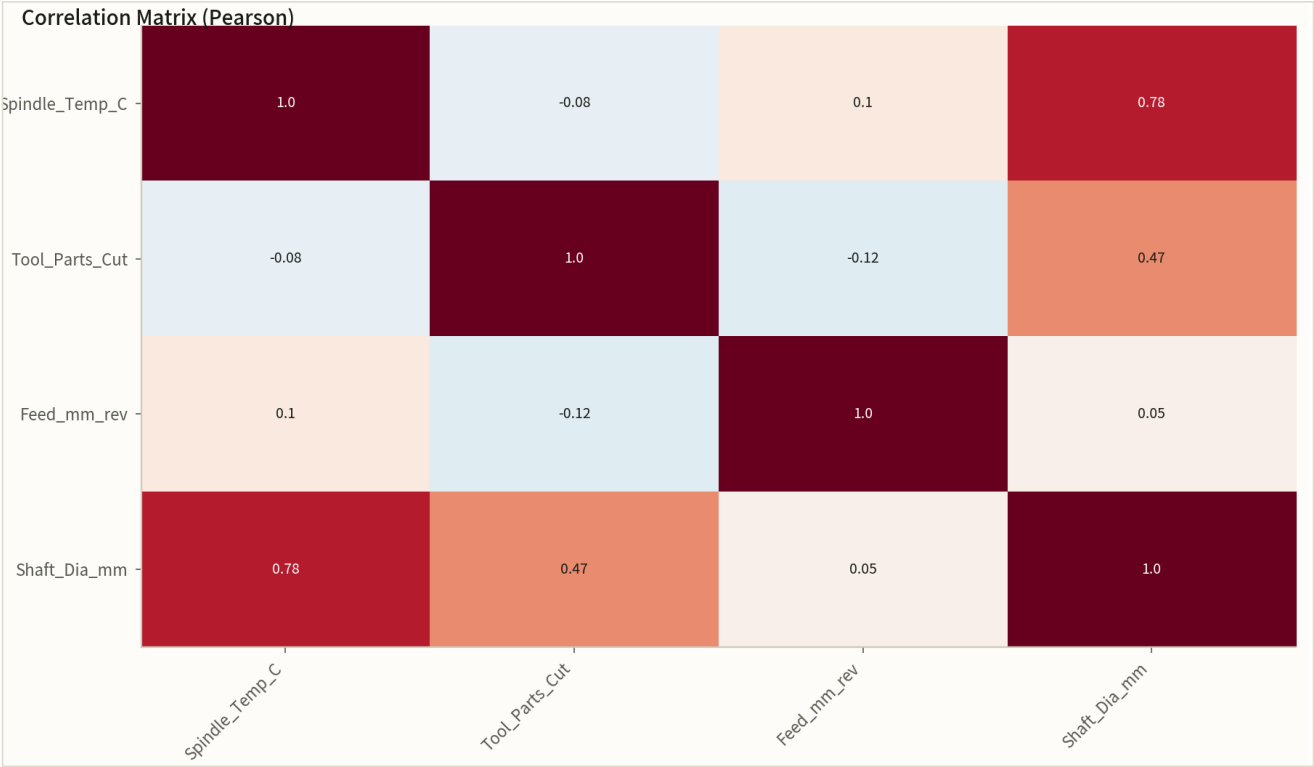


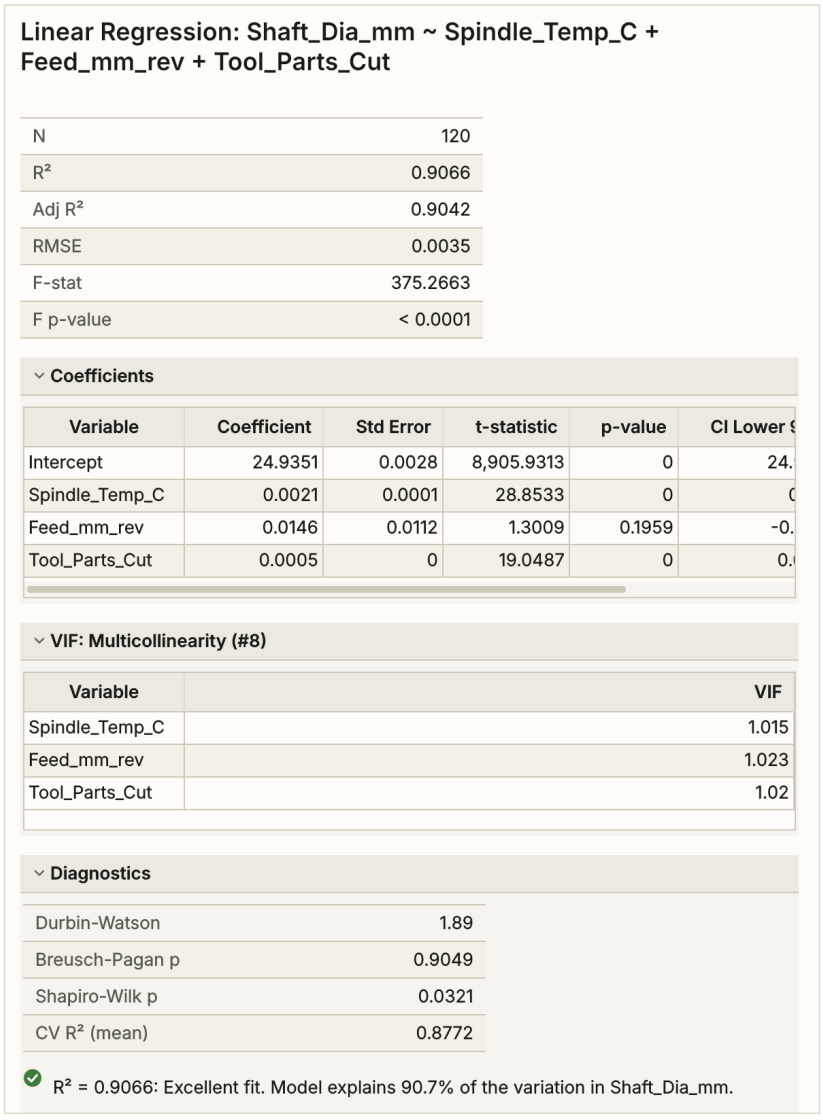
Figure 20.3 Every pair of columns at once. Dark red is strong positive, blue negative, pale none. The diagonal is each column with itself.

The bottom row is the one to read: the shaft diameter correlates 0.7834 with spindle temperature and 0.4736 with tool wear, and hardly at all with feed (0.0507). Correlation runs from -1 to +1; above about 0.7 (or below -0.7) is strong.

Regression: how much, and which inputs matter

Where: Predict > Regression > Linear

1. Choose the **response** (the result): Shaft_Dia_mm.
2. Choose the **predictors** (inputs): spindle temperature, parts cut and feed rate. Hold Ctrl to pick several.



Linear Regression: Shaft_Dia_mm ~ Spindle_Temp_C + Feed_mm_rev + Tool_Parts_Cut

N	120
R^2	0.9066
Adj R^2	0.9042
RMSE	0.0035
F-stat	375.2663
F p-value	< 0.0001

Coefficients

Variable	Coefficient	Std Error	t-statistic	p-value	CI Lower
Intercept	24.9351	0.0028	8,905.9313	0	24.
Spindle_Temp_C	0.0021	0.0001	28.8533	0	0
Feed_mm_rev	0.0146	0.0112	1.3009	0.1959	-0.
Tool_Parts_Cut	0.0005	0	19.0487	0	0.

VIF: Multicollinearity (#8)

Variable	VIF
Spindle_Temp_C	1.015
Feed_mm_rev	1.023
Tool_Parts_Cut	1.02

Diagnostics

Durbin-Watson	1.89
Breusch-Pagan p	0.9049
Shapiro-Wilk p	0.0321
CV R^2 (mean)	0.8772

R^2 = 0.9066: Excellent fit. Model explains 90.7% of the variation in Shaft_Dia_mm.

Figure 20.4 Linear regression of shaft diameter on the three inputs.

R^2 = 0.9066: the three inputs together explain 91% of the variation in diameter. The **coefficients** say how much the diameter changes per unit of each input, holding the others steady: 0.0021 mm per degree of

spindle temperature and 0.0005 mm per part the tool has cut. Their p-values are effectively zero: both are real effects. Feed rate's p-value is 0.1959: no evidence it matters.

VIF (variance inflation factor) checks whether inputs move together so much that their effects cannot be separated. Values near 1, as here, are ideal; above about 5 deserves attention.

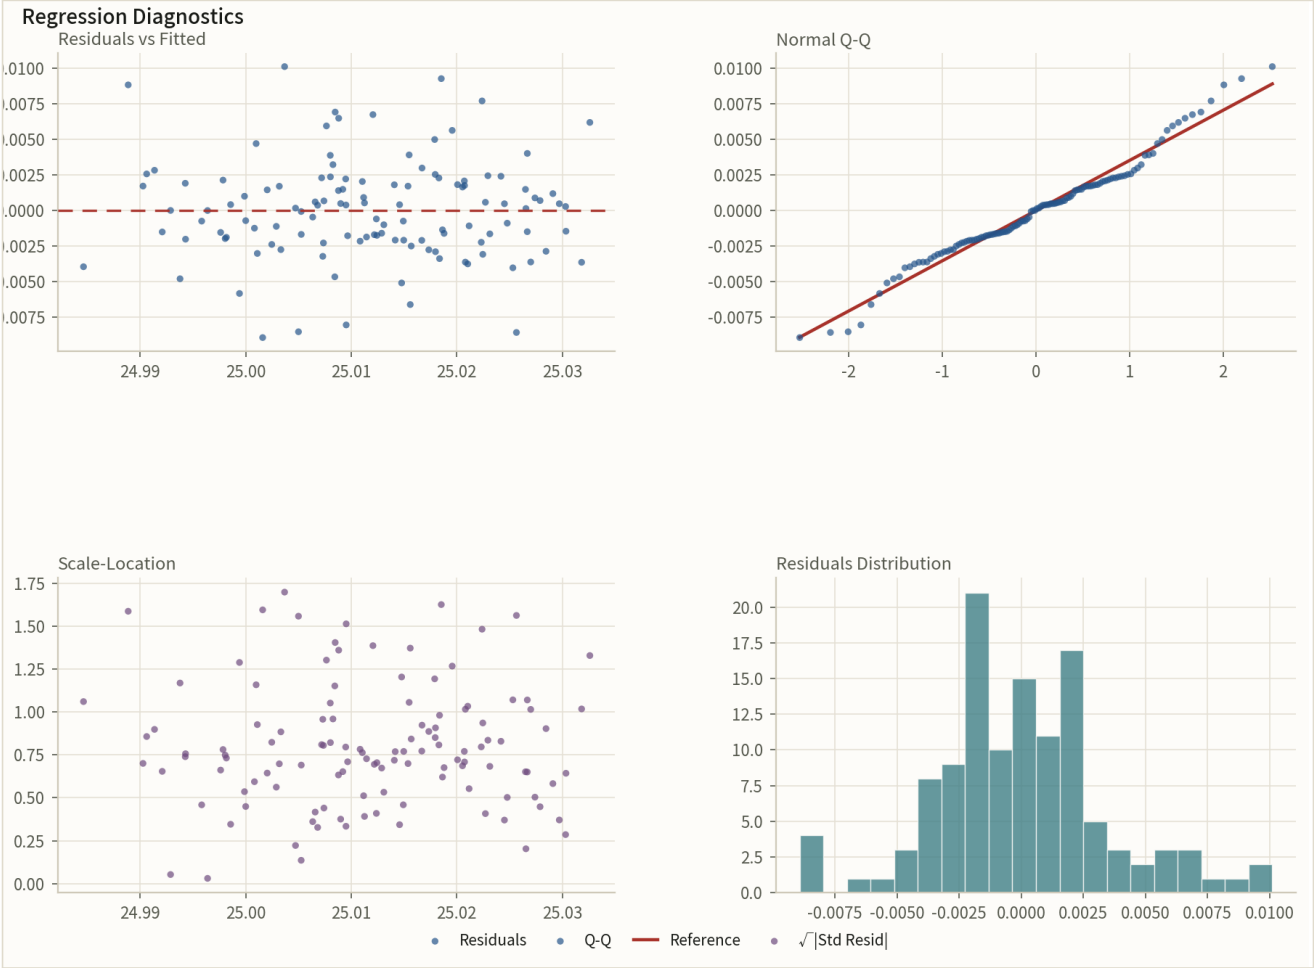


Figure 20.5 Diagnostics: residuals (what the model missed) should scatter evenly around zero with no pattern, and follow the line on the normal Q-Q plot.

In real life

The model turns into actions. With 0.0021 mm per degree, a spindle that warms 15 degrees through the morning moves the diameter 0.03 mm, a third of a 0.1 mm tolerance. Options: warm the spindle up before the first part, apply a temperature compensation offset, or tighten the tool-change interval, since wear adds another 0.02 mm over 40 parts.

Watch out

Regression describes the data you have. Don't predict far outside the range you measured (here 22 to 38 degrees).

A relationship in production data is a strong lead, not proof. Confirm it with a designed experiment (Chapter 22) before changing the process for good.

CHAPTER 21

Calculators: Z-scores and probabilities

Quick answers to everyday questions: how unusual is this reading, what share of parts will be over the limit?

Z-score

Where: Home > Analyze > Z-Score

A Z-score says how many standard deviations a value is from the mean. It turns any measurement into a common scale: Z = 2 is unusual, Z = 3 is rare for a normal process.

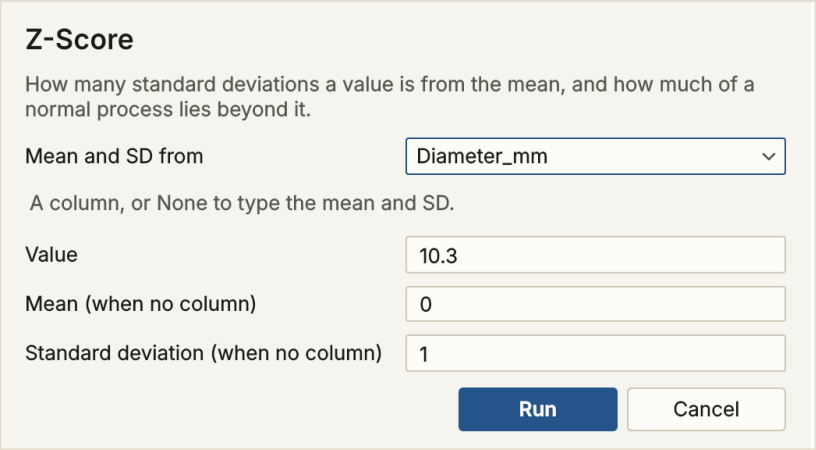


Figure 21.1 Take the mean and standard deviation from a column, or type them in.

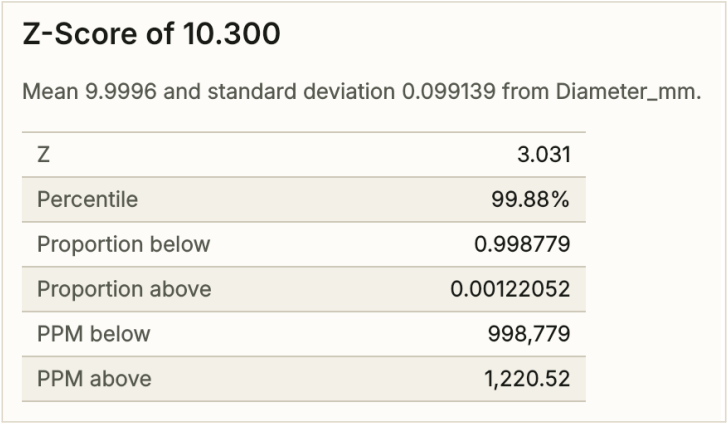


Figure 21.2 How unusual is a 10.300 mm bore?

A 10.300 mm bore is **Z = 3.031** standard deviations above the average. Only 1,220.52 parts per million would be larger, if the process is normal.

In real life

An operator measures a part at 10.30 and asks: is that just bad luck or has something changed? Z above 3 says this would happen by chance about once in a thousand parts. Check the next few parts; if they are high too, the process has moved.

Probability calculator

Where: Explore > Distributions > Probability

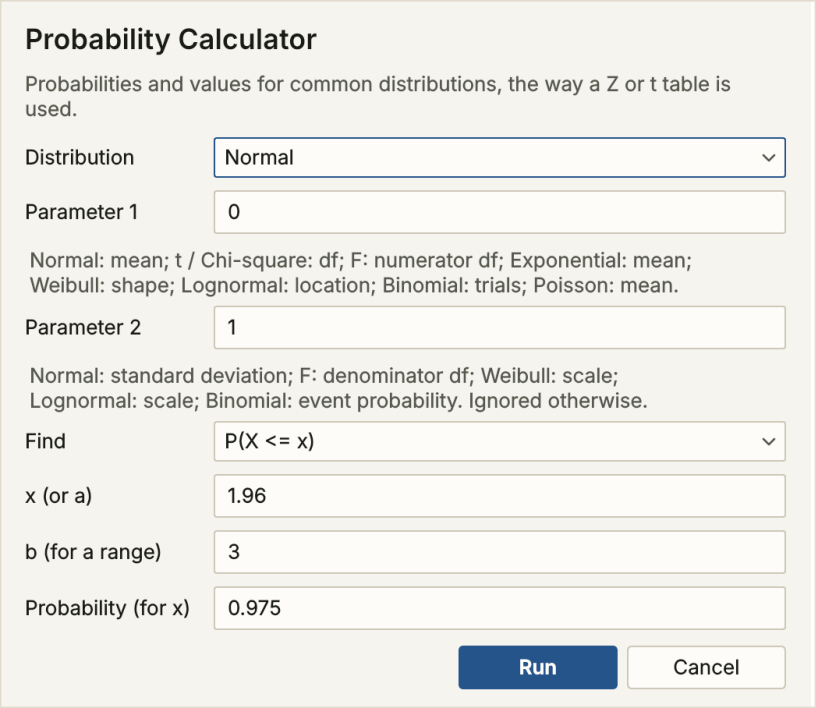


Figure 21.3 Choose a distribution, its parameters and what to find.

The calculator works like the statistics tables at the back of a textbook, for the normal, t, chi-square, F, exponential, Weibull, lognormal, binomial and Poisson distributions. It finds:

- **P(X <= x)** and **P(X >= x)**: the share below or above a value.
- **P(a <= X <= b)**: the share between two values (the in-spec share).
- **x for a probability**: the value below which a given share falls (the 99th percentile, for example).

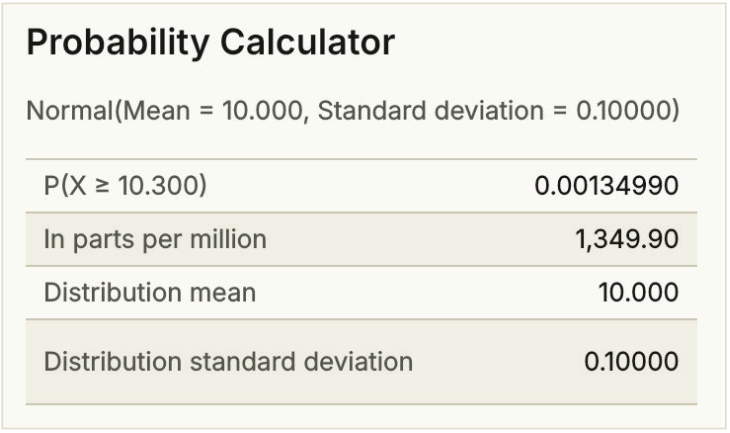


Figure 21.4 A normal process with mean 10.0 and standard deviation 0.1: what share is at or above 10.3?

Try it yourself

A filler runs at a mean of 502 g with a standard deviation of 1.5 g. The label says 500 g. What share of packs is under 500 g? (Normal, P(X <= x), x = 500.)

Bearings have a Weibull life with shape 2 and scale 1,500 hours. What share fail before 500 hours?

CHAPTER 22

Reliability and designed experiments

How long will it last, and which settings give the best result?

Weibull life analysis

Where: Reliability > Life Analysis > Weibull

Life data is different from dimensional data: it is time (or cycles, or miles) to failure, it is never negative, it is usually skewed, and some units have **not failed yet** when the test stops (suspended or censored units). The Weibull distribution handles all of that.

1. Choose the **time** column (hours, cycles).
2. Choose the **status** column and say which value means failed (here Failed = 1). The other units are treated as still running.
3. Click **Run**.

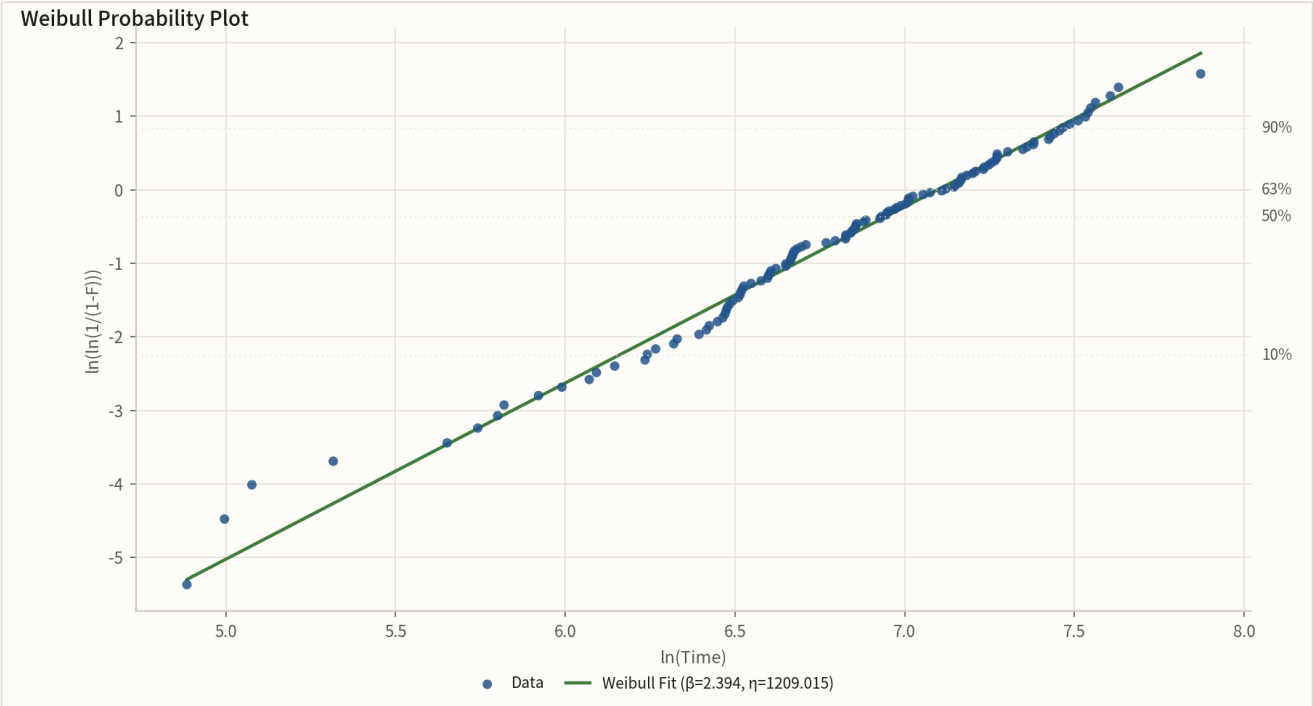


Figure 22.1 Weibull probability plot of 150 units on life test. Points close to the line mean the Weibull model fits.

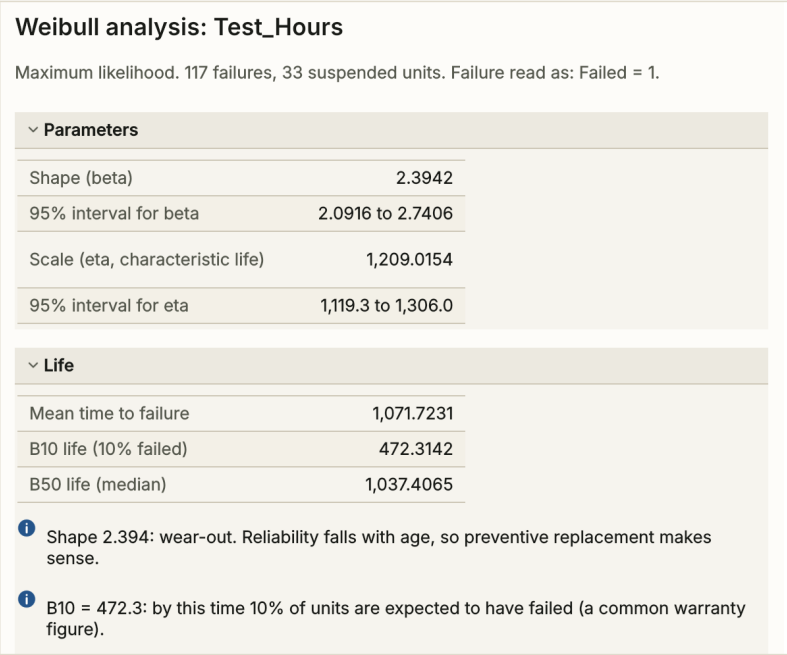


Figure 22.2 The Weibull parameters and life figures.

Number	Meaning	Here
Shape (beta)	Below 1: early failures (infant mortality, poor assembly). About 1: random failures. Above 1: wear-out.	2.3942: wear-out
Scale (eta)	The characteristic life: by this time 63.2% have failed.	1,209.0154 h
B10 life	The time by which 10% have failed: a common warranty and maintenance figure.	472.3142 h
Mean time to failure	The average life.	1,071.7231 h

In real life

Shape above 1 means the part wears out, so replacing it on a schedule before it fails makes sense: a preventive replacement at about the B10 life (470 hours here) avoids most failures. If the shape had been below 1, scheduled replacement would make things worse, because new parts are the ones failing; the fix then lies in assembly, burn-in or supplier quality.

Designed experiments (DOE)

Production data shows what goes with what. A **designed experiment** changes the inputs deliberately, in a planned pattern, so the effect of each one (and of combinations) can be separated cleanly with few runs.

Step	Command
Plan the runs	DOE tab: Full Factorial, Fractional, Screening, Response Surface and others create the run sheet as a new table

Step	Command
Run them in the listed (random) order and type the results in	the data sheet
See which factors matter	**DOE > Analyze > Effects Plot**
Test them	**DOE > Analyze > ANOVA**
Find the best settings	**DOE > Analyze > Profiler** and Contour Profiler

The example **DOE process optimization** varies temperature, pressure and speed and records the yield.

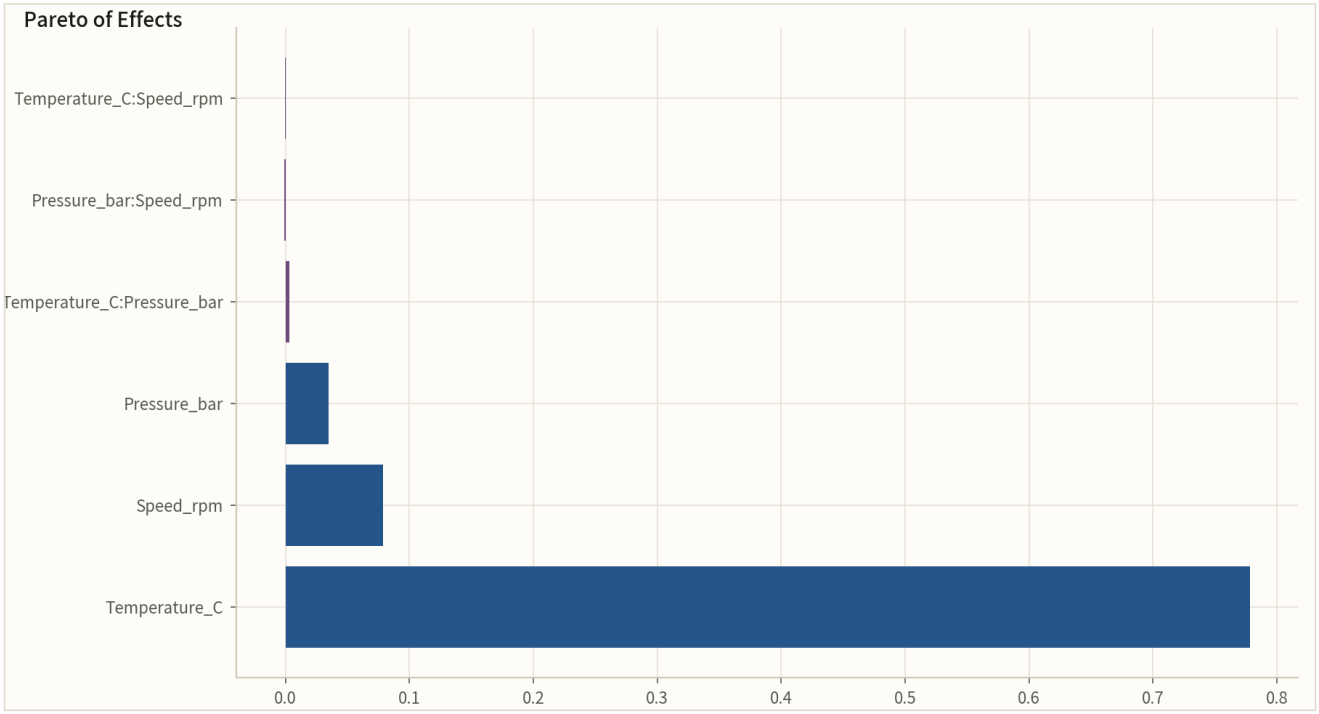


Figure 22.3 Pareto of effects on yield: temperature dominates; speed and pressure barely register; the interactions are negligible.

DOE ANOVA: Yield_Pct (interactions)

R^2	0.9116
Adj R^2	0.8674
F-statistic	20.6221
F p-value	0.0003

ANOVA Table

	sum_sq	df	F	PR(>F)
Temperature_C	964.9888	1	229.6454	0
Pressure_bar	652.5781	1	155.2988	0
Speed_rpm	88.9473	1	21.1674	0.0018
Temperature_C × Pressure_bar	0.0182	1	0.0043	0.9491
Temperature_C × Speed_rpm	19.4938	1	4.6391	0.0634
Pressure_bar × Speed_rpm	0.0181	1	0.0043	0.9493
Residual	33.6167	8		

Coefficients

Term	Coefficient	Std Error	t-value	p-value
Intercept	0.0022	0.0011	2.021	0.0779
Temperature_C	0.3891	0.2071	1.8784	0.0971
Pressure_bar	0.0173	0.0092	1.877	0.0974
Speed_rpm	0.0393	0.0469	0.8377	0.4265
Temperature_C × Pressure_bar	0.0017	0.0256	0.0659	0.9491
Temperature_C × Speed_rpm	-0.0001	0.0001	-2.1539	0.0634
Pressure_bar × Speed_rpm	-0.0004	0.0058	-0.0656	0.9493

Significant Terms (p<0.05)

No term is significant at the 5% level.

Figure 22.4 The effects table behind the chart.

Raising temperature from its low to its high setting changes yield by 0.7781, about ten times more than speed (0.0786). The practical conclusion: control temperature tightly; speed and pressure can be set for throughput and cost.

> **In real life**
> A heat-treat shop suspects time, temperature and quench delay all affect hardness. Eight runs of a two-level full factorial (2 x 2 x 2) answer it in a day. Changing one factor at a time would take longer and would never reveal that temperature only matters when the quench delay is long (an interaction).

PART V

Reports and records

Turning results into the documents customers, managers and auditors ask for: the step-by-step report builder, exports, the Excel calculator, accounts and the audit trail.

CHAPTER 23

Building a report, step by step

Pick the report you need from a list of the common SPC reports. The program fills in everything it can find and asks a quick question for anything it cannot.

Where: Reports > Generate > Templates

Step 1: choose the report

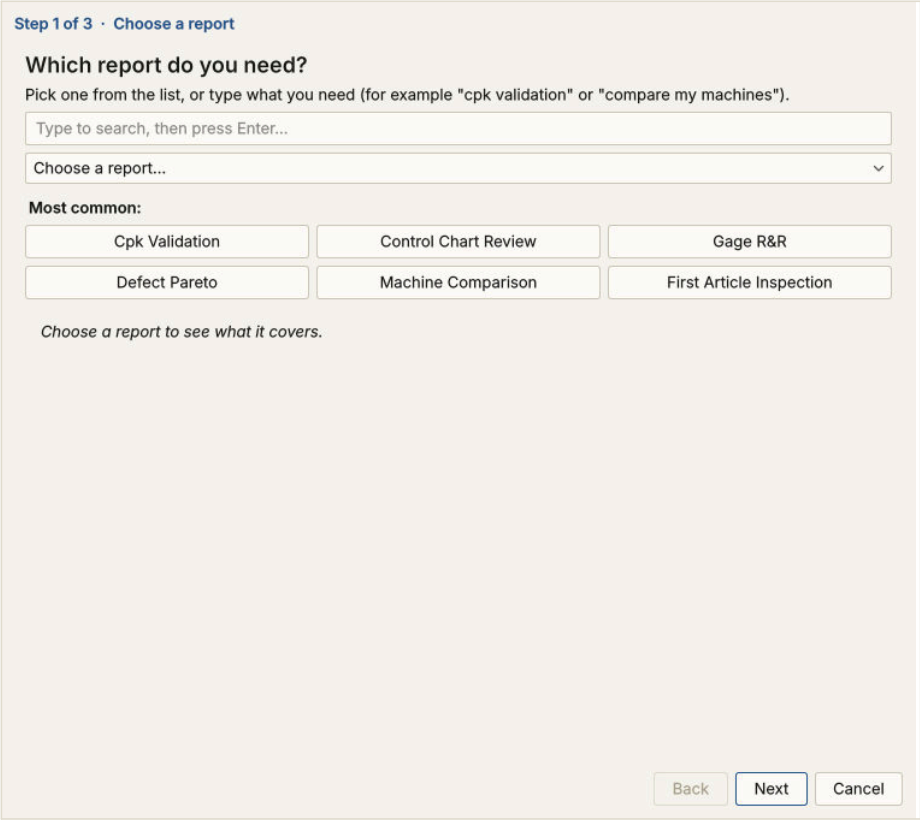


Figure 23.1 Step 1. The six most common reports are one click away; the drop-down holds all 23, grouped by kind.

There are three ways to choose:

- Click one of the **Most common** buttons.
- Open the **drop-down** and pick from the full list, grouped as Capability and validation, Control charts and monitoring, Measurement systems, Inspection and defects, Comparisons and improvement, Reliability and experiments, and Data and overview.
- **Type what you need** in the search box and press Enter: "cpk validation", "is my process capable", "compare my machines", "which defects happen most", "did the change work". The closest report is selected.

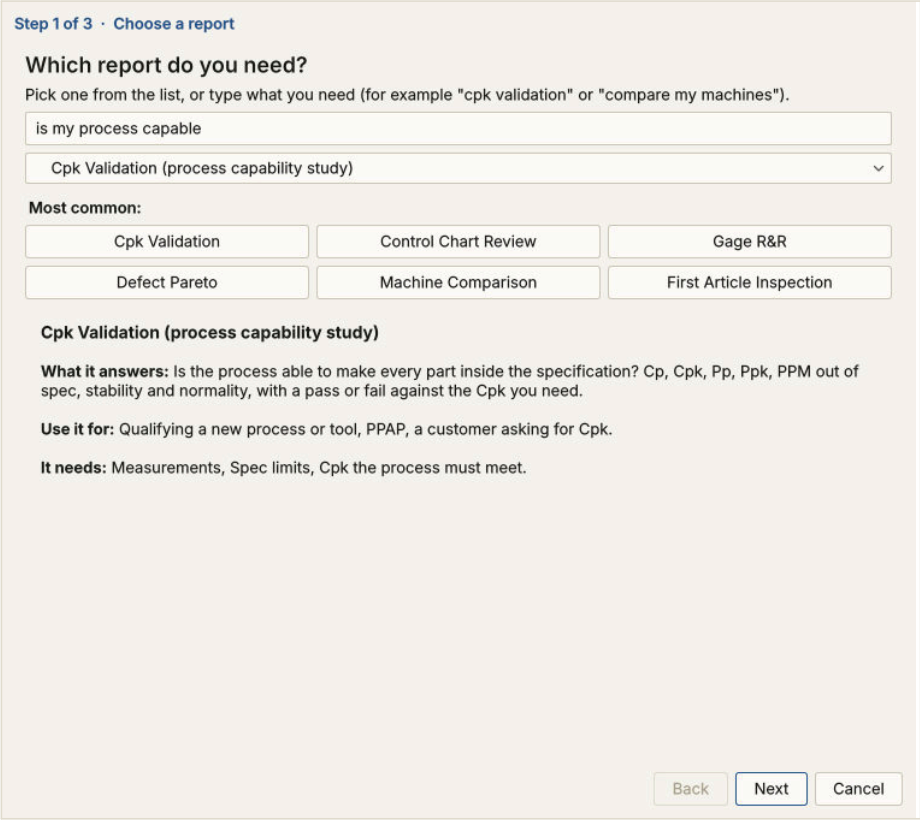


Figure 23.2 Typing "is my process capable" selects the Cpk Validation report. The card explains what it answers, when to use it and what it needs.

Step 2: Quick question!

The program looks through your data and settings for every answer the report needs: which columns are measurements, which column names the machine or operator, spec limits you already set, your user name. Only what it **could not** find is asked, at the top, under **Quick question!**

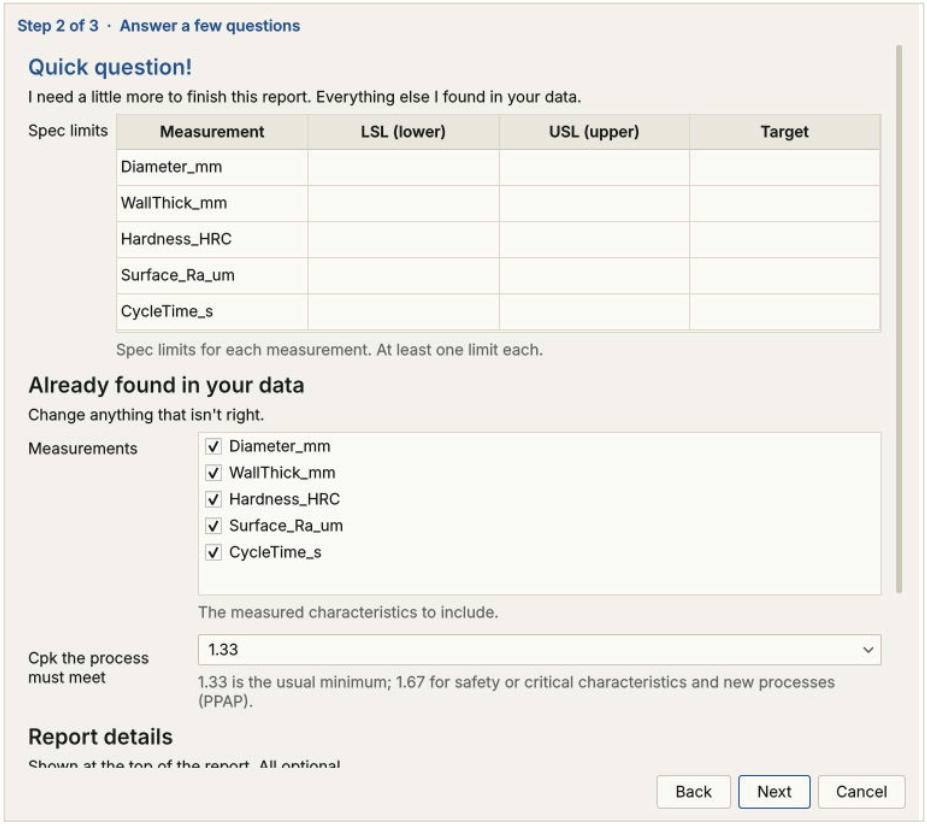


Figure 23.3 A Cpk Validation with no spec limits set yet: the only question is the limits. The measurements were found automatically.

Type the LSL and USL (and the target, if there is one) for each measurement. A one-sided spec, such as a maximum surface roughness, needs only one limit. Under **Already found in your data**, untick any measurement you don't want in the report, and choose the Cpk the process must meet: 1.33 for production, 1.67 for critical characteristics and new processes.

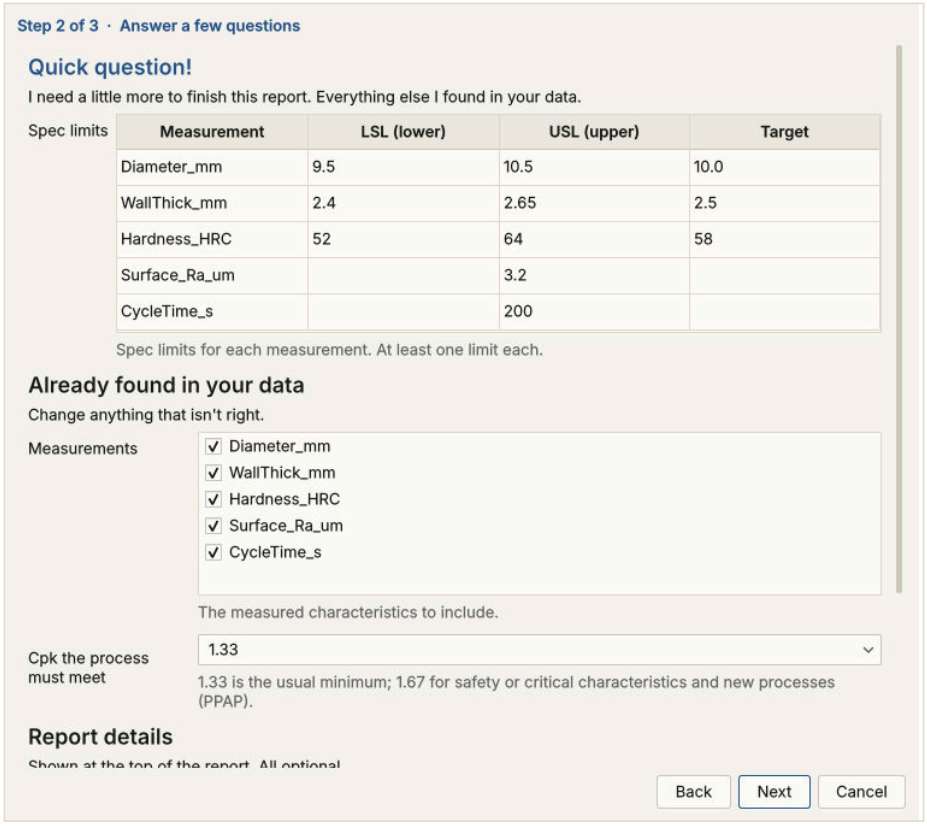


Figure 23.4 The limits typed in, with the part number and customer added under Report details.

Report details are printed in a panel at the top of the report: title, part number, customer, prepared by (your user name is filled in), and notes.

When a report needs nothing more, the page says so:

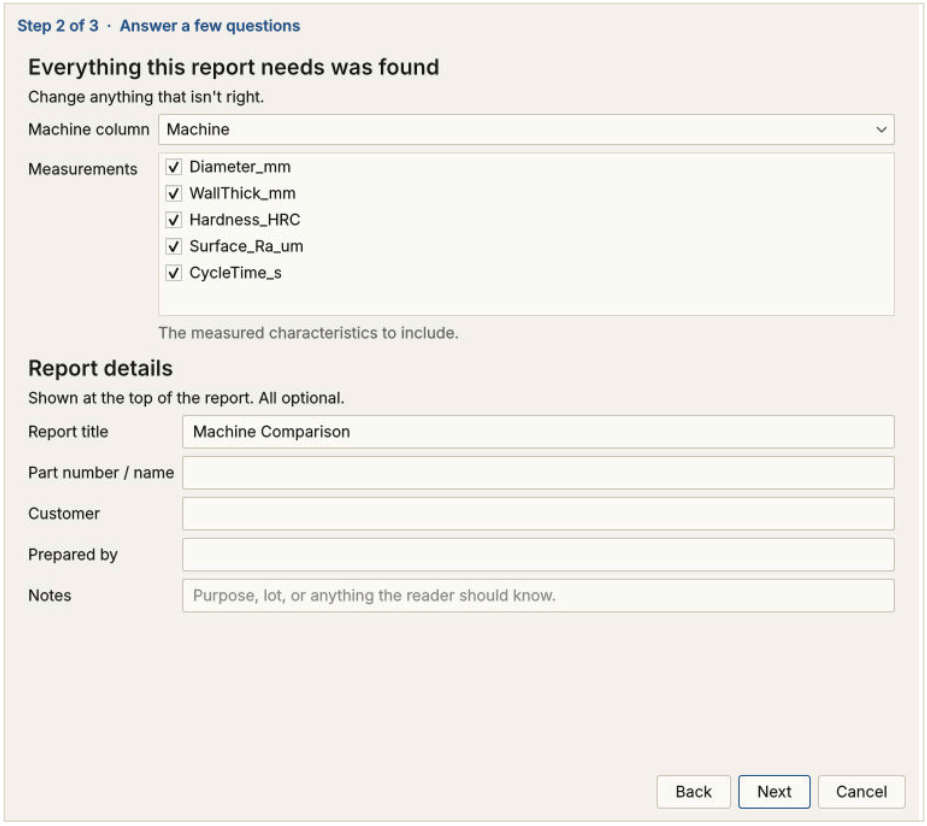


Figure 23.5 "compare my machines": the machine column and the measurements were all found, so there is no question to answer.

> **Tip**
> If an answer cannot work (a column name that isn't in the data, a lower limit above the upper one, a word where a number is needed), **Next** explains what to fix instead of building a broken report.

Step 3: review and build

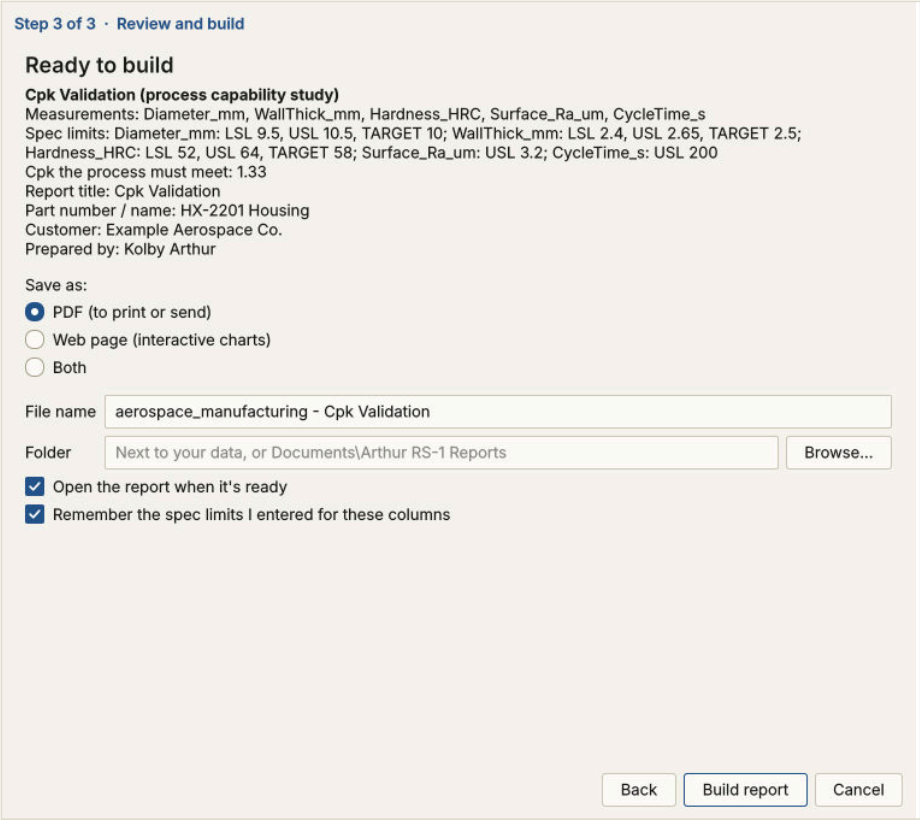
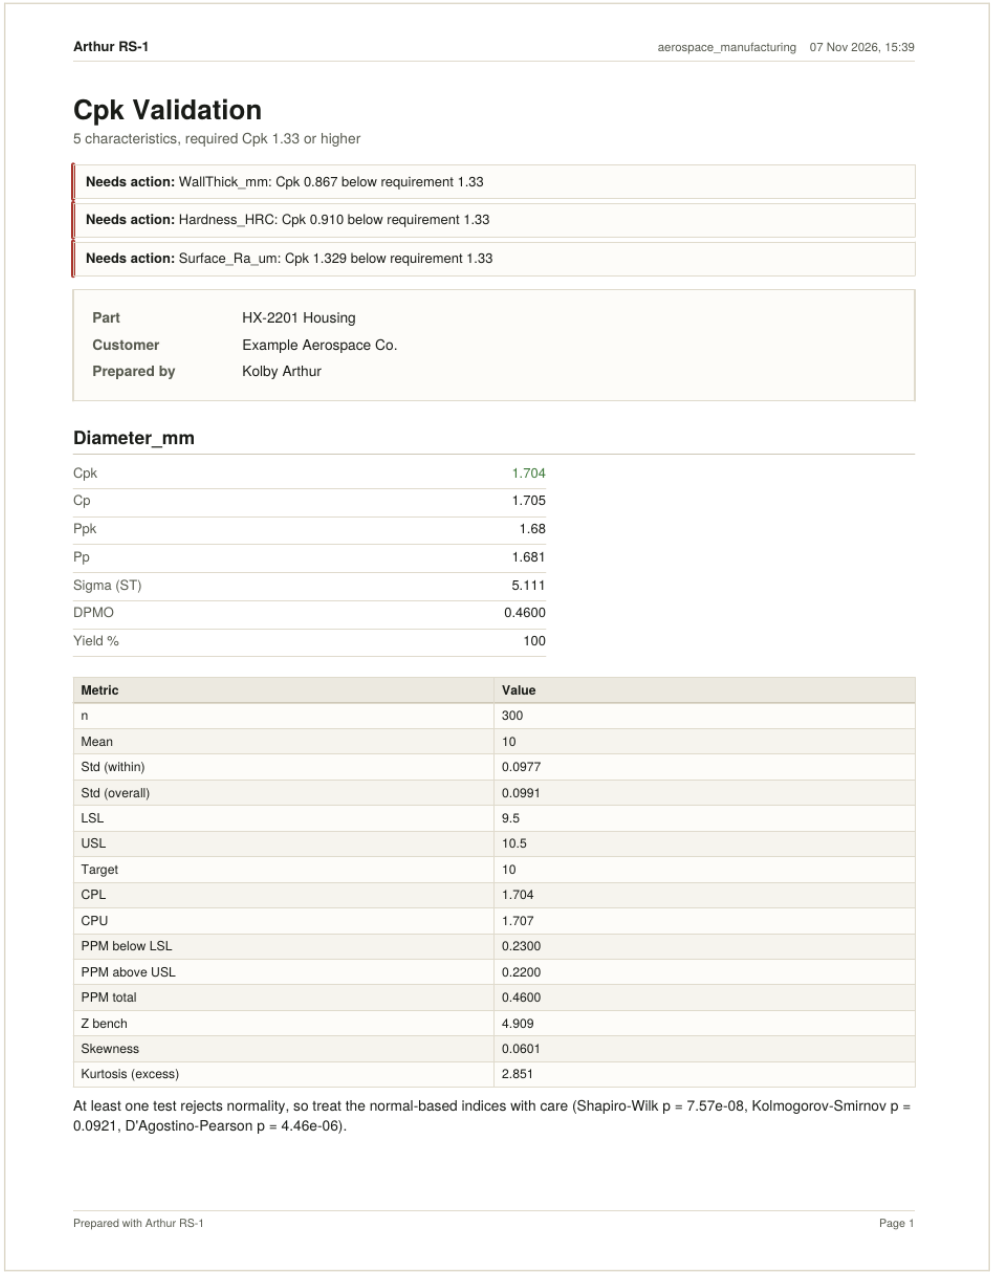


Figure 23.6 Step 3: everything the report will use, the file type, the file name and the folder.

- **PDF** to print, e-mail or attach to a PPAP; **Web page** for interactive charts you can zoom; or **Both**.
- The **file name** starts as the data's name plus the report's title. The **folder** defaults to the one your data came from, or Documents\RS-1 Reports.
- **Remember the spec limits** saves the limits you typed, so the next study or report on these columns doesn't ask again.

Click **Build report**. The report builds in the background (you can keep working) and opens when it is ready.

What you get



Arthur RS-1 aerospace_manufacturing 07 Nov 2026, 15:39

Cpk Validation

5 characteristics, required Cpk 1.33 or higher

Needs action: WallThick_mm: Cpk 0.867 below requirement 1.33

Needs action: Hardness_HRC: Cpk 0.910 below requirement 1.33

Needs action: Surface_Ra_um: Cpk 1.329 below requirement 1.33

Part	HX-2201 Housing
Customer	Example Aerospace Co.
Prepared by	Kolby Arthur

Diameter_mm

Cpk	1.704
Cp	1.705
Ppk	1.68
Pp	1.681
Sigma (ST)	5.111
DPMO	0.4600
Yield %	100

Metric	Value
n	300
Mean	10
Std (within)	0.0977
Std (overall)	0.0991
LSL	9.5
USL	10.5
Target	10
CPL	1.704
CPU	1.707
PPM below LSL	0.2300
PPM above USL	0.2200
PPM total	0.4600
Z bench	4.909
Skewness	0.0601
Kurtosis (excess)	2.851

At least one test rejects normality, so treat the normal-based indices with care (Shapiro-Wilk p = 7.57e-08, Kolmogorov-Smirnov p = 0.0921, D'Agostino-Pearson p = 4.46e-06).

Prepared with Arthur RS-1 Page 1

Figure 23.7 Page 1 of the Cpk Validation report: the verdicts first, then the report details, then each characteristic.

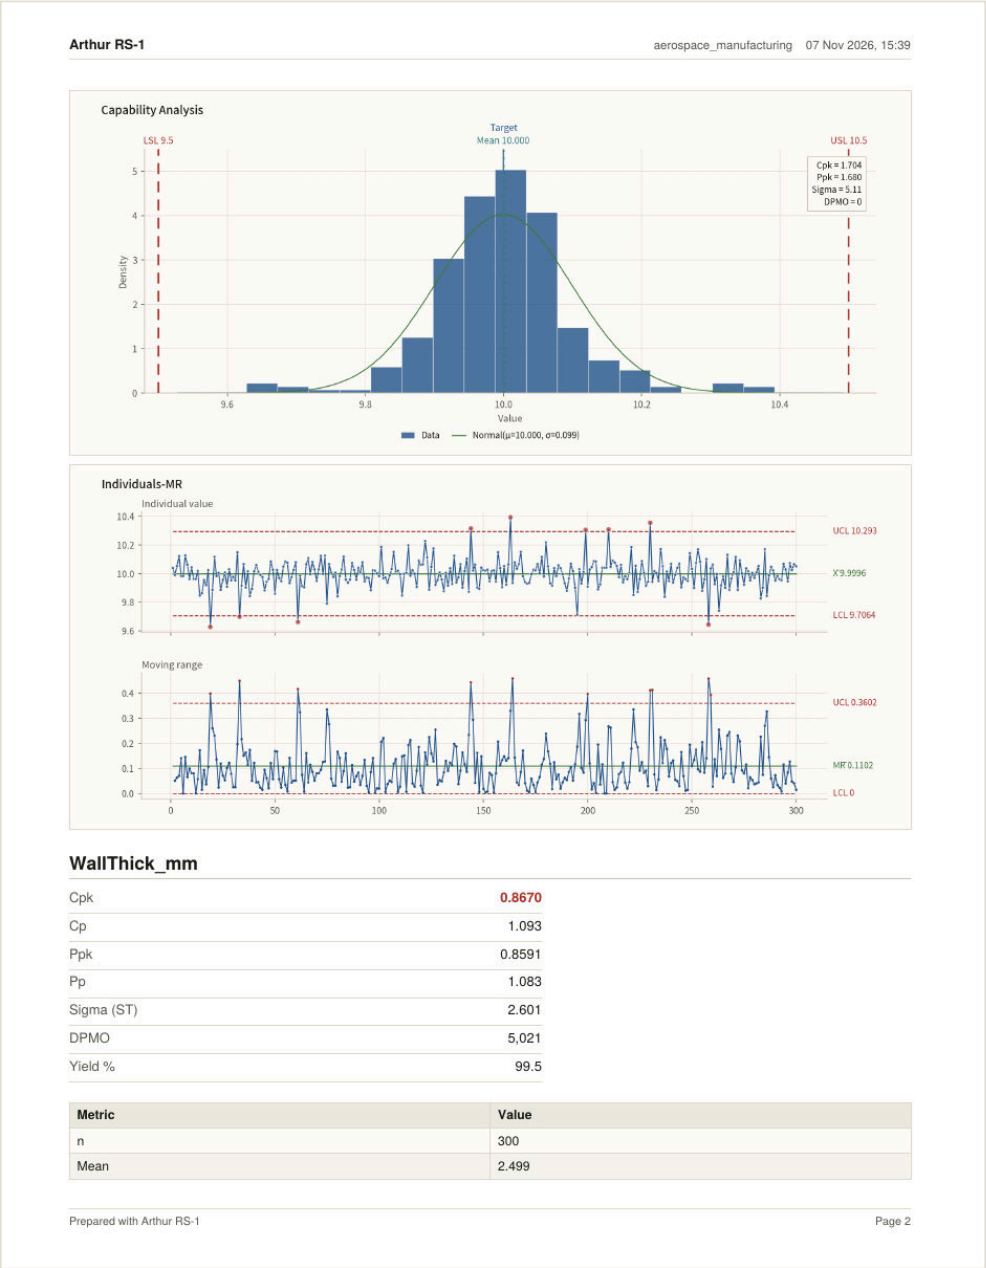


Figure 23.8 Inside: the capability histogram and I-MR chart for each characteristic, with its numbers.

The 23 reports

Group	Reports
Capability and validation	Cpk Validation; First Article Inspection; Executive Quality Summary
Control charts and monitoring	Control Chart Review; Trend and Drift; Process Dashboard
Measurement systems	Gage R&R; Attribute Agreement; Gage Bias and Linearity
Inspection and defects	Incoming Inspection / Lot Acceptance; Defect Pareto; Yield and DPMO

Group	Reports
Comparisons and improvement	Machine Comparison; Operator Comparison; Shift Quality Report; Before / After Improvement
Reliability and experiments	Reliability / Weibull Life; Designed Experiment (DOE); What Drives the Result (regression)
Data and overview	Normality and Outliers; Distribution Identification; Correlation Screen; Data Quality Audit

Which report for which job

Customer asks for Cpk on a new part: Cpk Validation, minimum 1.67, with the part number and customer filled in.

Monthly quality review: Executive Quality Summary for management, Defect Pareto and Yield and DPMO for the team.

Closing an 8D or DMAIC project: Before / After Improvement with a column marking the two periods.

New gauge on the line: Gage R&R, with the tolerance entered.

CHAPTER 24

Exporting results

Any results you have on screen can go out as a PDF, Excel workbook, PowerPoint deck or web page that look like one family.

Export what is on screen

Command	Produces
Reports > Export > PDF	A printable report of the session's results.
Reports > Export > Excel	A workbook: a cover sheet with contents, then one sheet per result with its tables and charts.
Reports > Export > PowerPoint	One slide per result, ready to present.
Reports > Export > HTML	A single web page that opens in any browser.
Reports > Export > Interactive HTML	A web page whose charts can be zoomed and hovered.

Each export first asks which **sections** to include:

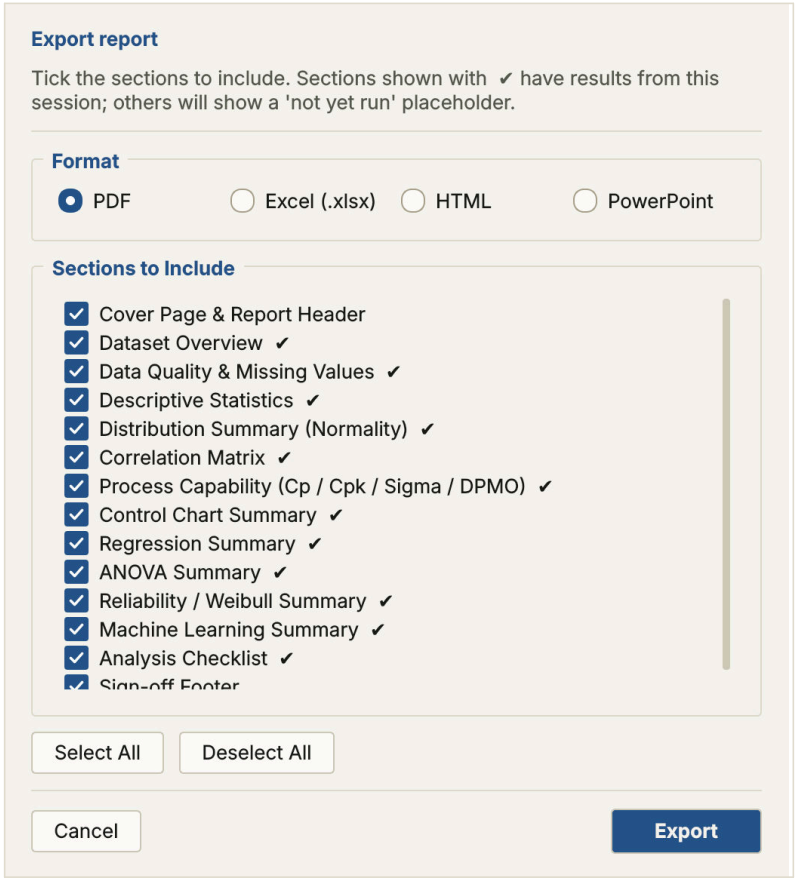


Figure 24.1 Choose the sections and the format.

All four formats share the PDF's look: the same colors, fonts, page header and footer, so a workbook and a slide deck from the same study read as one set of documents.

Report packs

Command	For
Reports > Generate > Executive	A short summary for management.
Reports > Generate > Manufacturing	Capability, control and MSA for production.
Reports > Generate > DMAIC	A Six Sigma project laid out as Define, Measure, Analyze, Improve and Control.
Reports > Generate > Reliability	Life data and reliability metrics.

Single charts

Every chart's toolbar has a **save** button (a picture file) and **Copy** (paste straight into an e-mail, Word or PowerPoint).

A statement of the standard

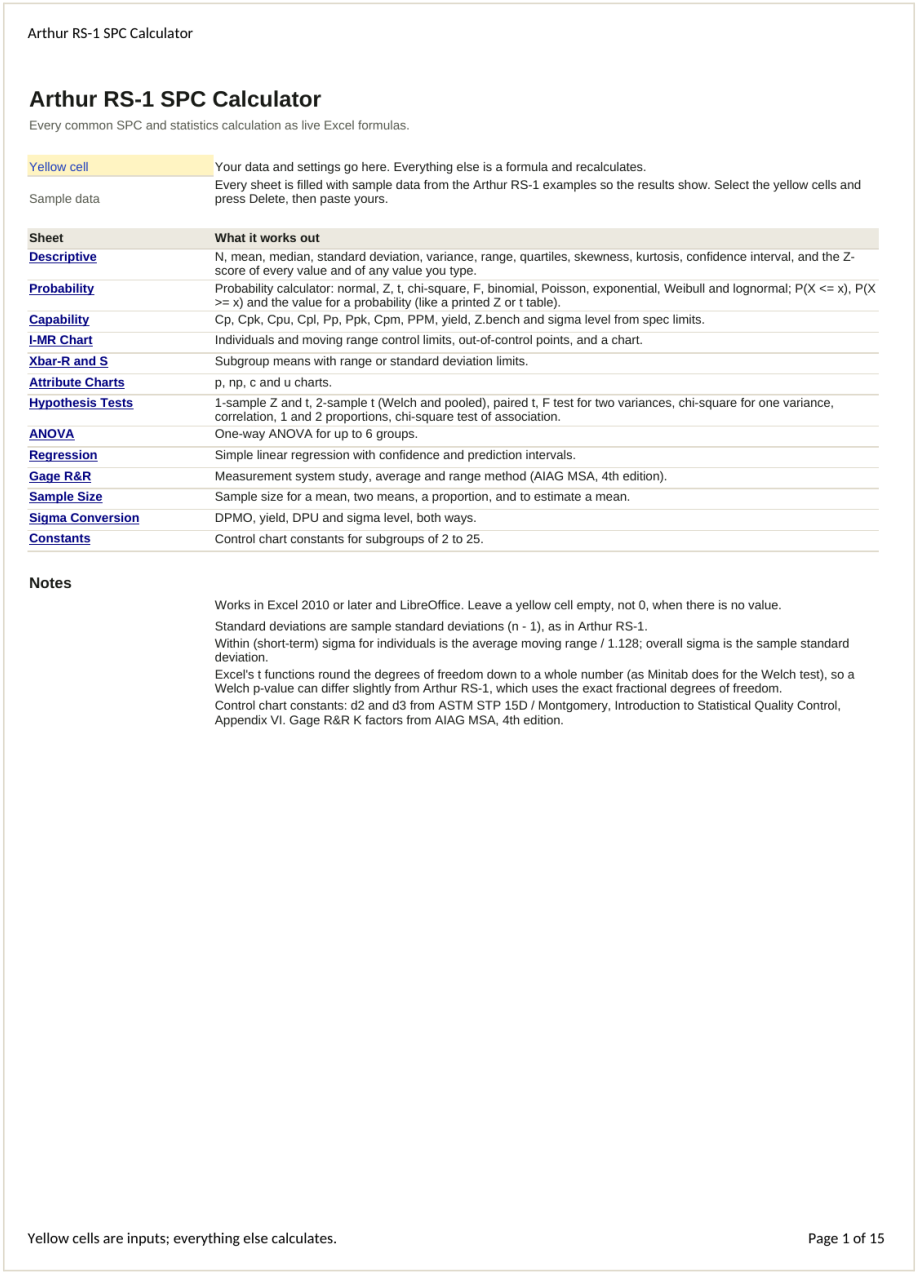
If your work is done under ISO 9001, AS9100 or IATF 16949, choose the standard once from the program menu (Tools > Compliance and traceability > Standard for reports). Every report then carries the matching statement.

CHAPTER 25

The Excel SPC calculator

A companion workbook that does every everyday SPC calculation with visible Excel formulas, for anyone without the program, and for checking the program's arithmetic by hand.

RS-1 SPC Calculator.xlsx works in Excel 2010 or later and in LibreOffice. Every sheet works the same way: **yellow cells are inputs**; everything else is a formula you can click on and read.



Arthur RS-1 SPC Calculator

Arthur RS-1 SPC Calculator

Every common SPC and statistics calculation as live Excel formulas.

Yellow cell	Your data and settings go here. Everything else is a formula and recalculates.
Sample data	Every sheet is filled with sample data from the Arthur RS-1 examples so the results show. Select the yellow cells and press Delete, then paste yours.

Sheet	What it works out
Descriptive	N, mean, median, standard deviation, variance, range, quartiles, skewness, kurtosis, confidence interval, and the Z-score of every value and of any value you type.
Probability	Probability calculator: normal, Z, t, chi-square, F, binomial, Poisson, exponential, Weibull and lognormal; P(X <= x), P(X >= x) and the value for a probability (like a printed Z or t table).
Capability	Cp, Cpk, Cpu, Cpl, Pp, Ppk, Cpm, PPM, yield, Z.bench and sigma level from spec limits.
I-MR Chart	Individuals and moving range control limits, out-of-control points, and a chart.
Xbar-R and S	Subgroup means with range or standard deviation limits.
Attribute Charts	p, np, c and u charts.
Hypothesis Tests	1-sample Z and t, 2-sample t (Welch and pooled), paired t, F test for two variances, chi-square for one variance, correlation, 1 and 2 proportions, chi-square test of association.
ANOVA	One-way ANOVA for up to 6 groups.
Regression	Simple linear regression with confidence and prediction intervals.
Gage R&R	Measurement system study, average and range method (AIAG MSA, 4th edition).
Sample Size	Sample size for a mean, two means, a proportion, and to estimate a mean.
Sigma Conversion	DPMO, yield, DPU and sigma level, both ways.
Constants	Control chart constants for subgroups of 2 to 25.

Notes

Works in Excel 2010 or later and LibreOffice. Leave a yellow cell empty, not 0, when there is no value.

Standard deviations are sample standard deviations (n - 1), as in Arthur RS-1.

Within (short-term) sigma for individuals is the average moving range / 1.128; overall sigma is the sample standard deviation.

Excel's t functions round the degrees of freedom down to a whole number (as Minitab does for the Welch test), so a Welch p-value can differ slightly from Arthur RS-1, which uses the exact fractional degrees of freedom.

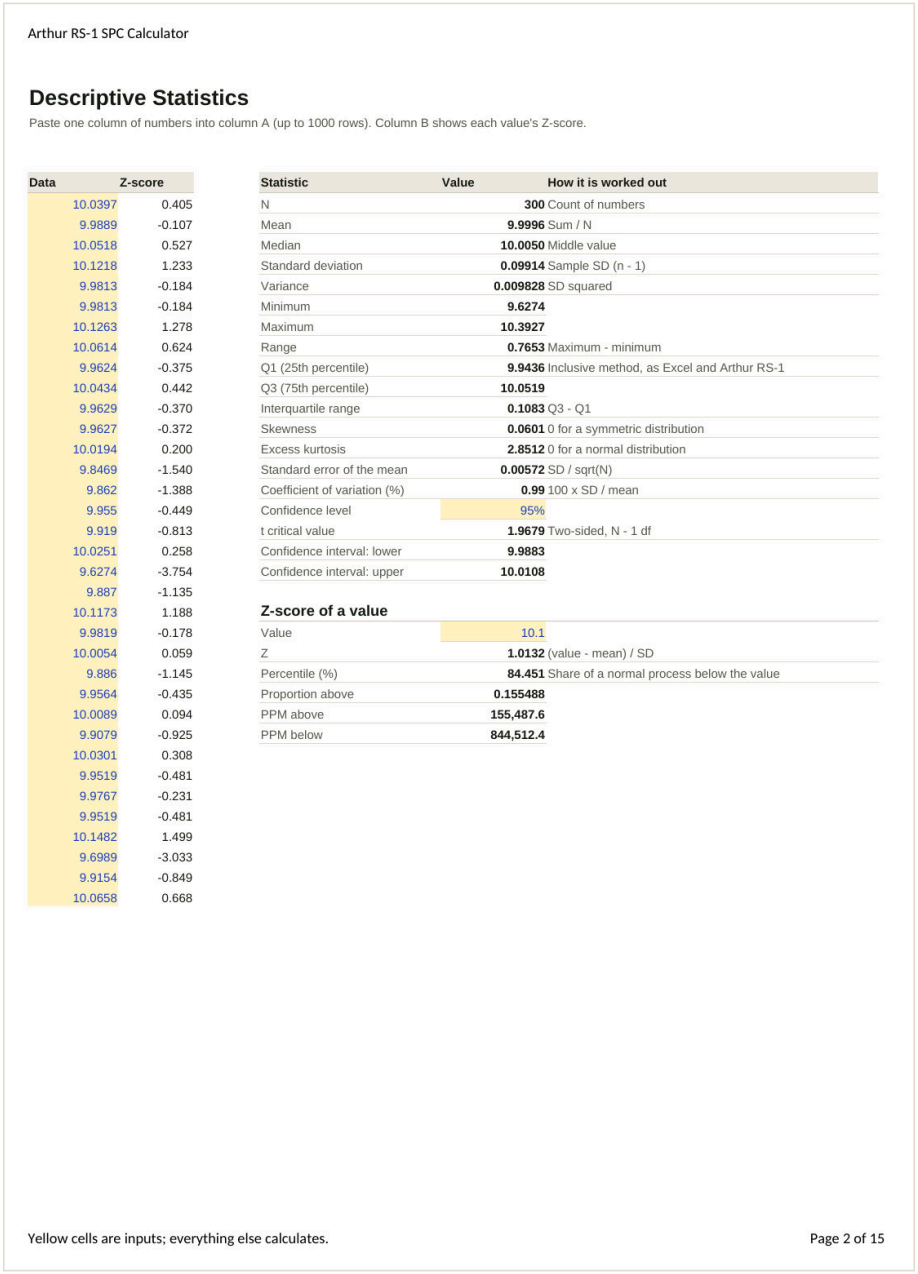
Control chart constants: d2 and d3 from ASTM STP 15D / Montgomery, Introduction to Statistical Quality Control, Appendix VI. Gage R&R K factors from AIAG MSA, 4th edition.

Yellow cells are inputs; everything else calculates. Page 1 of 15

Figure 25.1 The Start Here sheet lists every sheet and what it works out.

Sheet	Works out
Descriptive	Mean, median, standard deviation, quartiles, skewness, kurtosis, confidence interval, and the Z-score of any value
Probability	Normal, t, chi-square, F, exponential, Weibull, lognormal, binomial and Poisson probabilities
Capability	Cp, Cpk, Cpu, Cpl, Pp, Ppk, Cpm, PPM and Z bench
I-MR Chart	Individuals and moving range limits, out-of-control points, a chart
Xbar-R and S	Subgroup means with range or standard deviation limits
Attribute Charts	p, np, c and u charts

Sheet	Works out
Hypothesis Tests	1-sample Z and t, 2-sample t (Welch and pooled), paired t, F, chi-square tests, proportions
ANOVA	One-way ANOVA for up to six groups
Regression	Simple linear regression with confidence and prediction intervals
Gage R&R	Average and range method (AIAG MSA 4th edition)
Sample Size	For a mean, two means, a proportion and a margin of error
Sigma Conversion	DPMO, yield, DPU and sigma level, both ways
Constants	Control chart constants for subgroups of 2 to 25



Arthur RS-1 SPC Calculator

Descriptive Statistics

Paste one column of numbers into column A (up to 1000 rows). Column B shows each value's Z-score.

Data	Z-score
10.0397	0.405
9.9889	-0.107
10.0518	0.527
10.1218	1.233
9.9813	-0.184
9.9813	-0.184
10.1263	1.278
10.0614	0.624
9.9624	-0.375
10.0434	0.442
9.9629	-0.370
9.9627	-0.372
10.0194	0.200
9.8469	-1.540
9.862	-1.388
9.955	-0.449
9.919	-0.813
10.0251	0.258
9.6274	-3.754
9.887	-1.135
10.1173	1.188
9.9819	-0.178
10.0054	0.059
9.886	-1.145
9.9564	-0.435
10.0089	0.094
9.9079	-0.925
10.0301	0.308
9.9519	-0.481
9.9767	-0.231
9.9519	-0.481
10.1482	1.499
9.6989	-3.033
9.9154	-0.849
10.0658	0.668

Statistic	Value	How it is worked out
N	**300**	Count of numbers
Mean	**9.9996**	Sum / N
Median	**10.0050**	Middle value
Standard deviation	**0.09914**	Sample SD (n - 1)
Variance	**0.009828**	SD squared
Minimum	**9.6274**	
Maximum	**10.3927**	
Range	**0.7653**	Maximum - minimum
Q1 (25th percentile)	**9.9436**	Inclusive method, as Excel and Arthur RS-1
Q3 (75th percentile)	**10.0519**	
Interquartile range	**0.1083**	Q3 - Q1
Skewness	**0.0601**	0 for a symmetric distribution
Excess kurtosis	**2.8512**	0 for a normal distribution
Standard error of the mean	**0.00572**	SD / sqrt(N)
Coefficient of variation (%)	**0.99**	100 x SD / mean
Confidence level	95%	
t critical value	**1.9679**	Two-sided, N - 1 df
Confidence interval: lower	**9.9883**	
Confidence interval: upper	**10.0108**	

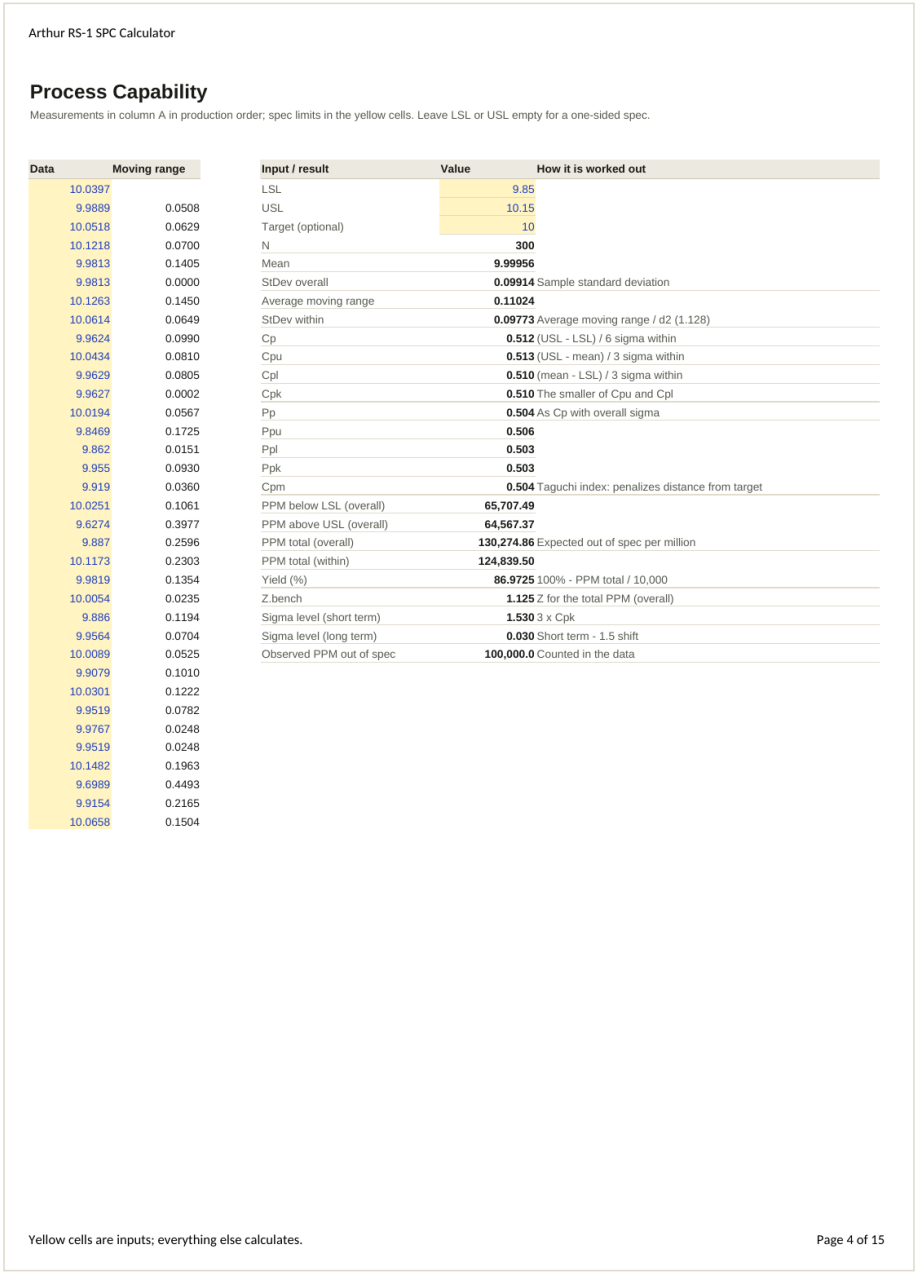
Z-score of a value

Value	10.1	
Z	**1.0132**	(value - mean) / SD
Percentile (%)	**84.451**	Share of a normal process below the value
Proportion above	**0.155488**	
PPM above	**155,487.6**	
PPM below	**844,512.4**	

Yellow cells are inputs; everything else calculates.

Page 2 of 15

Figure 25.2 Descriptive: paste a column of data into the yellow cells; the statistics and each value's Z-score follow.



Arthur RS-1 SPC Calculator

Process Capability

Measurements in column A in production order; spec limits in the yellow cells. Leave LSL or USL empty for a one-sided spec.

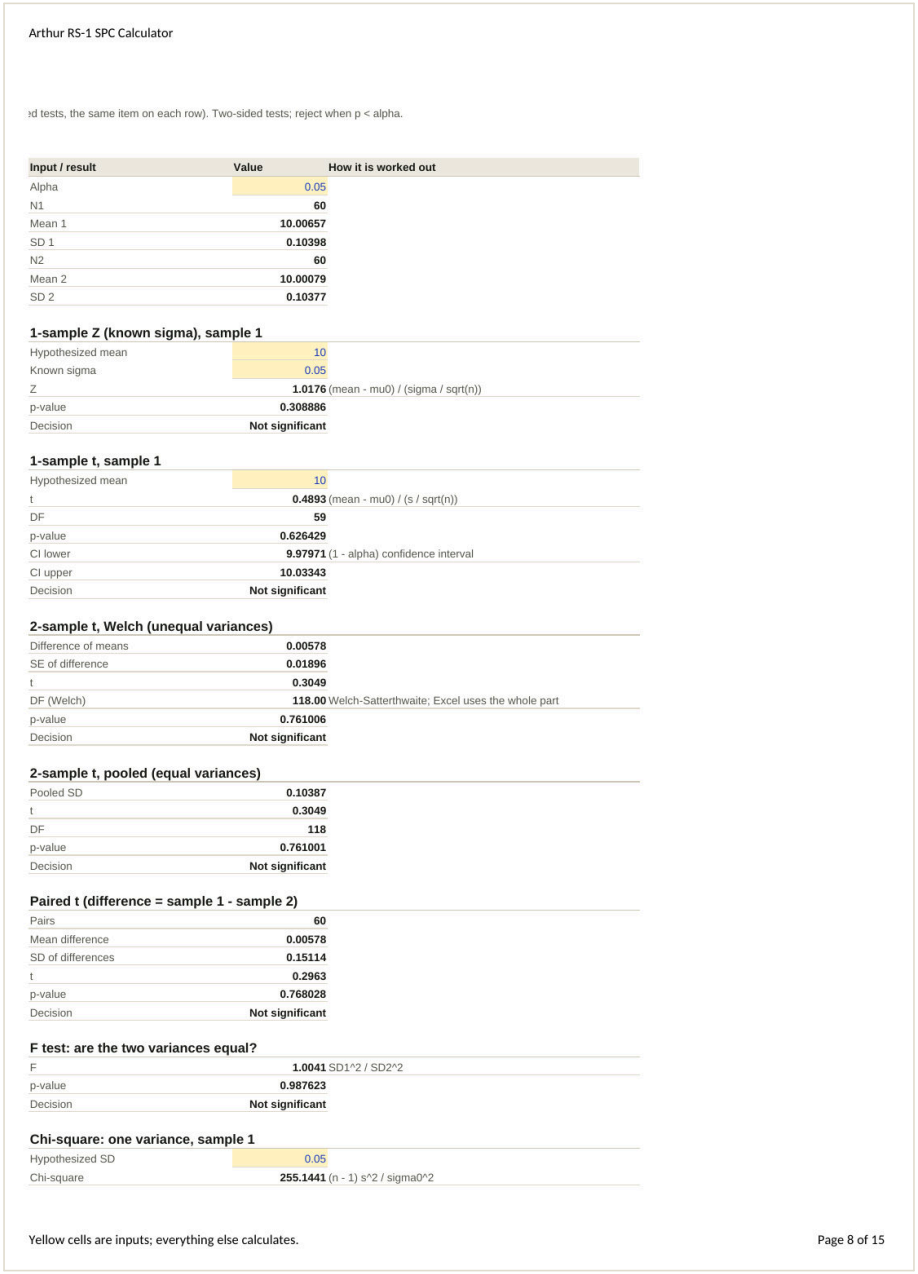
Data	Moving range
10.0397	
9.9889	0.0508
10.0518	0.0629
10.1218	0.0700
9.9813	0.1405
9.9813	0.0000
10.1263	0.1450
10.0614	0.0649
9.9624	0.0990
10.0434	0.0810
9.9629	0.0805
9.9627	0.0002
10.0194	0.0567
9.8469	0.1725
9.862	0.0151
9.955	0.0930
9.919	0.0360
10.0251	0.1061
9.6274	0.3977
9.887	0.2596
10.1173	0.2303
9.9819	0.1354
10.0054	0.0235
9.886	0.1194
9.9564	0.0704
10.0089	0.0525
9.9079	0.1010
10.0301	0.1222
9.9519	0.0782
9.9767	0.0248
9.9519	0.0248
10.1482	0.1963
9.6989	0.4493
9.9154	0.2165
10.0658	0.1504

Input / result	Value	How it is worked out
LSL	9.85	
USL	10.15	
Target (optional)	10	
N	**300**	
Mean	**9.99956**	
StDev overall	**0.09914**	Sample standard deviation
Average moving range	**0.11024**	
StDev within	**0.09773**	Average moving range / d2 (1.128)
Cp	**0.512**	(USL - LSL) / 6 sigma within
Cpu	**0.513**	(USL - mean) / 3 sigma within
Cpl	**0.510**	(mean - LSL) / 3 sigma within
Cpk	**0.510**	The smaller of Cpu and Cpl
Pp	**0.504**	As Cp with overall sigma
Ppu	**0.506**	
Ppl	**0.503**	
Ppk	**0.503**	
Cpm	**0.504**	Taguchi index: penalizes distance from target
PPM below LSL (overall)	**65,707.49**	
PPM above USL (overall)	**64,567.37**	
PPM total (overall)	**130,274.86**	Expected out of spec per million
PPM total (within)	**124,839.50**	
Yield (%)	**86.9725**	100% - PPM total / 10,000
Z.bench	**1.125**	Z for the total PPM (overall)
Sigma level (short term)	**1.530**	3 x Cpk
Sigma level (long term)	**0.030**	Short term - 1.5 shift
Observed PPM out of spec	**100,000.0**	Counted in the data

Yellow cells are inputs; everything else calculates.

Page 4 of 15

Figure 25.3 Capability: readings in production order on the left, limits in the yellow cells, every index with the formula explained.



Arthur RS-1 SPC Calculator

ed tests, the same item on each row). Two-sided tests; reject when p < alpha.

Input / result	Value	How it is worked out
Alpha	0.05	
N1	60	
Mean 1	10.00657	
SD 1	0.10398	
N2	60	
Mean 2	10.00079	
SD 2	0.10377	

1-sample Z (known sigma), sample 1

Hypothesized mean	10	
Known sigma	0.05	
Z	1.0176	(mean - mu0) / (sigma / sqrt(n))
p-value	0.308886	
Decision	Not significant	

1-sample t, sample 1

Hypothesized mean	10	
t	0.4893	(mean - mu0) / (s / sqrt(n))
DF	59	
p-value	0.626429	
CI lower	9.97971	(1 - alpha) confidence interval
CI upper	10.03343	
Decision	Not significant	

2-sample t, Welch (unequal variances)

Difference of means	0.00578	
SE of difference	0.01896	
t	0.3049	
DF (Welch)	118.00	Welch-Satterthwaite; Excel uses the whole part
p-value	0.761006	
Decision	Not significant	

2-sample t, pooled (equal variances)

Pooled SD	0.10387	
t	0.3049	
DF	118	
p-value	0.761001	
Decision	Not significant	

Paired t (difference = sample 1 - sample 2)

Pairs	60	
Mean difference	0.00578	
SD of differences	0.15114	
t	0.2963	
p-value	0.768028	
Decision	Not significant	

F test: are the two variances equal?

F	1.0041	SD1^2 / SD2^2
p-value	0.987623	
Decision	Not significant	

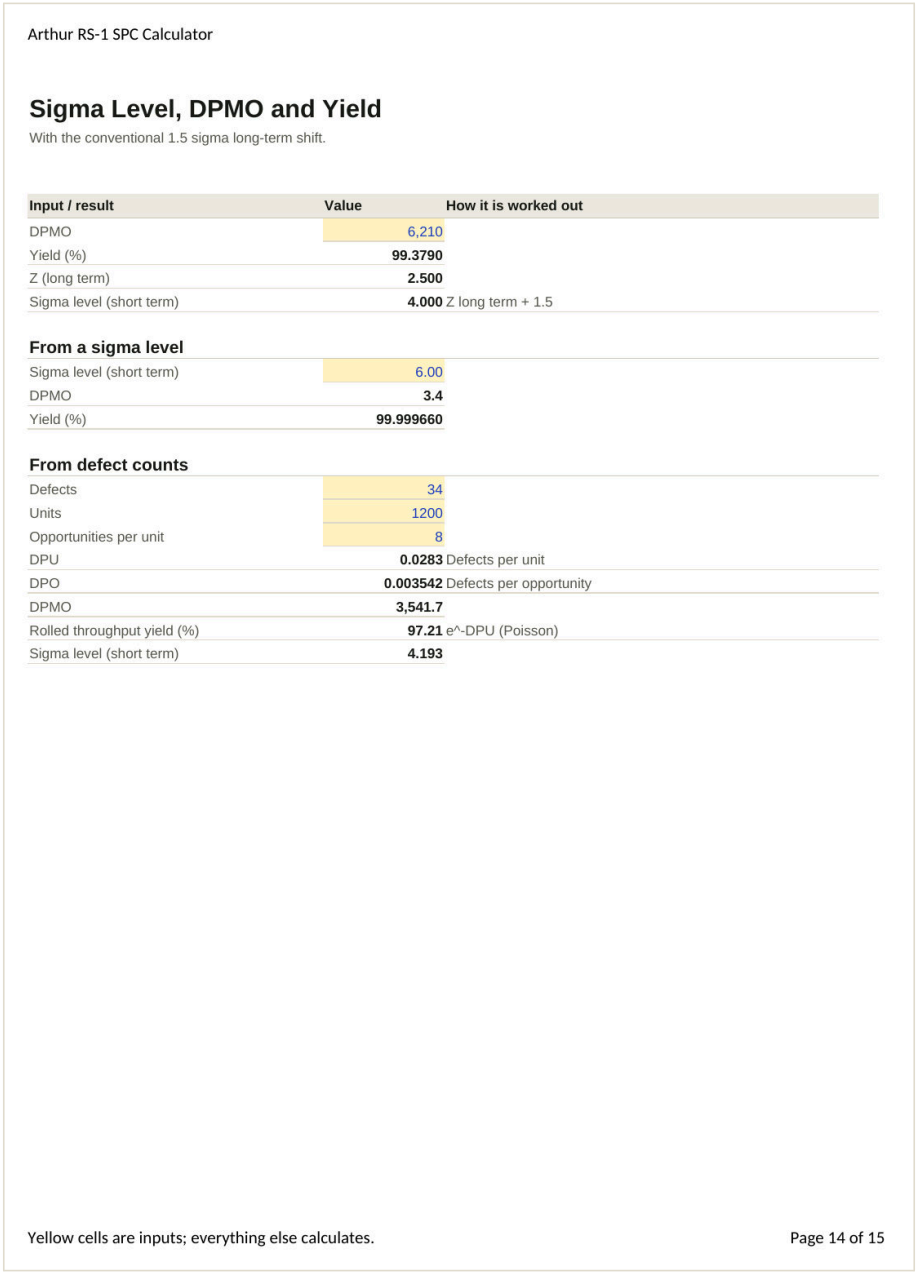
Chi-square: one variance, sample 1

Hypothesized SD	0.05	
Chi-square	255.1441	(n - 1) s^2 / sigma0^2

Yellow cells are inputs; everything else calculates.

Page 8 of 15

Figure 25.4 Hypothesis Tests: each test in its own block, with its decision.



Arthur RS-1 SPC Calculator

Sigma Level, DPMO and Yield

With the conventional 1.5 sigma long-term shift.

Input / result	Value	How it is worked out
DPMO	6,210	
Yield (%)	**99.3790**	
Z (long term)	**2.500**	
Sigma level (short term)	**4.000**	Z long term + 1.5
From a sigma level		
Sigma level (short term)	6.00	
DPMO	**3.4**	
Yield (%)	**99.999660**	
From defect counts		
Defects	34	
Units	1200	
Opportunities per unit	8	
DPU	**0.0283**	Defects per unit
DPO	**0.003542**	Defects per opportunity
DPMO	**3,541.7**	
Rolled throughput yield (%)	**97.21**	e^-DPU (Poisson)
Sigma level (short term)	**4.193**	

Yellow cells are inputs; everything else calculates.

Page 14 of 15

Figure 25.5 Sigma Conversion: from DPMO to sigma level and back.

> **Tip**
> The workbook ships filled with the aerospace example data, so every number can be compared with the program's results in this guide. Paste your own data over the yellow cells.

> **Watch out**
> Excel's own functions round degrees of freedom differently in a few tests (for example Welch's t-test), so the last digit of a p-value can differ slightly from RS-1. The decision will be the same.

CHAPTER 26

Accounts, settings and the audit trail

Who did what, and when: the records that make results defensible.

Accounts

- Each person should have their own account (Chapter 2). The first account on a computer is the administrator.
- New accounts are made from the sign-in window with **Create an account**.
- Every account has a one-time **recovery code** for resetting a forgotten password. Keep it somewhere safe.

The audit trail

RS-1 records every significant action: data opened, analyses run, limits saved and unlocked, causes recorded, reports built. Each entry carries the user and the time.

Program menu > Tools > Compliance and traceability >	Does
Audit trail	Shows the log.
Verify the audit log	Proves the log has not been altered: every entry is chained to the one before it, so changing or deleting any entry breaks the chain. The check reports PASS or FAIL, and gives a **seal** to record outside the program.
Standard for reports	The compliance statement printed on reports.

> **In real life**
> For an audit, run Verify the audit log, print the result with its seal, and file it with the reports. The auditor can see the records are complete and unaltered.

Settings and layout

The gear button on the top row opens:

Item	Use
Settings...	Seven pages: Appearance, Numbers (decimals, p-values), Analysis, Charts, Data and files, Ribbon and panels, and Updates.
Text size...	Larger or smaller text everywhere. Numbers are never squeezed out of sight at large sizes; panels wrap or scroll instead.

Item	Use
Customize the ribbon…	Add, remove, rename and reorder tabs, groups and buttons. Reset the ribbon returns to the built-in layout.
Lock panels	Stops the columns and results panels being dragged by accident.
Reset the layout (Ctrl+L)	Puts every panel back where it started.

Saving your work

- **Ctrl+S** (Save data as…) saves the data sheet as CSV or Excel.
- Save session… (the arrow beside the save button) saves the data, spec limits, column notes and results together, to reopen later exactly as you left them.
- RS-1 also autosaves in the background and offers to recover work after a crash or power cut.

CHAPTER 27

Troubleshooting

When something does not work, the program says why in plain words. Here are the messages you are most likely to meet and what to do.

Message (or symptom)	What it means	Do
Please load a dataset first	No data is open.	Open a file or an example (Chapter 4).
At least one spec limit (LSL or USL) is required	A capability study needs the drawing's limits.	**Quality > Capability > Set Spec Limits**, or type them in the study's form.
A column chosen as numeric contains text	Some cells hold words or units ("10.2 mm", "n/a").	Clean the column: **Data > Clean**, or retype the offending cells.
Some rows have missing values in the chosen columns	Empty cells in a column the analysis needs.	**Data > Clean > Fill Missing**, or choose other columns.
The same column was chosen for two roles	For example the part column also chosen as the operator.	Choose a different column for one role.
There are more groups than rows	The grouping column is almost unique per row (an ID).	Choose a grouping column with few values.
That is not a database address Arthur can read	The SQL import address is malformed.	Use the form **sqlite:///C:/path/file.db** or your server's address.
A red point you cannot explain	A special cause, or a typing or gauge error.	Check the reading first, then the context columns (Chapter 8).
Every point is red on a p chart	Very large samples make the limits too narrow.	Use Laney p' (Chapter 12).
Cpk looks good but customers complain	The process is not stable, or the gauge is poor.	Look at the control chart and Ppk (Chapter 15); run a Gage R&R (Chapter 17).
The report seemed to do nothing	Reports build in the background.	Watch the status bar; the report opens (or a box offers to open its folder) when done.

When something really is broken

If an unexpected error occurs, RS-1 keeps running, shows a message saying what went wrong, and writes the technical details to its log. **Home > Help > Support** shows how to send the log to Reliability Solutions. Nothing is sent without you choosing to.

> **Tip**
> Press **F1** on any command to open its on-screen guide page: what it does, what data it needs and an example.

Reference

A glossary of the terms in this guide, and every button on the ribbon with where to find it.

CHAPTER 28

Glossary

Term	Meaning
Alpha (α)	The risk of a false alarm you accept in a test, usually 0.05.
ANOVA	Analysis of variance: tests whether the means of several groups differ.
Attribute data	Counts or pass/fail results rather than measurements.
Baseline	A known good period whose readings set the control limits.
Bias	The average difference between a gauge's readings and the true value.
B10 life	The time by which 10% of units are expected to fail.
Capability	How well a stable process fits inside its tolerance.
Center line	The average plotted on a control chart.
Common cause	The everyday random variation built into a process.
Confidence interval	A range that contains the true value with a stated confidence, such as 95%.
Control limits	Lines 3 sigma either side of the center line, calculated from the data. They show how far a stable process normally strays.
Cp	Tolerance width divided by six within-sigma: potential capability if centered.
Cpk	Distance from the mean to the nearer spec limit divided by three within-sigma: actual capability.
CUSUM	Cumulative sum chart: running totals of deviations from target, for small, lasting shifts.
Defect / defective	A defect is a flaw; a defective is an item with one or more flaws that fails.
DNOM	Deviation from nominal: a short-run chart plotting each part's distance from its own nominal.
DOE	Design of experiments: changing inputs in a planned pattern to measure their effects.
DPMO / PPM	Defects per million opportunities / parts per million.
EWMA	Exponentially weighted moving average chart: a smoothed line that catches small, lasting shifts.
Gage R&R	A study of measurement error: repeatability (the gauge) and reproducibility (the people).
Histogram	Bars showing how many readings fall in each size range.
I-MR chart	Individuals and moving range chart, for one reading at a time.
In control / stable	Showing only common-cause variation: predictable.
Kappa	Agreement beyond chance between inspectors: 1 is perfect, 0 is chance.
LSL / USL	Lower and upper specification limits, from the drawing.
Mean	The average.

Term	Meaning
Moving range	The difference between a reading and the one before it.
ndc	Number of distinct categories a gauge can separate parts into; 5 or more is needed.
Normal distribution	The symmetric bell-shaped curve many measurements follow.
Out of control	Breaking one of the alarm rules: a sign of a special cause.
p-value	The chance of seeing a result at least this extreme if nothing were really going on. Small (below 0.05) means probably real.
Pareto	Categories sorted from most to least frequent, with a running total.
Phase	A period of the process (before / after a change) given its own limits.
Pp, Ppk	Like Cp and Cpk but with the overall sigma: long-term performance.
Pre-control	Green / yellow / red zones from the tolerance with simple run or stop rules.
R^2	The share of the variation in a result explained by a regression model.
Range	Largest minus smallest value in a group.
Rational subgroup	A small group of parts made under the same conditions.
Regression	A model of how a result changes with one or more inputs.
Repeatability	Variation when one person measures the same part several times with one gauge.
Reproducibility	Variation between different people measuring the same parts.
Run chart	Readings plotted in time order around the median.
Sigma (σ)	The standard deviation; also used for the sigma level of a process.
Sigma level	How many standard deviations fit between the mean and the nearer limit.
Special cause	A specific event outside the normal process: worn tool, new lot, wrong setting.
Specification (tolerance)	The range of acceptable values set by the design.
Standard deviation	The typical distance of readings from their mean.
Subgroup	A small group of readings taken together and charted as one point.
Weibull	A flexible distribution for life data; its shape tells early-life, random or wear-out failure.
Within / overall sigma	Short-term spread (from ranges) and total spread (plain standard deviation).
Xbar-R / Xbar-S chart	Charts of subgroup means with their ranges or standard deviations.
Z-score	How many standard deviations a value is from the mean.
Zones A, B, C	Bands 2-3, 1-2 and 0-1 sigma from the center line, used by the pattern rules.

CHAPTER 29

Command finder

Every button on the built-in ribbon, tab by tab. Remember you can also press Ctrl+E and type a command's name.

The Home tab

Group	Button	What it does
Get Data	**Examples**	Open one of the bundled example datasets.
Get Data	**Paste**	Paste a copied Excel or spreadsheet table directly into Arthur.
Get Data	**Live Data**	Keep a dataset in step with an Excel file or a Google Sheet.
Get Data	**Recent**	Reopen a file you used recently.
Analyze	**Automatic Analysis**	Let Arthur read the data and run the analyses it calls for.
Analyze	**Summary**	Compute N, mean, median, std, min, max, range, Q1, Q3, skewness, kurtosis for all numeric columns.
Analyze	**Z-Score**	How many standard deviations a value is from the mean.
Analyze	**Capability**	Measure how well your process fits within specification limits.
Analyze	**Control Chart**	Control chart for individual measurements with no subgroups.
Analyze	**2-Sample t**	Compare the means of two groups.
Analyze	**ANOVA**	Test whether means differ across 3+ groups. Includes Tukey HSD post-hoc.
Analyze	**Which Analysis?**	Find the right analysis for your question.
Measure	**Gage R&R**	Measurement System Analysis -- quantify how much variation comes from the measurement process itself.
Measure	**Attribute MSA**	How consistently inspectors judge pass/fail.
Measure	**Bias / Linearity**	Does the gauge read high or low, and does that change across the range?
Help	**Guides**	These guides.
Help	**Validation**	Evidence that Arthur's statistics are correct.
Help	**User Guide**	The printable user guide.
Help	**Support**	Contact details and a way to send a problem report.

The Explore tab

Group	Button	What it does
Understand	**Summary**	Compute N, mean, median, std, min, max, range, Q1, Q3, skewness, kurtosis for all numeric columns.
Understand	**Distribution**	Histogram, box plot and summary of one column.
Understand	**Profile Report**	A one-page profile of every column.
Understand	**Missing Values**	Visual dashboard showing which columns have missing data and how severe it is.
Understand	**Column Details**	Details of the selected column.
Relationships	**Correlation**	Pairwise Pearson correlations between all numeric variables, with strong pairs highlighted.
Relationships	**Rank Correlation**	Spearman or Kendall correlation between columns.
Relationships	**Covariance**	Covariance and correlation matrices for selected numeric columns.
Relationships	**Fit Y by X**	Scatter plot with multiple fitted curves and 95% density ellipse overlaid simultaneously.
Relationships	**Crosstabs**	Contingency table of two categorical variables with chi-square test.
Relationships	**Scatter Matrix**	Every pair of columns as scatter plots in one grid.
Distributions	**Normality**	Test whether data follows a normal distribution (Shapiro-Wilk).
Distributions	**Probability Plot**	Probability plot with the Anderson-Darling normality test.
Distributions	**Fit Distribution**	Fit and rank multiple probability distributions by AIC, overlay top 3 on histogram.
Distributions	**Quantiles**	Detailed quantile breakdown of a numeric column.
Distributions	**Probability**	Probabilities and values for normal, t, chi-square, F, binomial, Poisson and other distributions.
Distributions	**Z-Score**	How many standard deviations a value is from the mean.
Distributions	**Confidence Interval**	Compute 95% confidence interval for the population mean.
Distributions	**Tolerance Interval**	An interval that covers a stated share of the population.
Distributions	**Bootstrap CI**	Non-parametric confidence interval for any statistic, regardless of distribution.
Patterns	**Outlier Tests**	Grubbs' test or the generalized ESD test for outliers.
Patterns	**Flag Outliers**	Flag unusual values using Z-score, IQR, and Isolation Forest methods.
Patterns	**Runs Test**	Check whether a sequence is random.
Patterns	**Frequency (FFT)**	Decompose a signal into frequency components using Fast Fourier Transform.
Time Series	**Run Chart**	Values over order with a median reference line.
Time Series	**ACF / PACF**	Autocorrelation and Partial Autocorrelation plots for diagnosing ARIMA order.

Group	Button	What it does
Time Series	**Decompose**	Separate a time series into trend, seasonal, and residual components.
Time Series	**Smoothing**	Smooth a series and forecast ahead.
Time Series	**ADF Test**	Augmented Dickey-Fuller test for a trend-free (stationary) series.
Time Series	**KPSS Test**	Test whether a series is stationary.
Time Series	**Cross-Correlation**	Does one series lead another?

The Quality tab

Group	Button	What it does
Capability	**Capability**	Measure how well your process fits within specification limits.
Capability	**Capability Report**	Stability, normality and capability of one measurement on one page.
Capability	**Set Spec Limits**	Enter the specification limits for a column.
Capability	**View Spec Limits**	List the spec limits entered for this dataset.
Capability	**Non-Normal Cpk**	Capability analysis for non-normal data using best-fit distribution and percentile method.
Capability	**Batch Cpk**	Run Cpk on ALL columns with saved spec limits at once -- full quality dashboard.
Capability	**Goal Plot**	Capability overview of all characteristics on one plot.
Capability	**Sigma Level**	The process sigma level and defect rate.
Capability	**Yield**	First-pass yield from pass/fail results.
Attribute Capability	**Binomial**	Capability for defectives (pass/fail) data.
Attribute Capability	**Poisson**	Capability for defects-per-unit data.
Measurement	**Gage R&R**	Measurement System Analysis -- quantify how much variation comes from the measurement process itself.
Measurement	**Type 1 Gage**	Is the gauge's repeatability small enough for the tolerance?
Measurement	**Attribute MSA**	How consistently inspectors judge pass/fail.
Measurement	**Cohen's Kappa**	Measure agreement between two raters on categorical classifications.
Measurement	**Bias / Linearity**	Does the gauge read high or low, and does that change across the range?
Sampling	**Acceptance Sampling**	Design a single sampling plan or draw its OC curve.
Sampling	**Sample Size**	Calculate required N for desired statistical power and effect size.
Problem Solving	**Pareto**	Identify the vital few causes responsible for most defects (80/20 rule).
Problem Solving	**Fishbone**	A cause-and-effect (Ishikawa) diagram built from your columns, with the causes ranked when you choose a response.
Problem Solving	**Variability**	Multi-level nested box plot showing which grouping factor contributes most variation.

Group	Button	What it does
Problem Solving	**Process Screening**	Summary of every measurement on one page.
Problem Solving	**Process Map**	Draw a process flowchart from a list of steps.
Problem Solving	**Monte Carlo**	Propagate input uncertainty through a formula to predict output distribution.

The SPC tab

Group	Button	What it does
Control Charts	**Control Chart**	Control chart for individual measurements with no subgroups.
Control Charts	**Xbar-R / S**	Monitor subgroup means and variation over time. Out-of-control points flagged in red.
Control Charts	**p, np, c, u**	Control charts for counts of defectives or defects.
Control Charts	**Laney p'**	p chart corrected for overdispersion.
Control Charts	**EWMA / CUSUM**	Detect small process shifts (0.5-2σ) that standard charts miss.
Control Charts	**Moving Average**	Control chart of moving averages for small shifts.
Control Charts	**G Chart**	Control chart for days or units between rare events.
Control Charts	**Hotelling T^2**	Multivariate control chart monitoring several variables jointly.
Control Charts	**Short Run**	Several part numbers made on one process, on one chart.
Control Charts	**Pre-Control**	Run or stop, reading by reading, from green, yellow and red zones.
Shop Floor	**Data Entry**	Operators enter readings one at a time; each is judged on the spot.
Shop Floor	**Dashboard**	Every characteristic's status on one page, worst first.
Limits and Causes	**Record a Cause**	Note why a point was out of control and what was done about it.
Limits and Causes	**Cause Log**	Every recorded cause and corrective action.
Limits and Causes	**Saved Limits**	Control limits locked for each column.
Limits and Causes	**Unlock Limits**	Remove the saved control limits for a column.
Chart Rules	**SPC Alarms**	Check current dataset against last control chart's limits using all 8 Western Electric rules.
Chart Rules	**Alarm Rules**	Choose which Nelson rules control charts apply.
Chart Rules	**Clear Selection**	Clear rows selected on a chart.

The Compare Groups tab

Group	Button	What it does
Means	**2-Sample t**	Compare the means of two groups.
Means	**1-Sample t**	Test whether a mean equals a target value.
Means	**Paired t**	Paired t-test for before/after or two methods on the same subjects.
Means	**1-Sample Z**	Test a mean when the process standard deviation is known.
Means	**ANOVA**	Test whether means differ across 3+ groups. Includes Tukey HSD post-hoc.
Means	**Welch ANOVA**	Compare three or more means without assuming equal variances.
Means	**Two-Way ANOVA**	Effects of two factors and their interaction on a measurement.
Comparisons	**Tukey HSD**	Which groups differ, after ANOVA.
Comparisons	**Dunnett**	Compare every group with one control group.
Comparisons	**ANOVA Checks**	Check normality and equal spread before ANOVA.
Variances	**Equal Variances**	Compare the spread of three or more groups.
Variances	**2 Variances**	Compare the spread of two groups.
Variances	**1 Variance**	Test whether the standard deviation equals a target.
Nonparametric	**Mann-Whitney**	Compare two groups without assuming normality.
Nonparametric	**Kruskal-Wallis**	Compare three or more groups without assuming normality.
Nonparametric	**Mood's Median**	Compare group medians; robust to outliers.
Nonparametric	**1-Sample Wilcoxon**	Test a median against a target (symmetric data).
Nonparametric	**Sign Test**	Test a median against a target with no shape assumption.
Nonparametric	**Friedman**	Compare treatments applied to the same blocks.
Proportions & Counts	**2 Proportions**	Compare the defect rate of two groups.
Proportions & Counts	**1 Proportion**	Test a defect or pass rate against a target.
Proportions & Counts	**Crosstabs**	Contingency table of two categorical variables with chi-square test.
Proportions & Counts	**Goodness of Fit**	Test whether category counts match expected shares.
Proportions & Counts	**Fisher Exact**	Exact test of association for a 2x2 contingency table.
Proportions & Counts	**McNemar**	Compare paired pass/fail results.
Proportions & Counts	**Poisson Rate**	Test a defect rate per unit, or compare two rates.
Proportions & Counts	**Relative Risk**	Relative risk and odds ratio with 95% confidence intervals.
Equivalence	**Equivalence**	Prove that two processes or methods are equivalent within a pre-specified margin.
Equivalence	**Non-Inferiority**	Tests whether a new treatment is not worse than a control by more than a margin.
Equivalence	**Superiority**	Tests whether a treatment is superior to control by more than a margin.

Group	Button	What it does
Advanced	**MANOVA**	Compare groups on several measurements at once.
Advanced	**Repeated Measures**	Compare conditions measured on the same subjects.
Advanced	**Mixed Model**	Regression accounting for repeated measurements or grouped data structure.
Advanced	**Propensity Matching**	Estimates a causal treatment effect by balancing confounders.

The Predict tab

Group	Button	What it does
Regression	**Linear**	Model a continuous output as a linear function of predictors. Reports coefficients, R^2, VIF, and diagnostic plots.
Regression	**Polynomial**	Fit a curve (quadratic, cubic, ...) to two columns.
Regression	**Stepwise**	Automatically select the best predictor subset using AIC or BIC.
Regression	**Logistic**	Predict a pass/fail (or other category) outcome.
Regression	**Spline**	Flexible curve fitting with B-splines.
Regression	**Nonlinear**	Fit exponential, growth, power or your own formula.
Regression	**Predict Y**	Enter new X values and get a predicted Y from your last fitted regression model.
Model Selection	**Best Subsets**	Rank regression models by every combination of predictors.
Model Selection	**Variance Inflation**	Variance inflation factors for predictors.
Model Selection	**Residuals**	Check the last regression's residuals.
Model Selection	**Cross-Validate**	How well the model predicts data it has not seen.
Method Comparison	**Bland-Altman**	Agreement between two measurement methods.
Method Comparison	**Orthogonal Regression**	Regression when both methods have measurement error.
Regularized	**Ridge**	Regression that stays stable when predictors overlap.
Regularized	**LASSO**	Regression that picks the predictors that matter.
Regularized	**Elastic Net**	Regularized regression preventing overfitting with automatic variable selection.
Advanced	**GLM**	Regression for count data, proportions, or positive-skewed continuous outcomes.
Advanced	**GLMM**	Generalized Linear Mixed Model: fixed effects plus random effects for grouped data.
Advanced	**PLS**	Regression for many correlated predictors or when predictors outnumber observations.
Advanced	**Discriminant**	Classify observations into groups based on numeric predictors.
Forecast	**ARIMA**	Automatically fit the best ARIMA model and forecast future values.

Group	Button	What it does
Forecast	**Smoothing**	Smooth a series and forecast ahead.
Forecast	**Decompose**	Separate a time series into trend, seasonal, and residual components.

The Reliability tab

Group	Button	What it does
Life Analysis	**Weibull**	Fit Weibull distribution to failure data. Estimate MTTF, B10, B50, reliability.
Life Analysis	**Kaplan-Meier**	Non-parametric survival function from censored time-to-event data.
Life Analysis	**Life Distribution**	Fit Weibull, lognormal, exponential and log-logistic models to life data, counting suspended units, and rank them by AIC.
Life Analysis	**3P Weibull**	Weibull with threshold γ for data with a guaranteed minimum life before first failure.
Life Analysis	**Life Data**	Weibull analysis of failure and suspension times.
Life Analysis	**Hazard Function**	How the failure rate changes with age.
Metrics	**MTBF / MTTF**	Mean time between failures, and availability.
Metrics	**Block Diagram**	Computes system reliability from component reliabilities.
Metrics	**Reliability Growth**	Is reliability improving during development testing? Duane and Crow-AMSAA growth models.
Calculators	**Stress-Strength**	Reliability from strength and load distributions.
Calculators	**Availability**	Steady-state availability of a unit or system.
Test Planning	**Demonstration Test**	Sample size to demonstrate a reliability with confidence.
Test Planning	**MTBF Demonstration**	Test time to demonstrate an MTBF, or the MTBF a test showed.
Test Planning	**Sample Size**	Calculate required N for desired statistical power and effect size.
Advanced	**Arrhenius ALT**	Model how temperature accelerates failure to predict life at normal use conditions.
Advanced	**Degradation**	Predict failure from a measurement that wears over time.
Advanced	**Fatigue S-N**	Woehler S-N curve fitting: stress amplitude vs cycles to failure.
Advanced	**Competing Risks**	CIF for each failure cause when multiple modes compete.
Advanced	**Recurrence**	Analyze how often repairable systems fail using Mean Cumulative Function.

The DOE tab

Group	Button	What it does
Design	**Screening**	Find the few important factors among many.
Design	**Full Factorial**	All combinations of all factors at all levels -- estimates every effect and interaction.
Design	**Fractional**	A fraction of the full factorial with known confounding.
Design	**Response Surface**	CCD or Box-Behnken design for fitting quadratic models and finding optimal settings.
Design	**Definitive Screening**	Screen many factors AND estimate quadratic effects in 2k+1 runs.
Design	**Taguchi**	An orthogonal array for robust design.
Design	**Mixture**	Experiments on the proportions of a blend.
Custom Designs	**D-Optimal**	D-optimal design with optional linear constraints on factor settings.
Custom Designs	**Custom Design**	A design with exactly the number of runs you can afford.
Custom Designs	**Augment**	Add runs to an existing experiment.
Custom Designs	**Split-Plot**	DOE for hard-to-change and easy-to-change factors using restricted randomization.
Custom Designs	**Bayesian Optimization**	Suggests the next experiments via Gaussian Process surrogate + Expected Improvement.
Custom Designs	**Latin Hypercube**	Space-filling design for computer experiments and simulation models.
Planning	**Sample Size**	Calculate required N for desired statistical power and effect size.
Planning	**Equal Variances**	Compare the spread of three or more groups.
Analyze	**Effects Plot**	Which factors move the response, and by how much.
Analyze	**ANOVA**	ANOVA table and coefficients for an experiment.
Analyze	**Profiler**	Explore predicted responses by moving the factors.
Analyze	**Contour Profiler**	Predicted response over two factors as contours.
Analyze	**Surface Plot**	The response surface in 3D.
Analyze	**Desirability**	Optimize multiple responses simultaneously using Derringer-Suich desirability functions.

The Visualize tab

Group	Button	What it does
Distribution	**Histogram**	The distribution of one column as bars.
Distribution	**Box Plot**	Median, quartiles and outliers, optionally by group.
Distribution	**Violin**	Box plot with the full shape of each group.
Distribution	**Dot Plot**	Stacked-dot distribution plot.

Group	Button	What it does
Comparison	**Mean Plot**	Group means with 95% confidence-interval error bars.
Comparison	**Error Bars**	Group means with error bars (SE, SD or a 95% confidence interval).
Comparison	**Variability**	Multi-level nested box plot showing which grouping factor contributes most variation.
Comparison	**Overlay**	Several columns plotted against one X axis.
Relationships	**Scatter**	One measurement against another, with an optional trend line.
Relationships	**Bubble**	Scatter plot with a third column as bubble size.
Relationships	**3D Scatter**	Three measurements in a rotatable 3D plot.
Relationships	**Scatter Matrix**	Every pair of columns as scatter plots in one grid.
Relationships	**Heatmap**	A grid colored by value: correlations or a two-way table.
Relationships	**Contour**	Lines of equal response across two columns.
Relationships	**Parallel**	Each row as a line across several axes.
Over Time	**Run Chart**	Values over order with a median reference line.
Over Time	**Time Series**	Columns plotted in time order.
Over Time	**Area**	Filled line chart for one or more series.
Composition	**Pareto**	Identify the vital few causes responsible for most defects (80/20 rule).
Composition	**Mosaic**	Tile plot where area is proportional to cell frequency.
Composition	**Treemap**	Nested rectangles sized by a value.
Composition	**Ternary**	Three-part compositions on a triangle.
Builder	**Graph Builder**	Pick any chart type and its columns in one place.
Builder	**Geo Map**	Choropleth map by region code, or scatter map by lat/lon.
Builder	**Word Cloud**	Visualizes word frequency in a text column.

The Machine Learning tab

Group	Button	What it does
Classify / Predict	**Random Forest**	Ensemble of decision trees -- robust, handles non-linearity, reports feature importances.
Classify / Predict	**Decision Tree**	Predict a target from other columns: decision tree.
Classify / Predict	**Gradient Boost**	Predict a target from other columns: gradient boosting.
Classify / Predict	**Neural Net**	Predict a target from other columns: neural network.
Classify / Predict	**SVM**	Predict a target from other columns: support vector machine.
Classify / Predict	**Naive Bayes**	Predict a target from other columns: naive bayes.
Classify / Predict	**Discriminant**	Classify observations into groups based on numeric predictors.

Group	Button	What it does
Validate	**Holdout**	Explicit train/test split with a complete holdout evaluation report.
Validate	**Precision-Recall**	Precision-Recall and ROC curves for 4 classifiers compared on the same test set.
Validate	**Cross-Validate**	How well the model predicts data it has not seen.
Cluster & Reduce	**K-Means**	Group rows into clusters of similar values.
Cluster & Reduce	**Hierarchical**	Group similar observations using tree-based clustering with dendrogram.
Cluster & Reduce	**PCA**	Reduce many correlated variables to uncorrelated principal components.
Cluster & Reduce	**Factor Analysis**	Find latent factors underlying correlations between variables.
Cluster & Reduce	**t-SNE**	2D nonlinear embedding that reveals clusters in high-dimensional data.
Cluster & Reduce	**UMAP**	Manifold embedding that preserves both local and global structure.
Cluster & Reduce	**MDS**	2D map where distances approximate original high-dimensional distances.
Text	**Text Explorer**	Analyze free text: term frequency, TF-IDF, bigrams, and LDA topic modeling.
Text	**Sentiment**	Scores text as positive, negative, or neutral.
Text	**Text Regression**	Predicts a numeric value from text.
Text	**Topics (LSA)**	Extracts hidden topics from a text column.
Advanced	**SEM**	Fit CFA and full SEM models with latent variables, path coefficients, and fit indices.
Advanced	**Latent Class**	Identifies hidden (latent) subgroups in your data.
Advanced	**Correspondence**	Maps rows and columns of a contingency table into a shared 2D space.
Advanced	**Multiple Correspondence**	MCA maps 3 or more categorical variables into a shared 2D space.
Advanced	**Association Rules**	Find co-occurrence patterns in transactional data (market basket analysis).
Advanced	**Bayesian**	Update beliefs about a proportion or mean using data and a prior.
Advanced	**Item Analysis**	Assess internal consistency of a multi-item scale.

The Data tab

Group	Button	What it does
Import	**Open File**	Load a dataset from disk (CSV, Excel, JSON, Parquet, SQLite, TSV).
Import	**Paste**	Paste a copied Excel or spreadsheet table directly into Arthur.
Import	**Live Data**	Keep a dataset in step with an Excel file or a Google Sheet.
Import	**SQL Import**	Load a table from a database.
Import	**New Sheet**	A blank sheet to type or paste into, like a new Excel workbook.

Group	Button	What it does
Import	**Concatenate**	Stack two datasets with the same columns.
Datasets	**Switch**	Choose which open dataset is active.
Datasets	**Duplicate**	Make a working copy of the dataset.
Datasets	**Compare**	What differs between two datasets.
Datasets	**Revert**	Return to the data as first opened.
Datasets	**Close**	Close the active dataset.
Clean	**Missing Values**	Visual dashboard showing which columns have missing data and how severe it is.
Clean	**Remove Duplicates**	Delete repeated rows.
Clean	**Fill Missing**	Impute missing values using statistical methods.
Clean	**Validate**	Find rows that break a rule.
Clean	**Column Types**	See and change how each column is read.
Transform	**Formula**	Create a new column from a Python/numpy expression using existing column names.
Transform	**Derived Column**	Add a column calculated from others.
Transform	**Standardize**	Rescale numeric columns to a common scale before machine learning.
Transform	**Box-Cox**	Find the power transformation that makes a skewed, positive column closest to normal, and add it as a new column.
Transform	**Recode**	Replace a value throughout a column.
Transform	**Lag / Lead**	Add a column shifted up or down.
Transform	**Cumulative**	Running totals, means or counts.
Transform	**Parse Dates**	Turn text into dates and times.
Transform	**Split Column**	Split text at a delimiter into several columns.
Transform	**Extract Text**	Pull part of a text value into a new column.
Organize	**Filter**	Subset rows matching a condition and save as a new dataset.
Organize	**Filter Panel**	Keep rows that meet a condition.
Organize	**Row States**	Highlight or exclude specific rows matching a Boolean expression.
Organize	**Sort**	Sort the rows by a column.
Organize	**Aggregate**	Summarize by groups -- like a powerful pivot table.
Organize	**Join**	Combine two datasets on a shared key.
Organize	**Stack**	Turn several columns into one long column.
Organize	**Unstack**	Spread one column into several, by group.
Organize	**Transpose**	Swap rows and columns.

The Reports tab

Group	Button	What it does
Generate	**Templates**	Build one of the common SPC reports, step by step.
Generate	**Executive**	Generate a short summary for management.
Generate	**DMAIC**	A Six Sigma project report in the five DMAIC phases.
Generate	**Manufacturing**	Generate capability, control and MSA for production.
Generate	**Reliability**	Generate life data and reliability metrics.
Export	**PDF**	Export the results as a PDF.
Export	**Excel**	Export the results as a Excel workbook.
Export	**PowerPoint**	Export the results as a PowerPoint deck.
Export	**HTML**	Export the results as a web page.
Export	**Interactive HTML**	Export self-contained HTML with live, interactive Plotly charts.
Export	**Chart Image**	Save the charts from the results as images.

The View tab

Group	Button	What it does
Panels	**Columns**	Show or hide the column list on the left.
Panels	**Results**	Show or hide the results on the right.
Panels	**Data Table**	Show the data table in the middle.
Panels	**Column Details**	Details of the selected column.
Windows	**Full Screen**	Toggle full-screen mode.
Windows	**Tile**	Arrange open data windows side by side.
Windows	**Reset Layout**	Put every panel back where it started (Ctrl+L).
Windows	**Close All Results**	Close every result tab.
Display	**Text Size**	Make the text larger or smaller.
Display	**Lock Panels**	Stop panels being dragged or floated by accident.
About	**Guides**	These guides.
About	**About**	Version and license information.