LESSON BOOK

Formulating Probabilistic Engineering Questions

From physical systems and uncertain evidence to decisions, experiments and risk

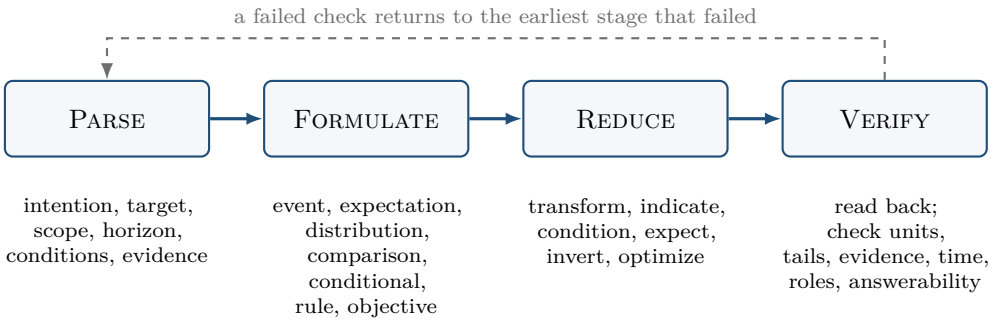
Draft · 345 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Formulating Probabilistic Engineering Questions

From Physical Systems and Uncertain Evidence to Decisions, Experiments, and Risk



Ricardo Marquez

Working manuscript, November 2026

How to read this book

This book teaches the work that comes before a numerical method: hearing an engineering question, identifying the probabilistic object that would answer it, writing that object mathematically, reducing it to a recognized general solution pattern, and checking that the final expression still answers the question that was asked. It does not teach how to implement an inference algorithm. Once a lesson has named the problem class and the inputs it needs, the problem is handed to the methods literature and to software, and Appendix E is that hand-off.

One practice runs through every lesson, and Figure 1 draws it: four stages, each named for what it produces.

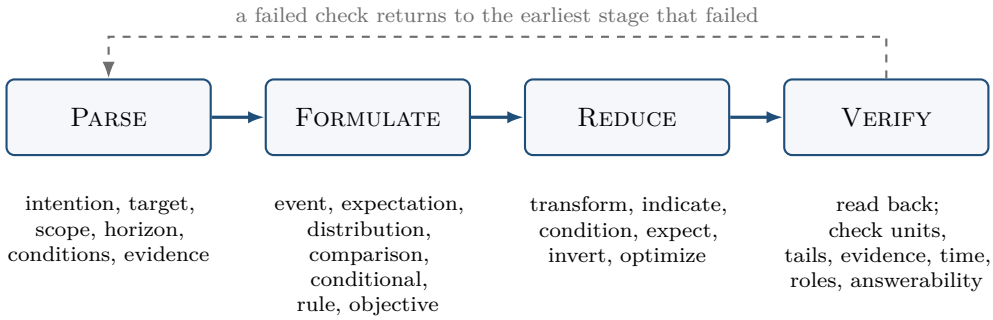


Figure 1. The four-stage practice. Each stage produces the objects listed beneath it; a check that fails in the last stage sends the formulation back to the earliest stage that failed.

Parse turns a concern into a request: an intention, a target, a scope, a horizon, the conditions that apply and the evidence in hand. Formulate writes the mathematical object that would supply the requested answer: an event, an expectation, a distribution, a comparison, a conditional, a decision rule or a design objective. Reduce exposes the operations and the model outputs that evaluating it requires, so that the problem class can be named and the model's interface made visible. Verify reads the expression back in engineering words and tests it against the original request.

The shape of a lesson. Every lesson has the same parts, so that a reader can find each one at speed. It opens with its objectives and prerequisites, then with the engineering question beside the nearby question it is most easily confused with, set in a blue box. The canonical formulation follows, then a derivation or reduction that exposes the reusable pattern, a pattern map or comparison table, and a sustained worked case drawn from the book's recurring systems. When a second case appears, it changes the mathematical structure rather than merely changing the engineering nouns. A compact procedure, in a green box, collects the steps. The exercises come in several modes, among them a reverse formulation, which turns an expression back into a sentence, and a critique and repair of a wrong formulation. A summary closes the lesson with the capability the reader should now have, in a gold box.

The systems. Six named systems recur so that the reader can accumulate engineering context instead of starting over with every formula. Riverbend Energy Center is the flagship case; North Fork Bridge, the Meridian Battery Program, Harbor Flood System, Ridgeway Wind Farm and Atlas Weld Line provide recurring contrasts. Each lesson stays mainly with one of them and uses at most one purposeful transfer. Their physical boundaries, owners, evidence conventions and stable notation are frozen in Appendix D. A later lesson may add evidence or a new decision, but it does not silently change the case.

The appendices. Appendix A holds a solution for every numbered exercise; each names the pattern, reduces the expression and explains the nearest wrong answer. Appendix B is the pattern atlas, one entry per lesson, with the cue that calls for the pattern, its canonical object and the check that catches the most common misformulation. Appendix C holds the four-stage worksheet, the expanded seven-step drafting sheet and the verification checklist. Appendix D freezes the system fact sheets and the notation, and Appendix E maps each problem class to the families of methods that solve it.

Contents

How to read this book ii

I Parse the Engineering Question 1

1 From an Engineering Concern to a Mathematical Question 2

2 Quantities, Events, Parameters, and Outcomes 14

3 From Engineering Language to Events and Logic 26

4 Scope, Evidence, Horizon, and Answerability 38

II Formulate Canonical Probability Statements 49

5 Thresholds, Margins, and Event Probabilities 50

6 Compound Events, Dependence, and Counts 60

7 Expectations and Measures of Uncertainty 70

8 Distributions, Quantiles, and Intervals 82

9 Comparisons, Differences, and Rankings 94

10 Conditioning, Evidence, and Inverse Questions 105

III Reduce to General Solution Patterns 116

11 The General Solution-Pattern Framework 117

12 Posterior Distributions and Engineering Quantities 128

13 Transformations, Marginalization, and Derived Distributions 138

14 Prediction, Paths, and First Passage 148

15 Explanations, Hypotheses, and Model Comparison 158

16 One Model, Many Questions 168

IV Formulate Decisions and Acceptance Rules 177

17 From Probabilistic Statements to Decisions 178

18 Acceptance, Assurance, and Probabilistic Requirements 190

19 False Acceptance, False Rejection, and Test Performance 200

20 Expected Loss, Intervention, and Constrained Choice 210

V Formulate Experiments and Information Questions 219

21 What Can a Measurement Actually Tell Us? 220

22 Information Gain and Decision Value 230

23 Designing the Next Experiment 240

VI Verify and Integrate 249

24 Verification, Failure Modes, and Repair 250

25 Integrated Engineering Assurance 261

A Solutions and Commentary 275

B Pattern Atlas 320

C Formulation and Verification Worksheets 325

D System Fact Sheets and Notation 329

E From Problem Class to Solution Literature 335

Part I

Parse the Engineering Question

Lesson 1

From an Engineering Concern to a Mathematical Question

Last month a packaging line stopped for forty minutes. The header pressure of the plant's compressed-air system sagged while compressor B was down on a cooling-water fault, the line's valves lost their supply, and the shift lost its output. At the Tuesday production meeting the production manager turns to the maintenance manager and asks the question this chapter is about: "Is the compressed-air system reliable enough?"

Everyone at the table understands the sentence. Nobody at the table could yet write down what would answer it. The maintenance manager has a vendor report on the desk that puts each compressor at 99.2 percent availability, and that number is true, and it does not answer the question either, because the line did not stop for want of a compressor. It stopped because both compressors drink from one cooling-water circuit. Two competent engineers could leave that meeting, study "reliability" for a week, come back with two different numbers, and neither would be wrong, because they would have answered two different questions.

Engineering work often begins this way, with a topic rather than a question: reliability, risk, uncertainty, degradation, safety. A topic directs attention. It does not say what answer is wanted. Before anyone chooses a model or a method, the concern has to become a request that another engineer reads the same way the first one did. This book teaches one practice for doing that, and Figure 1.1 draws it: four stages, each named for what it produces.

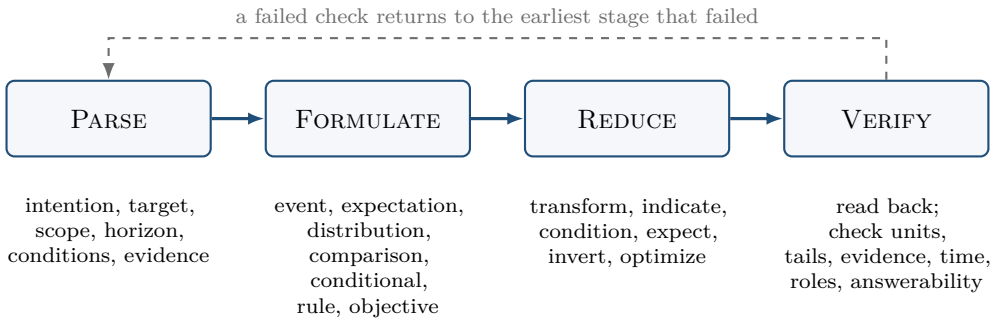


Figure 1.1. The four-stage practice. Parse produces the request, Formulate the mathematical object, Reduce the operations and model outputs that evaluate it, and Verify the read-back that tests it against the request. A failed check returns to the earliest stage that failed.

The practice ends with a mathematical object and a clear hand-off, not necessarily with a number. For some sentences, and the one above is one of them, the correct product of the work is not an expression at all but a statement of what is missing. That is an unusual thing to say in a technical subject, so it is said here, in advance.

Learning aim and prerequisites. After this lesson, you should be able to:

- tell an engineering topic from an answerable question;
- name the intention behind a request;
- keep the engineering intention, the mathematical question, the problem class and the solution strategy on separate lines;
- write a short question record in words; and
- say what kind of mathematical object could answer the request.

No probability notation is assumed. The one prerequisite is a willingness to stop when a material fact or a material choice is missing, which is harder in a meeting than it sounds.

1.1 The Riverbend concern

The question, and its nearest neighbor

Is the compressed-air system reliable enough? A nearby but different request, heard in the same meeting: how worn is compressor A, really?

The station is the one this book returns to most often, and Figure 1.2 draws it. Two compressors, A and B, feed one air header; either one, by its rating, can carry the ordinary production demand on its own, which is why the drawing claims a spare. Both compressors are cooled by the same cooling-water circuit and fed from the same electrical supply, the two shaded blocks, and a fault in either shaded block reaches both compressors at once. Production demand is served at the header, where the required flow and pressure are measured. The station's facts are frozen in Appendix D.

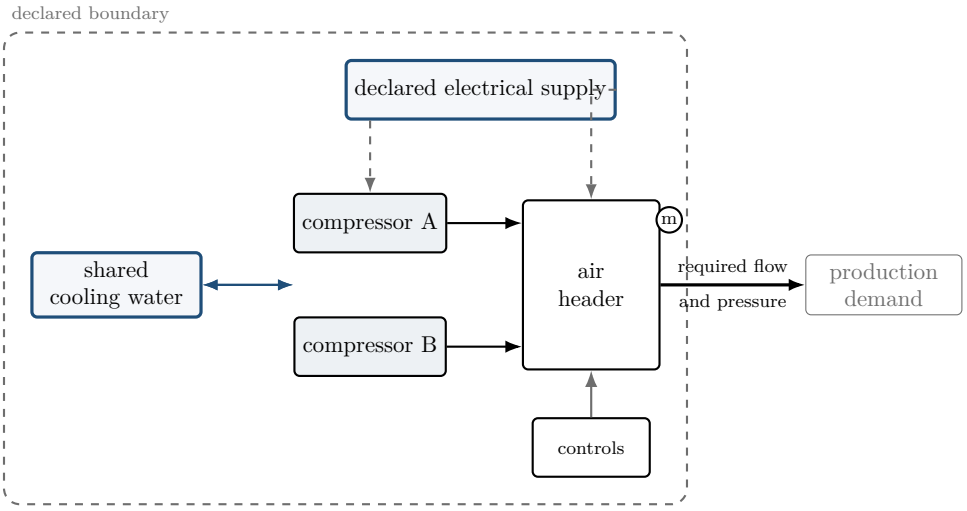


Figure 1.2. The compressed-air station. Two compressors feed one air header; both draw on one cooling-water circuit and one electrical supply, the shaded blocks, so a failure of either shared service reaches both compressors at once. Production demand is served at the header, where the required flow and pressure are measured, and enters the question as a condition.

Read against the drawing, the production manager's sentence names a system and a topic, and leaves open everything that an answer would turn on. It does not say:

- whether the manager wants to know how the system stands today, what it will do next week, how it compares with a proposed change, or what to do about it;
- what counts as success: pressure, flow, product quality, the length of an interruption, or some combination;
- which equipment lies inside the system, and in particular whether the cooling water and the electrical supply that failed last month are in or out;
- when the service must begin and how long it must last;
- which demand and operating conditions apply;
- which inspection, alarm and maintenance records count as evidence; or
- what probability, or what consequence, would be "enough."

The vendor's 99.2 percent is an answer to a question that nobody asked. It is the fraction of calendar hours over which a compressor of that model, tested on the vendor's own bench with its own cooling water, is in a state to run. It says nothing about the header, nothing about the shared circuit, and nothing about eight hours of a peak shift. The first correct response to the production manager is therefore not a formula and not a number. It is a request for the missing question record.

Suppose the manager, pressed, clarifies the request:

> Given the current inspection and operating records, what is the probability that the declared compressed-air system supplies at least the required flow and pressure throughout the next eight-hour peak-production shift, under the approved demand profile and starting state?

Now the requested object has a shape. It is a conditional probability of a future service event over an interval: the system, declared; the service, flow and pressure at the header; the interval, one shift; the conditions, the approved demand and the state the shift starts in; and the evidence, the records on hand. One thing is still missing. To say whether the result is "enough," the production manager must also supply an acceptance level, or a rule that weighs a lost shift against the cost of the actions that could prevent it. That is the manager's to supply, not the engineer's to assume, and Lesson 17 returns to it.

A nearby Riverbend question. The maintenance manager is interested in the same Thursday shift but asks, "How likely is compressor B to be fouled now?" The scope may still include B's cooler and the same inspection records, yet the target is present and the answer is a distribution or event probability for a latent condition. It does not require the future demand path, the behavior of compressor A or the header service definition. Putting the two questions side by side shows why a topic cannot select a model: both sound like "cooling reliability," but one asks what B is now and the other asks what the station will do next.

Section read-back. The cold concern became answerable only after the owner named service, system, time, conditions and evidence. The vendor availability number was not wrong; it belonged to a different unit, boundary and target.

1.2 The four levels

A disciplined hand-off keeps four levels separate. Figure 1.3 stacks them in the order a request descends through them, with what each level holds.

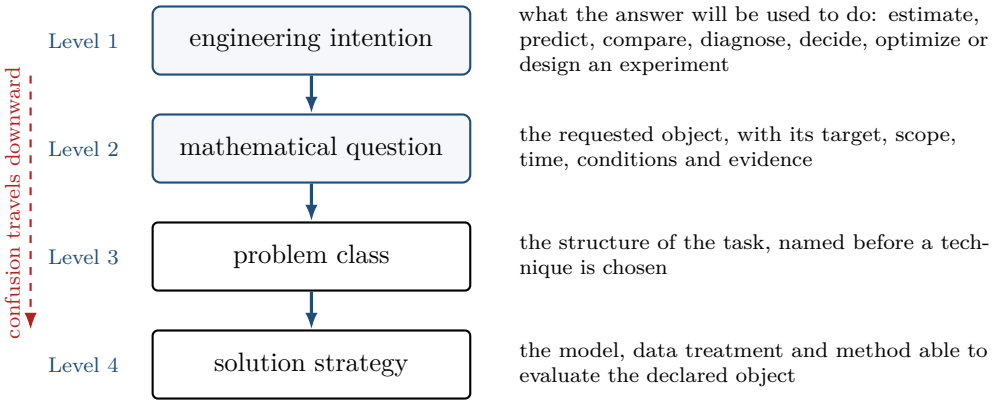


Figure 1.3. The four levels of a hand-off. The two shaded levels are the subject of this book; a request descends through all four, and a mismatch introduced at the top is carried, unseen, into every level beneath it.

1.2.1 Level 1: engineering intention

The intention states what the answer will be used to do: estimate a present quantity, predict a future outcome, compare alternatives, diagnose a cause, decide an action, optimize a design or design an experiment. The production manager's intention is to predict, and then to decide; the colleague who asked how worn compressor A is wants to estimate.

1.2.2 Level 2: mathematical question

The mathematical question names the requested object: an expected value, an event probability, a distribution, a comparison, a conditional probability, a decision rule or a design objective. It carries the target, scope, time, conditions and evidence down from the engineering record without losing any of them. For the clarified request it is a probability of one event, the shift served, given the records and the declared conditions.

1.2.3 Level 3: problem class

The problem class describes the general structure of the mathematical task: prediction, comparison, diagnosis, probability inversion, expected-loss choice or experimental design. Different fields use different names for the same structures. What matters is to recognize the structure before selecting a technique, because the technique is chosen for the structure and not for the words.

1.2.4 Level 4: solution strategy

The solution strategy chooses a model, a data treatment and a computational method able to evaluate the declared object. Simulation, quadrature, an optimization solver and a laboratory protocol all live here. They do not replace the question at Levels 1 and 2; they serve it.

Table 1.1 carries the station's question down the ladder. Confusion usually travels the same way. A team hears "risk," reaches for a simulation it already owns, runs it, and reports the average number of minutes of low pressure per shift. The production manager wanted the probability that a shift is lost at all. The number was computed correctly and answers a

question nobody asked, and because the levels were never written down, nobody can point to the line where the two parted. Keeping the levels visible makes that mismatch reviewable.

Table 1.1. The station's question carried through the four levels.

Level	The compressed-air station
1. Engineering intention	Will the air hold through the next peak shift, and should anything be done before it?
2. Mathematical question	$\mathbb{P}(A \mid$ the records, the approved demand, the starting state), with A the event that flow and pressure are held at the header throughout the shift
3. Problem class	Predictive probability of an interval service event for a system with shared services
4. Solution strategy	Not chosen here; named in Appendix E

1.2.5 Seven intentions, one system

The station is one physical object, and it supports at least seven legitimate questions, one per intention. Table 1.2 pairs each intention with a sentence someone at the Tuesday meeting might say, and with the kind of object that would answer it. The plant is unchanged from row to row. The requested object is not.

Table 1.2. The seven engineering intentions, each with one sentence about the station and the kind of object that answers it.

Intention	Said at the meeting	Object that answers it
estimate	How worn is compressor A?	a distribution for A's present condition, given its records
predict	Will the air hold through Thursday's peak shift?	the probability of the shift served
compare	Would a second cooling-water pump have prevented last month's stop?	the probability of service with and without it, on the same shift
diagnose	Was it the pump, the strainer or the relief valve?	a probability for each cause, given what was logged
decide	Do we run Thursday, derate the line, or bring B down for inspection first?	an action chosen by its consequences
optimize	What header set point keeps the risk of a lost shift under one in a hundred for the least energy?	a setting chosen by an objective, under a probability limit
design an experiment	Is a load test on Saturday worth the lost production?	a test chosen by what its readings, not yet taken, would change

The third column shows that the intentions lead to objects of different kinds, not merely to different probabilities. Estimating and predicting end in a probability. Comparing and diagnosing end in a probability spread over several alternatives at once. Deciding and optimizing end in a choice, and a probability enters only as one of its ingredients. Designing an experiment ends in a choice as well, but one judged over readings that do not exist yet, which makes it the most unusual object of the seven. A sentence that mixes intentions, such as "how good is the system and what should we do about it," holds more than one question, and each needs its own expression.

The list is also a question ledger for Riverbend. It prevents a meeting from quietly moving from one row to another. A present-condition estimate for compressor B does not, by itself, give the future path of the whole station; that path does not, by itself, compare a proposed pump with the installed system; and neither probability chooses an action until the owner states the consequences that matter. Reusing the same station makes those changes in the requested object visible because the physical facts have not changed.

A four-level mismatch. Suppose an analyst hears the comparison question about a second cooling-water pump and writes "Monte Carlo simulation" on the work order. The method is a Level 4 choice, while the response to compare, the common operating conditions and the direction of improvement still belong to Level 2. "Reliability comparison" is only a Level 3 label. Walking back up the ladder exposes the missing estimand before the team builds two simulations whose outputs cannot be compared. The lesson of the ladder is not that every project needs four documents; it is that a statement from one level cannot fill an empty line at another.

1.3 The short question record

Before any probability notation, the maintenance manager writes one page. It has seven fields, in words, and Table 1.3 gives the question each field asks.

Table 1.3. The seven fields of a short question record.

Field	Question to ask
intention	What will the answer be used to estimate, predict, compare, diagnose, decide, optimize or learn?
target	What quantity, event, comparison, explanation, action or experiment is requested?
scope	Which units, people, components and common services are inside the boundary?
horizon	When does the question begin, and is it about one time or an interval?
conditions	Which load, environment, configuration and starting state apply?
evidence	Which observations and records are already known?
missing	Which threshold, level, consequence, choice or model assumption must an owner supply?

A short record is not bureaucracy. It is the smallest piece of paper that lets two people find out whether they are solving the same problem, and it takes about ten minutes to write. For the station, the page reads:

Intention: predict service during the next peak shift, in order to decide whether to act before it.

Target: continuous delivery at or above the required flow and pressure at the header.

Scope: both compressors, the header, the shared cooling-water circuit, the controls and the declared electrical supply. The cooling water is inside the boundary; last month it was the cause.

Horizon: the next eight-hour peak-production shift, from its start.

Conditions: the approved demand profile and the equipment state at the start of the shift.

Evidence: the current inspection, alarm, maintenance and operating records, including last month's stop.

Missing: the acceptance level or the decision rule; what a short interruption counts as; and whether any external service is excluded.

The last field is the one most often left blank, and it is the one the production manager must fill.

The record should be revised in place as the conversation matures. If the manager later decides that a five-minute pressure dip is tolerable, the target changes from uninterrupted service to a duration-qualified event; that is a change to the question, not a note for the simulation team. If the electrical feeder is declared outside the study, the scope line must identify the feeder as a fixed external condition. If the demand forecast replaces the approved schedule, the conditions and evidence lines must show that future demand is now uncertain. Keeping those changes on the same page prevents a result from being compared with an earlier requirement as though their meanings had stayed fixed. The record is short because it carries decisions, not because its fields are optional.

1.4 The four-stage practice

The record is the product of the first stage. The other three carry it to an expression and back.

1.4.1 Parse

Hear the request at its engineering level. Separate the intentions if there are several, as there are here: a prediction, then a decision. Fill the short record and mark what is ambiguous rather than resolving it silently. The silent resolution is the common failure; it is how the vendor's 99.2 percent came to stand in for an answer.

For example, "Should we run Thursday?" contains a decision intention, while "Will the header hold Thursday?" contains a prediction that may inform it. Parse writes them on two lines instead of pretending that one probability is already the action.

1.4.2 Formulate

Choose the mathematical object that supplies the recorded answer, and define its target, boundary, direction, conditions, time and evidence. For the cold question a service event A is defined first in words, "required flow and pressure are held at the header throughout the declared shift," and only then written:

$$\mathbb{P}(A \,|\, \text{current evidence and declared conditions}).$$

Notation is useful only after the words have fixed what A contains. Written before them, it fixes nothing and looks as if it had.

As a check, replace "throughout" by "at the end" and read the event again. If the stakeholder notices no difference, the service event is not yet understood well enough to formalize.

1.4.3 Reduce

Expose the operations and the model outputs that evaluating the formulation requires. The cold question needs a model of the demand, of each compressor's capacity and of its starting and running behavior, of the shared cooling water, of the common electrical supply, and of the way the records change what is known about each of them. That is a requirements list for a model, not yet a choice of software, and it already says something the meeting did not know: no model that treats the two compressors as independent can answer this question, because the shared circuit is in it.

The reduction is therefore already a useful design review. It says which outputs and relations the digital twin must expose without drawing latent state or probability nodes onto the physical station graph.

1.4.4 Verify

Read the expression back, literally. Check that it says future service, throughout eight hours, for the whole declared system, under the stated conditions and evidence. Ask whether the available model and information can answer it. If not, repair the earliest stage that failed, or state what is missing and hand that back to the owner.

Verification ends with a one-sentence echo: "The result will be a probability of uninterrupted header service, not expected pressure, component availability or a recommended action." That sentence is short enough for the Tuesday meeting and strict enough to catch a changed question.

1.4.5 Pattern map

Table 1.4 lists what is heard in a meeting, what this lesson makes of it, and what the sentence does not yet justify.

Table 1.4. What is heard, what this lesson produces from it, and what is not yet justified by it.

What you hear	Product of this lesson	Not yet justified
"study reliability"	a request for the intention and the question record	a probability model
"will it hold Thursday?"	a future target, its conditions and horizon	a simulation method
"would a second pump have helped?"	a declared response, a direction and a common scope	a ranking from separate averages
"what was it?"	candidate explanations and the evidence target	an exhaustive list of causes
"what do we do?"	actions, consequences and their owner	an optimum without a feasible set
"run a load test"	what the result could teach or change	a test design before the target is named

1.4.6 One concern, several answer types

The phrase “Is it reliable enough?” can lead to several legitimate questions. The cases below keep the compressed-air station fixed and change only the stakeholder’s intention.

Case 1: estimate. The maintenance manager asks for the present probability that compressor B’s cooler is fouled, given the inspection record. The requested answer is a conditional probability about a current latent condition. No future shift or action has yet been introduced.

Case 2: predict. The production manager asks for the probability that pressure and flow remain acceptable throughout tomorrow’s peak shift. The target is a future path event with a declared start and horizon. A present fault probability is relevant evidence but not the answer.

Case 3: compare. The energy manager asks whether the proposed control policy will consume less energy than the current policy under the same demand conditions. This requires a paired contrast and a rule for what counts as meaningful improvement, not two unrelated energy summaries.

Case 4: decide and design. The owner asks which maintenance action has least expected consequence, while the commissioning engineer asks which safe test would most improve that choice. The first request outputs an action; the second outputs an experiment. Both need more than a generic topic label such as “risk analysis.”

Lab read-back. For each case, name the stakeholder, requested output, scope, horizon and information available. Then state the mathematical object before naming any algorithm. A question is not ready while one of those fields remains an unacknowledged owner choice.

1.5 Transfer case: North Fork Bridge

The highways authority’s first message about North Fork is shorter than the production manager’s question: “We need an uncertainty study of bridge corrosion.” It names a system and a topic. The bridge dossier in Appendix D identifies the named asset and connection C-17, but the message does not say whether the authority wants the present remaining section, the probability of an unacceptable state during the next year, a comparison of inspection plans, or a decision about traffic restrictions. Those are four different intentions. A consultant who begins by fitting a corrosion-rate distribution has selected one possible solution strategy before the authority has selected its question.

The authority clarifies the operational concern. Connection C-17 was flagged in the last inspection; the bridge must remain in service for twelve months; and the asset manager wants to know whether its structural margin will stay positive throughout that interval under the declared traffic and exposure conditions. The mathematical question is therefore not “What is the corrosion rate?” A rate is one input to the requested answer. The target is the event that the connection’s margin remains positive from the stated starting date through the end of the twelve-month horizon, conditional on the inspection record and declared conditions. At the four levels, the record now reads: predict continued service in order to

prepare a management choice; request a conditional probability of an interval event; recognize a predictive degradation-and-load problem; and leave the model and computation open until their required outputs are known.

The bridge exposes a common transfer error. The authority also owns records for a family of nominally similar connections. Those records may help learn corrosion behavior, but a population statement such as "two percent of connections reach the limit within a year" is not automatically an answer about C-17. The named connection has its own observed dimensions, inspection signals, environment and loading history. The analysis needs an explicit step that transfers information from the population to this unit. Otherwise the subject changes midway through the work: the owner asks about C-17 and receives a fleet frequency.

The short question record makes the remaining gaps plain. The target is positive structural margin throughout twelve months. The scope includes C-17 and the loads and deterioration mechanisms declared by the bridge owner; whether adjacent members or scour lie inside the boundary must be said. The conditions include the traffic scenario, environmental exposure and any planned repairs. The evidence includes the named inspection, previous measurements and the population record only to the extent that a model connects it to C-17. Still missing may be the exact margin definition, the treatment of temporary overloads, the acceptance level and the actions available if the result is unsatisfactory.

Reading the record back is the decisive test. "We will return the probability that connection C-17 retains positive structural margin at every time during the next twelve months, given the declared evidence and conditions." If the owner instead wanted the distribution of present section loss, the question has changed from prediction to estimation. If the owner wanted to choose between a six-month inspection and immediate repair, the output must be a comparison or an action rather than the single probability. North Fork therefore teaches the same discipline as Riverbend without pretending the two systems share a model: preserve the engineering intention, state the requested object, and choose a solution strategy only after the record is complete.

1.5.1 North Fork question clinic

The authority can now test three requests against the same C-17 fact sheet. First: "How much section has C-17 lost?" This is still incomplete because it does not state the answer time or evidence, but it points toward a present quantity. A repaired version asks for the distribution of present section loss given the named inspection and calibration record. Its answer has physical units and does not by itself claim that the connection is safe.

Second: "Will C-17 be safe next year?" This points toward prediction but leaves endpoint versus throughout service open. A repaired version asks for the probability that structural margin remains positive at every time through the next twelve months under declared traffic and exposure conditions. Its answer is dimensionless, and evaluating it needs a relation from present condition to future deterioration and load.

Third: "Should we repair now or inspect again in six months?" This is a decision question. The two earlier answers may inform it, but they cannot choose between the actions until the authority supplies feasible options, consequences, constraints and what the later inspection could reveal. The mathematical output is an action selected by a rule, not merely a probability.

These three versions are useful because their nouns scarcely change. All mention C-17, corrosion, inspection and the coming year. The intention and requested object do change, and those changes determine which model outputs are required. A long report that never records them would remain less useful than the one-page question record.

The model team receives the record as an interface contract. It can say whether it represents the named connection, current measurement uncertainty, future margin paths and action effects. If it supplies only population corrosion rates, the mismatch is visible before calculation. If the authority changes the horizon or tolerates a temporary restriction, the record is versioned and the previous probability is not compared as though its event were unchanged.

Transfer read-back. Riverbend and North Fork share the four-stage practice, not a solution model. The compressed-air answer concerns two service paths over a shift; the bridge answer concerns one connection's margin path over twelve months. What transfers is the discipline of fixing intention, target, scope, time, conditions and evidence before selecting a computational strategy.

1.6 A compact formulation procedure

Compact procedure

1. Remove the topic words that are standing in for a question.
2. Name the engineering intention, and separate the intentions if there are several.
3. Fill the target, scope, horizon, conditions, evidence and missing fields.
4. State the requested answer in one sentence.
5. Name the kind of mathematical object that could supply that answer.
6. Keep the problem class and the solution strategy on separate lines.
7. Read the record back to the owner before building a model.

1.7 Exercises

1.1. One concern, two questions. Riverbend's owner says, "Is the compressed-air system reliable enough?" Write one present-estimation question and one future-prediction question, and name the object each requests.

1.2. Intention contrast. For compressor B's cooling fault, write an estimation question, a diagnostic question and a decision question. Explain why the three answers cannot silently substitute for one another.

1.3. Reverse formulation. Write the meeting-language question answered by a conditional probability that Riverbend maintains both pressure and flow throughout the next eight-hour peak shift.

1.4. Four levels. Classify these Riverbend phrases: "decide whether to derate Thursday's line," "probability of uninterrupted header service during the shift," "predictive path-event problem," and "evaluate with a dynamic simulation."

1.5. Missing information. The manager asks, "Can Riverbend ride through a cooling-water trip?" Fill the short question record while leaving explicit blanks for trip timing and duration, demand, success criterion, evidence and acceptance.

1.6. North Fork topic repair. Rewrite "study bridge corrosion" as one question about C-17's present margin and one about positive margin throughout the next twelve months. Name the output type of each.

1.7. Named unit versus population. The authority has a one-year exceedance frequency for similar connections. State what additional question must be answered before it can inform C-17.

1.8. Horizon switch. For C-17, contrast probability of positive margin at the twelve-month endpoint with probability of positive margin throughout the twelve months. Give a path that separates the questions.

1.9. Method-first repair. A Riverbend team proposes "run a Monte Carlo reliability study." State what remains unspecified and give one complete target the simulation could evaluate.

1.10. Read-back verification. Write a short question record for North Fork's twelve-month service concern, then read back the exact answer, scope, evidence and time interval the analysis must return.

1.8 Summary and capability check

A topic is not a question. The forty-minute stop and the vendor's 99.2 percent are both facts, and neither is an answer, because no one had yet said what was being asked. Begin with the intention, then record the target, scope, horizon, conditions, evidence and what is missing. Keep the mathematical question separate from its problem-class label and from the strategy used to solve it. The book's four stages turn that record into a formulation, expose the model it needs, and test the result against the sentence it came from.

Capability check

You have met the lesson's capability when you can hear an unfamiliar concern, state in one sentence what answer is requested, name the kind of mathematical object that could provide it, and stop before inventing information the owner has not supplied.

Lesson 2

Quantities, Events, Parameters, and Outcomes

At ten past six, fifty minutes before the peak shift starts, the shift operator at the compressed-air station is on the phone to the production manager. The manager has one question: "What is the header pressure?" The operator reads the control-room display and gives the number on it. Both of them hang up believing that the question has been answered, and neither has noticed that three different things could have been meant. The manager may have wanted the pressure that is in the header at this moment, which nobody has seen directly, only through a transmitter that was last calibrated some months ago. The manager may have wanted to know whether the pressure will still be there at two in the afternoon, when the packaging line is drawing hardest. Or the manager may have wanted a yes or a no: is the pressure above the limit the line needs. The display answers none of these. It shows a reading.

An engineering noun does not determine its mathematical role. "Pressure" at the station names, at different moments in the same conversation, a recorded reading, an uncertain present state, a future path, and a pass or fail against a limit. "Unloader setting" on compressor A names a recorded operating condition in last month's logs and a deliberate choice when the maintenance manager proposes changing it. A sound formulation assigns the role from the question, the information available, and the time at which the answer is wanted. This lesson builds the vocabulary for doing that, and ends with the one-page ledger in which the roles are written down.

Learning aim and prerequisites. After this lesson the reader can take a sentence about a plant, separate what is fixed, what is uncertain, what has been observed, what has not yet happened and what someone can choose, keep a quantity apart from an event defined on it, name the primitive objects a derived quantity needs, and write the result as an object ledger before any probability is written. Lesson 1 is the only prerequisite; this lesson introduces the object vocabulary the rest of the book uses.

2.1 The question fixes the roles

The question, and its nearest neighbor

What is the header pressure now, given the calibrated transmitter reading? A nearby but different question is: what is the probability that header pressure stays above its service limit throughout the next shift?

Write P_0 for the true header pressure at ten past six, Y_P for the number on the display, and $P(t)$ for the pressure path through the shift that has not started yet. In the first question P_0

is the target and Y_P is evidence about it. In the second question $P(t)$ is a future outcome, and the thing asked about is the event

$$A_P = \{P(t) \geq p_* \text{ for every } t \in [t_0, t_0 + H]\},$$

read as: the pressure is at or above the service limit p_* at every instant from the start of the shift, t_0, to its end, $t_0 + H$, with H equal to eight hours. The number the operator read out answers neither question on its own. It answers the first only once the relation between the reading and the true pressure is declared, and the second only through a model that carries the present state forward. The production manager is tempted by the display because it is the only number in the room; a reading feels like a fact about the header, and it is in truth a fact about the transmitter.

The same physical pressure therefore appears in several roles, and the notation should make the roles visible instead of leaning on the noun. A reading, a present state, a future path and a verdict against a limit are four objects, and a sentence that is true of one is not thereby true of the others.

2.1.1 A working object taxonomy

Table 2.1 lists the roles a plant noun can take in a formulation, what each role means, and how it is usually written.

Table 2.1. Object roles, their meaning in a formulation, and a typical representation.

Object role	Meaning in a formulation	Typical representation
fixed given	declared value or condition for this question	threshold p_*, horizon H
uncertain quantity	numerical value not fixed by present information	state X, parameter Θ
parameter	feature indexing a model or population	wear rate Θ of the compressor model
state	condition of a particular system at a stated time	header pressure P_0 at ten past six
outcome	quantity or path not yet realized at the answer time	demand $D(t)$ through the shift
observation	recorded result available as evidence	reading $Y_P = y_P$
action or design	choice controlled by a decision maker	unloader setting a, load test d
hypothesis	declared explanation or model possibility	H_j: the pump, the strainer, the relief valve
derived quantity	function of other objects	margin $M = P_0 - p_*$
event	set of possible states or outcomes	$\{M \geq 0\}$

The categories make no claim about the nature of things. They are bookkeeping roles, and their one job is to stop a formulation from exchanging what is observed for what is unknown, what may happen for what has happened, and what can be chosen for what was merely recorded. Every row in the table can be filled from the station alone, which is the point: one plant, ten roles.

Worked role walk-through. Take the phrase "cooling-water temperature." At 06:10, the true temperature T_0 is a latent state and the transmitter value $Y_T = y_T$ is an observation. The calibration bias B may be a parameter over the shift. The temperature path $T(t)$ after 07:00 is a future outcome. The fixed trip limit t_* is a requirement, the margin $t_* - T(t)$ is a derived quantity and $\{T(t) \le t_* \text{ throughout}\}$ is an event. A proposed pump command is an action. "Fouling" may be one hypothesis offered to explain the reading. No physical noun changed, yet nine roles appeared because the questions and times changed.

The takeaway is practical: never assign a symbol to the noun alone. Assign it to the noun together with role and time. If a meeting uses the same word for several rows, introduce separate symbols before any probability statement is written.

2.2 Following Riverbend's objects through time

The morning phone call becomes easier to audit if its objects are followed from the present into the shift. At ten past six, the true pressure and the condition of each compressor exist but are not known exactly. The display and maintenance records already exist and may be used as evidence. The demand, pressure and cooling-water paths after seven o'clock have not yet occurred. The service requirement is neither a reading nor a path: it is a set of future paths that satisfy the declared limits. The distinctions below are therefore a temporal narrative about one station, not a vocabulary list detached from practice.

2.2.1 Quantities and events are different objects

A quantity returns a value. An event returns true or false once the state is realized. At the station the capacity of whichever compressors are running is C and the header demand is D, and the margin is

$$M = C - D,$$

the capacity in hand less the flow being drawn, in the same units of flow. The margin is a derived quantity. The event

$$A = \{M \ge 0\} = \{C \ge D\}$$

is the set of states in which capacity covers demand. The expectation $\mathbb{E}[M \mid \mathcal{E}]$ answers a location question: how much spare flow, on average, given the records. The probability $\mathbb{P}(A \mid \mathcal{E})$ answers how much belief the records support for the requirement being met. When the maintenance manager reports "we have a comfortable margin," the production manager hears the second and has been told the first. One cannot stand in for the other merely because both are built on the same M.

An event is a set in the space of possible states. The braces in $\{M \ge 0\}$ name that set. The probability operator then maps the set, under stated evidence and conditions, to a number between zero and one.

2.2.2 Parameters, states and future outcomes

A parameter describes a population, a mechanism or a model family. A state describes one system at a stated time. A future outcome is what will be realized after that time. The three are coupled and are not the same.

Compressors A and B are of one model, and the vendor's bench figure of 99.2 percent availability is a statement about that model on the vendor's cooling water: a parameter, written Θ for the model's wear or failure behavior. The present condition of compressor A, after its particular history and last month's events, is a state, $X(t_0)$. Whether A carries the line through Thursday's shift is a future outcome, E_f. A model may relate them as

$$p(e_f, x, \theta \mid \mathcal{E}) = p(e_f \mid x, \theta, C_f)\, p(x, \theta \mid \mathcal{E}),$$

read as: the joint density of the shift outcome, the present state and the parameter, given the records, factors into the outcome given state, parameter and the declared shift condition C_f, times the state and parameter given the records. The records can sharpen what is known about Θ and $X(t_0)$. Predicting E_f needs the first factor as well, a model that connects a worn compressor under the approved demand to what happens at the header. A high posterior probability that Θ is favorable is not, by itself, the probability that Thursday's shift is served. The vendor's number was true and did not answer the production manager's question for exactly this reason.

Worked variation: a favorable fleet parameter. Suppose the updated records put ninety-five percent posterior probability on a low wear-rate regime for this compressor model. That is a statement about Θ. Compressor A may nevertheless start Thursday in a poor state after its own thermal history, and the shared cooling circuit may still remove both machines. To obtain the probability of service, the analysis must average the future service relation over the joint uncertainty in parameter and present state. The example repeats the section's lesson: population behavior, named asset condition and future outcome are linked objects, not synonyms.

2.2.3 Observations and latent quantities

An observation is the recorded result, not the physical quantity it is about. The cooling-water temperature transmitter at the station has a bias B and a reading error ε, so that

$$Y_T = T + B + \varepsilon,$$

the displayed temperature equals the true temperature plus the bias plus the error. Reading $Y_T = y_T$ does not make T equal to y_T. The evidence is the recorded value together with its provenance: which transmitter, when calibrated, where in the circuit. The target may be the true water temperature, the bias itself, or what the water will do during the shift. A measurement relation is needed to connect any of them to the reading.

The distinction catches loose plant language. When the shift operator says "cooling water is at 31 degrees," the sentence may report a display, state a belief about the true temperature, or fix a scenario value for a calculation. The question record must say which, because last month's fault sat exactly in that circuit.

Worked variation: a perfect display. Even if the displayed value has negligible random noise, provenance still matters. A transmitter mounted upstream of the restriction may measure its location perfectly while failing to represent the temperature reaching the compressor coolers. "Perfect measurement" describes a relation between a reading and its local state; it does not prove that the local state is the question's target. The ledger therefore records location and measurement relation as well as numerical uncertainty.

Section read-back. At the answer time, Riverbend has latent present quantities, realized records and future paths. Parameters describe mechanisms or populations; events are sets built from quantities; and observations connect to states only through declared measurement relations.

2.3 Actions, hypotheses and information time

Last month's operating logs record the unloader setting on compressor A as $U = u$. In that record U is an observed operating condition, chosen at the time by an operator reacting to the header. When the maintenance manager now asks what would happen if A were deliberately held at u through the peak shift, u is an action. The first question concerns the distribution of outcomes that went with that setting in the past; the second needs a model of what the setting does. The two differ whenever operators, controllers or the physics itself choose the setting in response to the state of the system, which is what controls at a compressor station do all day.

An action may also create evidence. The Saturday load test is a design choice d; its readings are $\widetilde{\mathcal{E}}$, uncertain until the test is run, and they must not be written into $\mathcal{E}$ in advance.

Worked variation: the same setting twice. On a high-demand afternoon, the controller may choose $u = 80\%$ because the header is already sagging. Historical shifts with $U = 80\%$ can therefore have worse outcomes than shifts with a lower setting even if imposing eighty percent would improve pressure under common conditions. The recorded setting is evidence about what operators did in response to the state. The proposed setting is an action whose effect must be modeled. Matching the numeric value does not match the probabilistic role.

2.3.1 Hypotheses are declared possibilities

After the forty-minute stop, three explanations were on the table: the cooling-water pump, the strainer, the relief valve. A hypothesis H_j is a member of such a stated comparison set,

$$\mathcal{H} = \{H_1, H_2, H_3\},$$

and a posterior probability such as $\mathbb{P}(H_2 \mid \mathcal{E}, \mathcal{H})$, the probability that the strainer was the cause given the logs and given the list, is conditional on that list. It does not show that the list is complete, and it does not show that the strainer alone is an adequate explanation. Writing the list is a modeling commitment, and someone owns it.

Worked variation: an omitted explanation. If every listed hypothesis predicts high water temperature but the observed temperature stayed ordinary, the posterior probabilities still sum to one within $\mathcal{H}$. A large value for the strainer hypothesis can mean only that it is least inconsistent among an inadequate set. Diagnosis therefore needs both relative support and an adequacy check; the object ledger reminds the reviewer that H_j is a declared possibility, not a measured plant node.

2.3.2 Randomness is relative to the information state

"Random" means not fixed in the probability model at the stated information time. The demand the packaging line will draw at two in the afternoon is a future outcome at ten past six. By three it is in the historian and belongs in the evidence. The wear state of compressor A is a fixed fact about a physical machine and is uncertain to everyone in the plant. Every object ledger therefore names its answer time and its evidence state. Writing $\mathcal{E}$ for evidence in hand and $\widetilde{\mathcal{E}}$ for a prospective result keeps a reading that has not been taken from being conditioned on as though it had.

This time stamp also prevents an action from drifting into the evidence set. At ten past six, tomorrow's unloader command is a choice. After tomorrow's shift it is a recorded action, but it remains a value selected in response to the information then available. Its later appearance in a historian does not erase that selection process. Role, time and provenance therefore travel together: a number can move from future to observed without becoming an uncontrolled physical fact.

Section read-back. Actions are chosen, hypotheses are declared and observations are realized at an information time. Their values may later occupy the same historian, but their origins remain different. A sound formulation retains those origins so that historical association, diagnostic support and deliberate intervention are not exchanged.

2.4 The object ledger

Before any probability is written, each symbol goes into a ledger. Table 2.2 is the station's ledger for the two questions in the morning phone call, plus the action the maintenance manager raised.

The ledger also records scope and provenance where the symbol alone cannot: which transmitter, which header, which shift. It becomes the notation table that later formulation and review work from, and it takes a few minutes to write, which is less than the phone call took.

The ledger is checked horizontally as well as vertically. Across a row, the symbol, role, time and units must describe one object. Calling Y_P a state would contradict its realized, transmitter-specific provenance; calling $P(t)$ observed would contradict its future status. Down the table, every derived object must be traceable to primitives already present. The event A_P, for example, needs the path $P(t)$, the limit p_*, the start t_0 and the horizon H. If either time quantity is absent, the event is not fully defined even though its name sounds complete.

The ledger also exposes missing model relations. A row for P_0 and a row for Y_P do not by themselves say how the reading informs the state; a measurement relation is required. Rows for P_0 and $P(t)$ do not say how the present becomes the future; dynamics and future demand

Table 2.2. The station's object ledger for the morning questions.

Symbol	Role	Time/status	Definition and units
P_0	state target	present, uncertain	true header pressure at ten past six, kPa
Y_P	observation	realized	header transmitter reading, kPa, with calibration date
p_*	fixed given	requirement	minimum service pressure at the header, kPa
$P(t)$	outcome path	future	true header pressure over the shift, kPa
A_P	event	future	pressure at or above p_* throughout the shift
u	action	chosen	commanded unloader setting on compressor A, percent

are required. A row for action u does not say how the command changes the path. These gaps are useful findings because they turn "we need a digital twin" into a specific list of interfaces the model must support.

Finally, the ledger is not a physical graph. The station drawing contains compressors, services, pipes, controls and the header because those objects are physically connected. Present pressure, future pressure, sensor readings and service events are coordinated mathematical views of that equipment. They belong in the ledger and time-series displays. Drawing them as additional plant nodes would confuse an inferential relationship with a physical connection.

A ledger review in practice. The production manager points to P_0 and asks why the display cannot simply replace it. The controls engineer points to Y_P and explains the measurement link. The maintenance manager points to $P(t)$ and asks where demand and the shared cooling circuit enter the forecast. The owner points to A_P and asks whether a two-minute excursion counts. Each question lands on a different row or missing relation. That is exactly what the ledger is for: it turns a vague argument about "the pressure number" into reviewable questions about state, evidence, future path and event definition.

The section's lesson can therefore be summarized as a chain. Observations inform latent present objects; present objects and conditions feed future outcomes; derived quantities summarize those outcomes; events select the outcomes that satisfy a requirement; and actions may alter the chain. The ledger does not calculate the answer. It makes the answer's required objects visible before calculation begins.

2.4.1 Two failed ledger reviews

In the first review, a dashboard label reads "Header pressure: 645 kPa" and the draft ledger contains only one symbol P. The team uses P for the display, the true present pressure and the future path. The measurement relation and time index disappear, so an observed number is later propagated as if it were a known initial state. Repair requires three rows—Y_P, P_0 and $P(t)$—plus the relations that connect them. The repair adds no new equipment; it restores distinctions the dashboard label erased.

In the second review, historical unloader position and tomorrow's commanded position share a row marked "input." That label is too weak. The historical value is a selected operating

observation, while the proposed value is an action. A prediction conditional on the observed input may reproduce operator selection; a prediction under the action must represent what imposing it does. Separating the rows identifies the additional action-effect relation and stops an association from being reported as an intervention.

Both failures arise from economy in notation rather than lack of data. The ledger is deliberately redundant with ordinary plant language because that redundancy makes role changes reviewable. A compact symbol table that merges state, reading, path and action saves a few lines and can invalidate an entire formulation.

2.5 Transfer case: Meridian Battery Program

Meridian operates a battery that must support a declared mission profile. Its physical one-line diagram contains cells, contactors, a converter and a load. That diagram answers questions about connection and power flow. It does not contain "capacity," "state of charge," "sensor bias" or "mission success" as pieces of equipment. Those belong in an object ledger because they are roles in a probabilistic question. Keeping the two representations separate is important: the one-line remains a physical graph, while the ledger states what is observed, latent, future, chosen and derived.

At the forecast origin, let X_0 be the battery's true state of charge and K its usable capacity over the mission. Terminal-voltage and current telemetry, $Y_{1:n} = y_{1:n}$, have been recorded. They are observations with timestamps, calibration histories and a measurement relation; they are not the state itself. Depending on the time scale and question, K may be treated as a slowly changing state of this battery or as a parameter held fixed over one mission. That choice must be written in the ledger. It is not settled by the everyday noun "capacity."

The future load $L(t)$ and state-of-charge path $X(t)$ are outcomes at the morning answer time. The end-of-mission reserve

$$R_H = X(t_0 + H) - x_*$$

is a derived quantity. It has the units of the chosen state-of-charge representation and can be positive or negative. Mission success is the event $A_H = \{R_H \geq 0\}$ if only the endpoint matters. If the protection logic requires the reserve never to cross the limit, the event instead is

$$A_H^{\text{path}} = \{X(t) \geq x_* \text{ for every } t \in [t_0, t_0 + H]\}.$$

The two events can differ when the converter curtails and the battery later recovers. A probability for one cannot be reported as the other merely because both are called "mission reliability." The quantity $\mathbb{E}[R_H \mid Y_{1:n} = y_{1:n}]$ is different again: a comfortable mean reserve can coexist with a material probability of a negative reserve.

Meridian is also considering two dispatch settings, a_0 and a_1. A setting found in last month's historian is an observed operating condition. The same setting proposed for tomorrow is an action. Historical outcomes conditional on a_1 need not represent what would happen if the owner deliberately imposed a_1, because operators may have selected it in response to load, temperature or alarms. The action question therefore needs the effect of the dispatch choice on the future state and load served, not merely a column in the historian.

Suppose an unexpected voltage sag appears. The team declares three possible explanations: H_1, reduced usable capacity; H_2, a biased voltage sensor; and H_3, an abnormal load transient. These are hypotheses, not observations and not exhaustive truth. The expression $\mathbb{P}(H_2 \mid Y_{1:n} = y_{1:n}, \mathcal{H})$ reports support for the sensor explanation relative to the declared set $\mathcal{H}$. It does not prove that the sensor is biased, and it does not decide whether to replace the sensor. That decision additionally needs the feasible actions and their consequences.

The ledger brings the narrative together:

Object	Role	Time/status	Meaning
K, X_0	parameter/state	present, latent	usable capacity and initial charge
$Y_{1:n}$	observation	realized	voltage and current telemetry with provenance
$L(t), X(t)$	outcome paths	future	mission load and charge trajectory
R_H	derived quantity	future	terminal reserve relative to x_*
A_H	event	future	declared mission-success set
a	action	chosen	dispatch setting imposed for the mission
H_j	hypothesis	declared	one explanation in $\mathcal{H}$

Read from top to bottom, the ledger shows what a model must connect. Telemetry must inform the present state and possibly capacity; the present objects and the declared load must generate the future path; the path determines reserve and the success event; an action changes that path; and hypotheses explain evidence only relative to their declared comparison set. It also prevents a digital-twin display from turning every number into a node on the physical graph. Topology, telemetry, latent state and future outcomes are coordinated views of one asset, not one undifferentiated diagram.

Worked variation: endpoint reserve or path protection. Suppose Meridian begins above the required charge, dips below x_* during a high-load segment, and ends above it after regenerative recovery. The terminal reserve R_H is positive and the endpoint event holds. The path event fails. Both statements are correct descriptions of the same realized mission. The owner's protection rule determines which event is relevant; the noun "reserve" does not.

Worked variation: telemetry after the mission. Before departure, the future current trace belongs to the outcome path. After return, the same trace is realized telemetry and can inform capacity. Its role has changed with the information time, while its physical provenance remains the same. A review conducted after the mission must not use the completed trace as though it had been available to the morning dispatch decision.

2.5.1 Three Meridian requests from one ledger

The operations engineer first asks, "What is the usable capacity now?" The target is K for this battery at the forecast origin. Voltage and current telemetry are evidence, initial charge is a related latent state and the measurement and battery relations determine how the readings inform K. There is no mission event yet. Reporting the probability of mission success would answer a larger question than the engineer asked.

The scheduler next asks, "Will the reserve stay above its protection limit through tomorrow's profile?" The ledger now needs the future load and state paths, the starting-state posterior and the protection event. Capacity is an input; it is not the answer. The same telemetry is reused, but a dynamic link must carry the present objects into the future. If the available twin returns only terminal state of charge, the throughout event is unsupported even when the initial estimate is excellent.

The owner finally asks, "Should we use conservative dispatch a_0 or aggressive dispatch a_1?" The two settings are actions. The analysis needs future paths under each action and consequences for unserved load, battery wear and mission failure. Historical records at the two settings are observations produced by earlier policies; they cannot silently replace the effect of imposing each action under common conditions. Again the ledger grows because the requested object changes, not because the battery gained new physical components.

These requests also suggest three coordinated digital-twin views. The physical one-line shows cells, converter, contactors and load. A time-series view shows telemetry and future bands against time. An object ledger shows which traces are observations, latent states or outcomes and which settings are actions. A decision view can later compare consequences. Combining all four into one node-link picture would make physical connection, temporal evolution, inference and choice look like the same kind of edge.

The reliability engineer adds a fourth request: "What explains last night's voltage sag?" The sag trace is now realized evidence. Reduced capacity, sensor bias and an abnormal load transient are hypotheses in a declared set. The capacity K remains a continuous latent quantity even when reduced capacity is one explanation; H_1 is not another name for K. A posterior distribution for K can inform support for H_1, while posterior support for H_1 still does not prove that the hypothesis list is adequate.

Write the four outputs beside each other. Present-capacity estimation returns a distribution with energy units. Mission assurance returns a dimensionless event probability. Dispatch choice returns an action label. Diagnosis returns probabilities over declared explanations. A single screen titled "battery health" can display all four, but the labels, conditioning sets and units must remain visible. Otherwise a high mission probability may be mistaken for a capacity estimate, or a likely explanation for a maintenance decision.

The object ledger supports this separation without duplicating the physical asset. Rows for telemetry, state, parameter, path, event, hypothesis and action can be reused across queries. Each query selects a target row and identifies which relations must be traversed from evidence to that target. This is a better foundation for a digital twin than adding state boxes to the one-line: the one-line answers where power flows, while the ledger answers what each probabilistic question treats as known, unknown, future or chosen.

Clinic read-back. For the first request, return a distribution for a present quantity. For the second, return a probability of a future path event. For the third, return an action under declared consequences. The nouns and evidence overlap, but the output type and required model relations do not.

Transfer takeaway. Riverbend pressure and Meridian charge use different physics, yet the object discipline transfers exactly. For every symbol, state whether it is known, observed, latent, future, chosen or derived; give its time, units and provenance; and only then decide whether the answer is a quantity, distribution, event probability, hypothesis comparison or action.

END OF THE PREVIEW

Formulating Probabilistic Engineering Questions

These are the opening pages. The complete book runs 345 pages.

The full PDF opens with a token. Write to **rmarquez123@gmail.com**, name the book, and I will send you one.

One token opens every book on the site.

Ricardo Marquez, Ph.D.

ricardo-marquez.com