ABL Digital Technologies

THE SHIFT PLANNING TOOL

A Practitioner’s Guide

What the tool does · Every input explained · How the maths runs
The interval distribution editor · The outputs and the Excel file
The three charts · A worked example end-to-end · Limits and tips

For the underlying queueing maths (Erlang B & C), see the companion "Erlang Formulas — Plain-Language Explainer" document.

THE SHIFT PLANNING TOOL

A Practitioner’s Guide

What the tool does · Every input explained · How the maths runs
The interval distribution editor · The outputs and the Excel file
The three charts · A worked example end-to-end · Limits and tips

For the underlying queueing maths (Erlang B & C), see the companion "Erlang Formulas — Plain-Language Explainer" document.

1 · What This Tool Does

The Shift Planning tool answers one question: given a forecast of contact volume, a service-level target, and a fixed pool of available staff, how should you deploy those staff across the day and the week — and what service level will you actually achieve?

It does five things, in order:

- Takes your forecast (monthly volume + day-of-week split + an interval distribution) and explodes it into 96 fifteen-minute volume buckets for each of the seven days.
- Runs the Erlang C formula on each interval to compute how many agents are required to meet the target service level, with an occupancy-cap floor and a minimum-staff floor applied.
- Sizes the workforce honestly — the "Required FTE" calculation uses the peak-hour headcount per day (the average of the top 6 required intervals), summed across the week, then divided by working days per FTE, then inflated for shrinkage.
- Assigns shifts greedily — one shift at a time — to cover the requirement, while respecting (a) the operating window (no shift starts before the operation opens or extends past close) and (b) the Available FTE budget (the week can use no more shifts than your roster supports).
- Predicts the service level you will achieve given your actual staffing, with a capacity-ratio penalty that crashes SL when staffing is short — and then exports a full styled Excel workbook with charts.

Everything runs in your browser. No data is sent anywhere. The page is a single HTML file with a JavaScript engine; the Excel workbook is built using the open-source ExcelJS library and Chart.js for the chart images.

This document explains the workflow, every input, every output, and where to look for what. For the deep mathematics of Erlang B and Erlang C, see the companion "Erlang Formulas — Plain-Language Explainer" document, which walks through the queueing maths from first principles.

2 · Before You Begin — Vocabulary You Will Need

Erlang (the unit of load)

One Erlang is the amount of traffic generated by one phone line being busy for one full hour. If you have 100 calls per hour and each one lasts 3 minutes, your load is (100 × 180) / 3600 = 5 Erlangs — the same as having five lines fully occupied for the entire hour.

Erlang C

A formula for the probability that a new caller has to wait, given the number of agents and the load. From that probability we derive a predicted service level, average speed of answer, queue length, and so on. The Erlang C model assumes calls arrive randomly, agents are identical, and callers wait patiently. It is over a hundred years old and remains the backbone of contact-centre staffing.

Service Level (SLA)

A target written as "X% of calls answered within Y seconds." A typical target is 80/20: 80% of calls answered within 20 seconds. The tool asks for both X (as a percentage) and Y (as seconds).

Shrinkage

The percentage of paid agent time that is not spent on calls — breaks, training, meetings, sick leave, attrition allowance. If shrinkage is 18%, then for every 100 agents on the roster only 82 are effectively available on the phone at any moment.

Occupancy

The fraction of time an agent is actively on a call. With load A Erlangs and N agents on the phone, occupancy is A / N. High occupancy (say above 85%) is operationally unsustainable — agents burn out, quality drops, attrition climbs. The tool lets you set an upper bound on occupancy as a sanity check on staffing.

FTE (Full-Time Equivalent)

One person on the roster working one full shift per working day. If each FTE works five days a week with two days off, then one FTE contributes five shifts to the weekly capacity. The tool asks for your available FTE — the total headcount on roster — and uses this as a hard cap on the assignment.

Interval

A 15-minute slice of the day. There are 96 intervals in a day (24 hours × 4 per hour). The tool computes Required and Provided at every interval, and the Excel sheets show the per-interval picture for each day of the week.

Shift

A continuous block of consecutive intervals during which an agent is on duty. The tool uses a fixed shift length (you choose how many hours) and assigns each shift a start time. A 9-hour shift covers 36 intervals.

Operating window

The contiguous span of intervals during which the operation is open (i.e. has non-zero forecast volume). The tool infers this from your interval distribution: the first and last intervals with positive volume. Shifts can only start such that they fit entirely inside the window — no shift extends past close.

3 · The Inputs — Every Field Explained

3.1 Service Level Parameters Card

Eight fields that define what "meeting service" means for your operation, plus the basic shift mechanics.

SLA Target (%)

The percentage of contacts that should be answered within the Service Time. Typical values: 80, 85, 90. The Erlang C formula uses this as the target the per-interval Required number must achieve.

Service Time (seconds)

The wait threshold inside which the SLA is measured. For "80/20" this is 20 seconds. For chat or back-office this can be much longer — 300, 600, or even 1200 seconds. Longer Service Time means more forgiving staffing, because callers are assumed to wait patiently.

AHT (Average Handle Time, seconds)

The average total time an agent spends on one contact — talk time plus after-call work. Affects Erlang load directly: load A = (calls/hour × AHT) / 3600. Higher AHT means more agents needed.

Shrinkage (%)

Total non-productive percentage of paid agent time. Used in two places: (a) when assigning a shift to a 15-minute interval, only (1 − shrinkage) of that body is treated as actually on the phone; (b) when computing Required FTE from on-call headcount need, the roster is inflated by 1 / (1 − shrinkage). Typical range: 25–35%.

Occupancy Cap (%)

The maximum acceptable occupancy at the interval level. Required staffing is bumped up to this floor if Erlang C alone would result in occupancy above the cap. A 65% cap is conservative; 85% is aggressive. Use 80–85% for cost-driven operations, 60–70% for quality-sensitive ones.

Min Staff per Interval

A floor that applies during operating hours regardless of volume. Useful for ensuring at least one agent is on duty even at the slowest period. Set to 1 for typical operations; set to 0 if you genuinely want zero staffing in zero-volume intervals.

Shift Length (hours)

How many hours each scheduled shift covers. Standard values are 8 or 9. The tool converts to intervals automatically: 9 hours = 36 intervals. The last legal shift start each day is (window close – shift length) so the shift ends at or before close.

Weekly Offs (days)

How many days off each FTE takes per week. Standard is 2 (a five-day working week). Affects (a) the weekly shift budget — Available FTE × (7 – Weekly Offs) — and (b) the Required FTE calculation.

Average Contact Center Working Days per Month

How many days per month the operation actually runs. The tool divides Monthly Volume by this number to get an average daily volume baseline, which is then redistributed across the seven days using the DoW split. 22 for Mon–Fri operations, 26 for Mon–Sat, 30 for 24×7.

Available FTE

Total headcount on roster who can be scheduled. This is a HARD CAP. The weekly shift budget is Available FTE × (7 – Weekly Offs). The greedy assignment will not exceed this — it stops when either all demand is met or the budget is exhausted. If the budget runs out before demand is met, the Coverage Status shows "■ Deficit" and you have a hiring decision to make.

3.2 Forecast Volume Card

Monthly Volume (total contacts)

Total contacts expected for the month. The tool splits this across the working days using the day-of-week percentages.

Day-of-Week Split (Mon–Sun, %)

Seven percentage fields. They need not sum exactly to 100 — the tool normalises. Use historical averages: for retail banking the typical pattern is Monday 24%, Tuesday 22%, Wednesday 17%, Thursday 14%, Friday 18%, Saturday 3%, Sunday 2%. Set Saturday and Sunday to zero for Mon–Fri-only operations.

3.3 Interval Distribution Card

Profile Preset (dropdown)

Five options that determine how each day’s volume is spread across the 96 fifteen-minute intervals:

- Office hours 08:00 to 20:00 bell curve — a smooth $\sin^2$ bell, zero outside the window. Suitable for office-hours operations.
- Extended 07:00 to 22:00 mild peak — same shape, wider window. For extended-hours retail or banking.

- 24×7 twin peaks — bimodal curve with peaks at 10:00 and 19:00 plus a non-zero overnight baseline. For 24-hour operations.
- Custom (textarea) — paste 96 numbers, applied uniformly to all seven days. Use this when you have a single per-interval profile from your historical data.
- Per-Day Custom — uses the saved 96 × 7 matrix from the Distribution Editor (see Section 5). Use this when each day has a distinctly different pattern (Saturday opens later, Sunday is half-day, etc.).

Custom Distribution textarea

Used only when the dropdown is set to "Custom". Paste 96 values separated by commas, spaces, or newlines. The tool normalises so the values sum to 100% of the daily volume.

Edit per-day distribution table button

Opens a modal overlay with a 96-row × 7-column editable table. Each cell is the percentage of that day's volume falling into that 15-minute interval. Pre-filled from whichever preset is currently selected — so if "Office hours 08:00 to 20:00" is selected, the editor opens with all seven days showing the office bell curve, and you can then customise any day's row. Save & Apply normalises each column to exactly 100% and switches the dropdown to "Per-Day Custom".

4 · The Maths At A Glance

4.1 From Monthly Volume to Per-Interval Volume

```
avgDailyVolume = MonthlyVolume / MonthlyDays
dailyVolume[d] = avgDailyVolume × 7 × dowPct[d]
intervalVolume[d,i] = dailyVolume[d] × intervalPct[d, i]
```

intervalPct is always a 7 × 96 matrix internally — for presets and single-column custom it is the same row broadcast to all seven days; for "Per-Day Custom" it is the saved matrix from the editor.

4.2 Required Staff per Interval (Erlang C)

For each interval, the tool computes the calls-per-hour rate from the 15-minute volume, then runs Erlang C to find the smallest N such that the predicted SLA reaches the target:

```
callsPerHour = intervalVolume × 4
A (load in Erlangs) = callsPerHour × AHT / 3600
Required(SL) = smallest N where slaAt(N, T, calls, AHT) ≥ SLA target
Required(OccCap) = ■ A / occCap ■
Required[i] = max( Required(SL), Required(OccCap), MinStaff )
```

4.3 Greedy Shift Assignment

Once every interval has a Required count, the tool builds a shift roster. The algorithm is greedy: each pass picks the (day, shift start) combination that fills the largest remaining deficit, adds one shift there, and repeats. Each scheduled shift contributes (1 − shrinkage) to the Provided count at each covered interval — because a real person on roster is only on the phone for that fraction of their shift, the rest being eaten by breaks etc.

```
per-shift contribution = (1 − shrinkage) per interval covered
weekly shift budget = Available FTE × (7 − Weekly Offs)
per-day cap = Available FTE
```

The loop stops when either every interval is covered (all deficits cleared) or the weekly budget is exhausted. Shifts can only start inside the operating window such that the entire shift fits before close.

4.4 Predicted Service Level

Per interval, the predicted SLA uses the Erlang C formula at the effective on-call count (Provided), with a capacity-ratio penalty when Provided is below Required:

```
if Provided ≤ A → SLA = 0 (system melted)
if Provided ≥ Required → SLA = Erlang C SLA at Provided
else → SLA = Erlang C SLA × (Provided − A) / (Required − A)
```

The capacity-ratio penalty is what makes the predicted SL crash when staffing is materially short of requirement. Without it the formula gave high values even at noticeable deficits because Required carries occupancy headroom that does not directly affect Erlang C.

4.5 Required FTE — Top-6 Peak Method

For sizing the roster (as distinct from the per-interval Required), the tool uses what the operations world calls a peak-hour method:

```
For each day d:
perDayPeak[d] = average of the 6 highest Required intervals that day
Sum across week = Σ perDayPeak[d]
÷ Working Days per FTE = on-call headcount needed
÷ (1 − Shrinkage) = Required FTE on roster
```

Why top 6? Six intervals at fifteen minutes each is about ninety minutes — the typical busy-hour-and-a-half. Sizing to the average of the top 6 captures the operation’s real peak demand. The simpler "sum of all required intervals divided by total slots" approach dilutes the peak by averaging in low-volume intervals and produces a number that looks fine in aggregate while every peak hour is short.

4.6 The FTE Deployment Math

Separate from "Required FTE", the tool also reports what was actually deployed inside the FTE budget:

```
Total Shifts Deployed = sum of shifts assigned across the week
FTE equivalent used = Total Shifts / Working Days per FTE
Capacity utilisation = Total Shifts / Weekly Shift Budget
On-call delivery = Total Shifts × 36 × (1 − shrinkage)
```

Coverage Status is "✓ Full" if the assignment covered every interval, "■ Deficit" if the budget ran out first. The number of uncovered interval-slots is reported alongside.

5 · The Per-Day Distribution Editor

Real operations rarely have the same intraday pattern on every day of the week. Saturday opens later. Sunday may run only a half-day. Monday peaks earlier than Friday. The Distribution Editor lets you capture these differences without abandoning the bell-curve preset as a starting point.

5.1 Opening the Editor

Below the profile dropdown is a button: "Edit per-day distribution table (96 × 7) →". Click it and a modal opens with a scrollable table — one row per 15-minute interval, one column per day, plus an interval label and a sticky header showing day names.

5.2 What the Cells Mean

Each cell is the percentage of that day's total volume falling into that 15-minute interval. The columns will each sum to roughly 100% (visible at the bottom of the editor in a "Total %" row). The Total row turns gold if a column deviates from 100% by more than half a percentage point — but normalisation on Save will correct this regardless, so you do not need to make the columns sum exactly.

5.3 Pre-filling From a Preset

When you first open the editor, every column is pre-filled from the currently selected preset. So if "Office hours 08:00 to 20:00" is selected, all seven columns show the same bell curve. Edit only what differs — flatten Saturday, zero out Sunday after midday, shift the Wednesday peak earlier, whatever you need.

5.4 Save & Apply

Save normalises each column to sum to exactly 100%, stores the 7 × 96 matrix, and switches the profile dropdown to a new option called "Per-Day Custom". From then on, the Compute Plan button uses this matrix. The Required staffing, the shift assignment, the charts, the heatmap, and the per-day Excel sheets all reflect the edits.

5.5 Reset and Cancel

Reset to Preset (with confirmation) re-seeds every cell from the currently selected preset, losing your edits. Cancel closes the modal without saving — any unsaved edits are discarded. Switching the dropdown back to a preset (e.g. "Extended") without re-opening the editor reverts the tool to using the preset; the saved matrix is still in memory and is restored if you switch back to "Per-Day Custom" later.

6 · The Outputs — On the Page

6.1 Required FTE Card

The honest answer to "how many people do you need on the roster?" Shows the top-6 peak chain in full:

- Sum of Daily Peaks (avg top 6 intervals per day)
- ÷ Working Days per FTE = On-call Headcount
- ÷ (1 − Shrinkage) = Required FTE
- Available FTE (echoed back)
- Gap vs Available FTE (red if short, green if surplus, gold if matched)
- Sizing method label

Below the chain is a math text block showing every day's peak (Mon, Tue, ..., Sun) so you can see where the load is concentrated.

6.2 Predicted Performance Card

Two numbers:

- Predicted SLA — volume-weighted across the week, with the capacity-ratio penalty applied to every under-staffed interval.
- Predicted Occupancy — total call-seconds divided by total effective on-call seconds.

6.3 FTE Deployment Math Card

What was actually deployed within the FTE budget:

- Available FTE (input)
- Working Days per FTE per week
- Total Shift Capacity = Available FTE × Working Days
- Total Shifts Deployed = actually-assigned shifts
- FTE Equivalent Used = Total Shifts / Working Days
- Capacity Utilisation = Deployed / Capacity
- Coverage Status (✓ Full or ■ Deficit)
- Uncovered Interval-Slots (the deficit count)

A math text block underneath shows the calculation step by step.

6.4 Per-Day Breakdown Table

Seven rows (one per day of week) plus a Total / Peak row. Each row shows daily volume, the peak interval and peak required, the number of shifts assigned that day, the day FTE equivalent, and the day's volume-weighted predicted SLA.

6.5 Required Staffing Heatmap

A 7 × 24 grid (days × hours) showing the maximum required staff in each hour, colour-coded by intensity. Useful for spotting the shape of the week at a glance — when the peaks are, which days are busiest, where the operating window opens and closes.

7 · The Excel Workbook — Eleven Sheets

Clicking Download Excel produces ABL_Shift_Plan.xlsx with all the detail behind the on-page summary. The workbook is professionally styled — gold and dark-brown headers, alternating row backgrounds, colour-coded variance, sticky column widths. Eleven sheets in order:

7.1 Summary

A single sheet with everything that matters at a glance:

- Title and timestamp
- INPUTS section — every input value echoed back
- REQUIRED FTE — PEAK-BASED (TOP-6 METHOD) — the chain with per-day peaks, sum, on-call headcount, Required FTE, Available FTE, Gap (red/green text), plus Peak Day Shifts for reference and Predicted SLA / Occupancy
- FTE DEPLOYMENT MATH — the chain with shifts deployed, FTE used, capacity, utilisation, and coverage status
- DAY-BY-DAY BREAKDOWN table

7.2 Charts

Three embedded chart images (rendered by Chart.js, then embedded as PNGs by ExcelJS — they appear identical to native Excel charts but are static). Each chart has a section header above it.

7.3 Daily_Volume_Profile

A 96 × 7 grid showing the per-interval volume for each day of the week, with the interval label in the first column.

7.4 Mon / Tue / Wed / Thu / Fri / Sat / Sun

Seven sheets, one per day. Each has a title banner with the day name, a subtitle showing daily volume and peak, then the per-interval table: Interval | Volume | Required | Provided | Variance | SL% | Shift Starts. The Variance column is colour-coded green when positive, red and bold when negative. A "Day Total" highlighted row at the bottom sums the relevant figures.

7.5 Shift_Pattern

Every assigned shift block in a clean table: Day | Shift Start | Agents on Shift | Shift End. Useful for translating the schedule into roster lines.

8 · The Three Charts

8.1 Volume Across the Month

A bar chart with one bar per working day in a typical month. Bars are colour-coded by day of week — Mondays in gold, Tuesdays in green, and so on. Useful for visualising how the same week pattern repeats four-or-five times across the month and where the heaviest days fall.

8.2 Volume Across the Week

A simple seven-bar chart of typical daily volume for each day of the week. Shows the relative weight of each day at a glance — whether Monday dominates, whether weekends matter, whether the distribution is balanced.

8.3 Volume Across Each Day

A multi-line chart with 96 intervals on the x-axis (every 15 minutes from 00:00 to 23:45) and one coloured line per day of the week. Shows the intraday shape clearly. When you use the per-day distribution editor to make Saturday different from Tuesday, the difference shows up here as two distinctly shaped lines.

9 · A Worked Example, End-to-End

Suppose you are sizing a retail-banking voice operation with the following parameters:

- SLA Target: 95% · Service Time: 600 sec · AHT: 290 sec
- Shrinkage: 18% · Occupancy Cap: 65% · Min Staff: 1
- Shift Length: 9 hours · Weekly Offs: 2 · Working Days/Month: 22
- Available FTE: 200
- Monthly Volume: 240,000 contacts
- DoW split: Mon 24%, Tue 22%, Wed 17%, Thu 14%, Fri 18%, Sat 3%, Sun 2%
- Distribution: Office hours 08:00–20:00 bell curve

9.1 What the Tool Computes Internally

```
avgDailyVolume = 240,000 / 22 = 10,909 contacts
Mon dailyVolume = 10,909 × 7 × 0.24 = 18,327 contacts
Sat dailyVolume = 10,909 × 7 × 0.03 ≈ 2,291 contacts
```

Each day's volume is then spread across 96 intervals using the sin² bell shape, with zero outside the 08:00–20:00 window. The peak interval is around 13:30–14:00 each day.

9.2 Per-Interval Required Calculation

At Monday's peak, intervalVolume might be ≈ 700 contacts in 15 minutes → callsPerHour ≈ 2800 → A = 2800 × 290 / 3600 ≈ 225.5 Erlangs. Erlang C for 95% SLA in 600 sec gives Required(SL) of roughly 235, occupancy cap of 65% requires ■225.5 / 0.65■ = 347, so Required[Mon, peak] = 347. (Occupancy cap binds in this case because of the long SLA window.)

9.3 Required FTE — Top-6 Peak Method

```
perDayPeak[Mon] ≈ avg of top 6 Required values ≈ 340
perDayPeak[Sat] ≈ avg of top 6 Required values ≈ 45
Sum across week ≈ 340 + 310 + 240 + 200 + 250 + 45 + 30 ≈ 1,415
÷ 5 working days per FTE = 283 on-call headcount needed
÷ (1 − 0.18) = ■283 / 0.82■ = 346 Required FTE
```

Gap = 346 − 200 = 146 short. The tool tells you you are 146 people underweight for this volume at this SLA — a clear, actionable hiring signal.

9.4 Shift Assignment

The greedy assigns shifts inside the 08:00–20:00 window (last legal start = 11:00). Each shift covers 36 intervals and adds 0.82 effective on-call agents per interval. The budget is 200 × 5 = 1,000 shifts for the week, with no day allowed more than 200 shifts. The algorithm hits the day cap on Mon, Tue, and Fri (heaviest days) — those days end up with 200 shifts each — and uses fewer on lighter days. Total deployed ≈ 950 shifts; about 50 budget slots unused on the lightest days but no spare capacity on the peak ones.

9.5 Predicted SLA

At Monday's peak intervals, Provided = 200 × 0.82 = 164 effective on-call, Required = 347, A = 225.5. So Provided is below A and the system is melted — SLA crashes to 0% at those intervals. Off-peak and lighter-day intervals do better. The volume-weighted overall predicted SLA might come out around 65–70%, well below the 95% target — consistent with being 146 FTE short.

9.6 What You Do With This

Three options on the table:

- Hire — recruit 146 more people, or as many as the business case will support.
- Loosen SLA — accept a 90% target instead of 95%, or 80% in 20 seconds instead of 95% in 600 seconds. Re-run and see how the gap narrows.
- Reduce AHT or shrinkage — if there are operational levers (better tooling to bring AHT down, less shrinkage from better roster management), the required headcount will drop proportionally.

10 · What This Tool Will Not Tell You

Erlang is a hundred-year-old model and the assumptions are tidy. Real contact centres are messy. Things the tool does not handle:

- Multi-skill agents. Erlang assumes everyone is identical and every queue is independent. If your agents are split into skill groups with priority routing, you need a discrete-event simulator.
- Real abandonment. The Erlang C model assumes infinite patience. The capacity-ratio penalty is a crude correction for the SL crash but does not predict abandons properly.
- AHT drift under load. When agents are slammed, real AHT inflates (rushed wrap-up, longer transfers). The tool assumes AHT is constant.
- Non-stationary intervals. The model assumes the arrival rate is constant within each 15-minute interval. If your volume changes inside an interval, use shorter intervals.
- Outbound, blended, back-office. Erlang is for inbound queued channels only.
- Rest-day constraints. The tool computes a weekly shift budget assuming each FTE can be placed on any day of the week subject to the weekly offs. It does not model named-individual rosters or rest-day patterns.

Use the tool for first-pass capacity planning, hiring decisions, and sensitivity analysis. For the final roster, hand off to a proper WFM system or a discrete-event simulation.

11 · Tips and Common Questions

Why is my Predicted SLA so low when Variance is small?

Because Required carries occupancy headroom on top of the Erlang C requirement. A small variance (Provided just below Required) can still mean Provided is significantly above the Erlang load A. The capacity-ratio penalty scales SL between full-target at Provided = Required and zero at Provided = A — so SL falls noticeably even at modest deficit, and crashes at large deficit. This matches the operational reality of running at high occupancy.

Why does the FTE Deployment Math say I have spare capacity when every day shows deficit?

It should not, with the new top-6 peak method. If you still see this, the operating window for the binding-deficit day is so narrow that the greedy cannot legally place more shifts there — for example, a 12-hour window with 9-hour shifts gives only 4 valid start positions, and each day caps at Available FTE. Check that your operating window is at least Shift Length + 1 hour wide.

Why does my Saturday show zero shifts assigned?

Either (a) its volume is genuinely zero, (b) its operating window is narrower than the shift length, or (c) the greedy ran out of budget before getting to Saturday. The budget exhaustion case is distinctive: you will see Coverage Status = "■ Deficit" with a large uncovered count.

Why does the editor pre-fill from a different preset than I expected?

The editor uses whatever profile is currently selected in the dropdown. If "Per-Day Custom" is selected but no matrix has yet been saved, it falls back to the office bell curve. Pick the preset you want before opening the editor.

Can I get a tighter SLA with the same FTE?

Probably not in a meaningful way. Erlang C is roughly linear in the tail — going from 95% to 99% takes substantial extra staff. The practical levers are usually elsewhere: lower AHT, lower shrinkage, longer permitted wait time. Run the tool with the alternative parameters and compare.

Where can I learn the maths in depth?

The companion document "Erlang Formulas — A Plain-Language Explainer" walks through Erlang B and Erlang C from first principles, with every formula derived in plain language and a full worked example. Available as a free download from the Tools page.