


Grounds

How We Know
and Why We Can Be Wrong



ARXO

Grounds

How We Know and Why We Can Be Wrong

PROLOGUE	4
How Do You Know?	
CHAPTER 1	10
Five Readings of One Part	
CHAPTER 2	17
A Measurement with a Biography	
CHAPTER 3	24
Peas Owe the Table Nothing	
CHAPTER 4	31
Pluto Went Nowhere	
CHAPTER 5	37
One Word, Three Definitions	
CHAPTER 6	42
It Can Be Said Shorter	
INTERLUDE	48
Knowledge Has a Date	


Contents

CONTINUED

CHAPTER 7	54
A Formula with Conditions	
CHAPTER 8	61
A Triangle on Paper	
CHAPTER 9	68
A Number That Passed the Test	
CHAPTER 10	74
The Same Three Letters	
CHAPTER 11	80
Three Signs Agree	
CHAPTER 12	87
An Experiment That Could Change Our Minds	
CHAPTER 13	94
A Dombyra Owes Fux No Exam	
INTERLUDE	101
The Word “Knowledge”	
EPILOGUE	107
An Answer Together with Its Grounds	

PROLOGUE

How Do You Know?



Around the dastarkhan — the laid table — they are discussing a trip to Burabay. The tea has been poured for the third time when someone asks:

“So how long does it take to get there?”

Answers come at once.

“Four hours,” says Uncle. “I drove it myself in May.”

“I heard three hours,” says Dana. “A colleague went last weekend.”

“The navigator says three forty,” says Timur, showing everyone his phone screen.

“Two hundred and fifty kilometres. At ninety an hour that is under three hours,” announces the younger brother. He has just done it in his head and is visibly pleased with himself.

Before reading on, decide: which answer would you rely on when planning the trip?

Each speaks with confidence, and each has something to point to: a personal trip, an acquaintance’s story, the navigator, a calculation. Behind one short answer stands personal experience; behind another, trust in someone else’s work. To compare the numbers, we must work out where each came from.



Four answers

“I saw it myself” sounds convincing. Uncle really did drive to Burabay and remembers how long the road took him. But in May

there were road works near Shchuchinsk, he stopped for lunch on the way, and he was never one to hurry. Four hours is how long that trip lasted. To say how long ours will take, we need to find out which circumstances will repeat.

“I was told” points to a person who is not at the table now. Did Dana’s colleague time it herself? Did she drive at night or on a Saturday afternoon? Did she count from leaving the city or from her own front door? She may remember her trip exactly. She may have said “about three hours”, and the hedge got lost in retelling. Her story will help us if we understand what circumstances she described and how precisely she described them.

“It says so” can also settle the argument. There is the number on the screen: three hours forty minutes. But this estimate comes with conditions. Which route is selected? From where, and when, are we meant to leave? What road data is the forecast built on? Such details are easy to miss when the number you want is already before your eyes.

“I worked it out” lets us follow the younger brother’s reasoning: he divided two hundred and fifty kilometres by ninety kilometres an hour. The arithmetic is easy to check. What remains is where the distance came from and why that speed was chosen. Can it be kept on average for the whole way? Are stops accounted for? The division may be flawless, and the trip may still take longer.

Now each answer faces a question of its own. What exactly did the person observe? Where did their source get it? What situation does the written word refer to? What data and assumptions went into the calculation?

Answering these questions lays bare the grounds of a claim — what we stood on in taking it for true. Grounds can be presented, checked, and discussed. Then we see more clearly how convincing an answer is and whether it fits our task.

QUESTIONS FOR THE GROUNDS

I saw it myself

What exactly did the person observe?

I was told

Where does the person they cite get it from?

It is written

Which situation does the writing refer to?

I calculated it

What data and assumptions went into the calculation?

Familiar grounds

Such conversations happen all the time. We recall what we saw ourselves, trust other people's stories, check textbooks and manuals. Much of what we know we never verified in person. Without other people's experience we would start from zero every time. So it pays to understand whom we trust, and on what point exactly.

Sometimes retelling changes the very source of confidence. Monopoly has a popular house rule: taxes and fees pile up in the middle of the board, and a player landing on Free Parking takes the accumulated money. Under the basic rules that square pays nothing. But if you were taught to play it differently since childhood, it is easy to mistake a house custom for a rule from the manual.

"We always played it that way."

A perfectly convincing account of how you know. To find out what the manual says, you will still have to open it.

A text itself has a history too. In George Orwell's *Animal Farm* the commandments on the wall change bit by bit, and the animals find it harder and harder to tell what was written before. The familiar inscription no longer matches what they remember.

Texts also change for perfectly open reasons: handbooks are reissued, definitions refined, tables corrected. So sometimes naming the book is not enough. We need to know which edition lies before us and whether the passage we need applies to our case.

We meet calculations every day: working out travel time, comparing prices, adding up expenses. A written-out calculation can be repeated and checked step by step. But correct arithmetic does not yet mean we picked a fitting method and the right data. A result stands together with the conditions under which it was obtained.

One claim of your own

Recall something ordinary you are sure of. That your house was built in 1968, say. That in the mornings another road gets you to work faster. Or that saying “baram” — “I am going” in Kazakh — is wrong.

Write your claim down in one sentence and keep it. In the epilogue we will come back to it. As you read, you will learn to ask sharper questions of its grounds — the same ones we just asked around the dastarkhan. Your opinion may stay the same. Then you will understand better why you count it as grounded.

What we will do

We will take apart claims that look obvious and see what they hang on. We will swap one instrument reading and repeat a calculation. We will pin down the definition of a planet and find out why the number of planets in the textbook changed. We will check the conditions of a familiar formula. We will weigh a report about a mushroom someone confidently called edible, and take apart how that moves the conclusion. Every story offers a chance to see what supports a claim and what casts doubt on it.

Gradually we will learn to tell five things apart: what exactly is claimed; what information is offered in support; what conditions are assumed; how the conclusion was reached; what could make us revise it. The last question pays to ask in advance, while the new result is still unknown and no answer waits to be defended.

Arxo, a system for tracing how a conclusion depends on stated rules and information, will assist these examinations. We will be able to follow the course of reasoning, change an initial condi-

tion, and find out what data an answer is missing. Observing, and judging how far the initial information can be trusted, still takes separate work: a successful calculation does not confirm them by itself.

The name Arxo recalls the Greek *archē* — beginning, source, first principle. With time we get used to what we know and stop noticing what holds it up. The question of grounds takes us back to its sources: experience, evidence, and accepted propositions. Why do we trust them? Are they enough? What is left to find out?

Where we start

On the table lies a small part. It was measured five times, and the five readings came out different. Dana is trying to work out which one to write down. Timur suggests taking the average.


It comes out at exactly ten millimetres. A fine round number.

What do we know about the part now?

PART I · HOW DO YOU KNOW?

CHAPTER 1

Five Readings of One Part





In a college workshop Dana is measuring a small steel part with a caliper. She needs its diameter: the part must fit a hole, and even a tenth of a millimetre may be one too many. The instrument reads 10.0 mm. Dana lifts it off and measures again. Now 10.2. The third time it comes out 9.8, the fourth 10.1, the fifth 9.9.

5

MEASUREMENTS

One part, and already five numbers.

“So which one is right?” Dana asks.

“Take the average,” Timur suggests. “Add them up and divide by five.”

The sum is 50.0 mm. Divide by five and you get exactly 10.0 mm. Timur is pleased: five different readings reduced to one round number. That is the one begging to be written into the logbook.

Before reading on, decide what you would write down: one number, all five, or something else.

What the average hides

Timur’s suggestion makes sense. The average takes all five readings into account and gives a handy starting point for estimating


the diameter. But before accepting it as the measurement result, look at what was lost in the adding and dividing.

Picture two more notebooks. In one, 10.0 mm stands written five times. In the other, the entries run 9.5, 10.5, 9.6, 10.4, 10.0 mm. The average is the same in all three notebooks. But the readings differ: in one series they agreed, in another they spread over a whole millimetre. Behind the entry “10.0 mm” that difference is invisible.

If you chose to keep all five numbers, you can always go back and check the calculation. That is useful. But comparing results also needs a short entry that shows the spread at a glance. The average alone is too little for that.

Let us find a number that shows how wide the readings spread. And then work out what it tells us about the average itself.

FIVE READINGS ON ONE SCALE



Readings ordered by size. One part.

How far the readings spread

Subtract the average from each reading. The differences are called deviations.

Reading, mm	Deviation from the average, mm	Squared deviation, mm ²
10.0	0	0
10.2	+0.2	0.04
9.8	-0.2	0.04
10.1	+0.1	0.01
9.9	-0.1	0.01

Add the deviations and you get zero. Positive and negative values cancel each other — as they will for any set of readings. So to estimate the spread we take the squared deviations instead: they are never negative and cannot cancel out in a sum. The last column totals 0.10 mm².

From here we follow the method of the international *Guide to the Expression of Uncertainty in Measurement*, known by its short name GUM. The sum of squares is divided by the number of readings minus one. In our case, by four.

Where did that one go? We computed the average from the same five readings, so their deviations are tied together. If four of them are known, the fifth is fixed: the sum must come out zero. The divisor “five minus one” accounts for that tie when estimating the spread.

We get $0.10 / 4 = 0.025$ mm². That is the sample variance. Take the square root to express the spread in millimetres: about 0.16 mm. This quantity is called the standard deviation.

It characterises the spread of readings in our series. It is handy for comparing series: the same calculation for the notebook with readings from 9.5 to 10.5 mm gives about 0.45 mm. Same average, but a standard deviation nearly three times larger. Now two numbers show the difference.

What we learned about the average

Dana writes the average of five readings into the logbook. If five more measurements are made, the numbers will most likely come out slightly different. The average will move too. How far can it wander from series to series?

The answer depends on how we measure. Suppose it is the same diameter, the measurement conditions do not change, and the random wobbles of the readings are independent of each other. Then averaging damps down those wobbles. In a real workshop these assumptions would have to be justified by looking into how the measurements were actually made — the five numbers alone give no such guarantee.

Next, divide the variance we found by the number of readings: $0.025 / 5 = 0.005 \text{ mm}^2$. Its root is about 0.07 mm. That is the estimated standard uncertainty of the average from the spread of readings. It characterises how the average might wobble when such series are repeated.

The two numbers have different jobs. The standard deviation of 0.16 mm describes the spread of single readings. The 0.07 mm helps estimate how that spread carries over to the average of five. This way of evaluating uncertainty — from a run of observations — is called Type A in the Guide.

Dana adds a line: “Five readings; the average is 10.0 mm; the standard uncertainty of the average from the spread is about 0.07 mm.”

The word “uncertainty” here is easy to misread as an admission of error. How far the result lies from the exact diameter, we do not know yet. We estimated how random spread affects the averaged result. This calculation will not catch a shift shared by all readings: if the instrument keeps overstating by the same amount, averaging preserves the shift.

What five numbers do not say

At the start of the story Dana needed to know whether the part would fit the hole. Did the computed average bring us closer to an answer?

Closer — but a decision still takes work with the instrument. Suppose its scale reads in 0.1 mm steps. On such a scale close values inside one division cannot be told apart. Even five identical entries would not mean the diameter is known exactly. We also need to find out whether the instrument was checked and whether its readings need a correction.

The part itself raises questions too. Where did Dana put the caliper's jaws? What if the cross-section is slightly oval? Then she measured different sizes in different directions. Before averaging them, we must decide which size we are after. For fitting a hole, one average may not be enough.

To estimate the other components of uncertainty, we can turn to the instrument's data sheet or its calibration certificate. Evaluation from such information is called Type B evaluation. The letters A and B mark ways of evaluating: they differ in where the information comes from and how it is handled. Some influences may already show in the spread of readings. In the combined evaluation each must be counted only once.

For now our 0.07 mm covers only this series. It is tempting to shorten the entry to " 10.0 ± 0.07 mm" and decide the diameter must lie between 9.93 and 10.07 mm. But the calculation gives no such guarantee. Even a complete standard-uncertainty evaluation by itself sets no boundary the value cannot cross.

Now Dana can say more precisely what she knows. The average of five readings is 10.0 mm. Their standard deviation is about 0.16 mm. Under the assumed conditions, the standard uncertainty of the average evaluated from the spread is about 0.07 mm. The calculation can be repeated from the notebook. But to check its conditions and draw conclusions about the diameter, we will have to go back to the instrument and the part.

The round number stayed the same. But now we understand how it was obtained, what stands behind it, and what is still missing for the task.


An estimate can be defended even when the exact value is unknown.

In the next chapter Dana will receive a laboratory report on measuring the same part. It will hold a number with many digits after the point, another instrument's name, and a specialist's signature. Will that mean more is known?

PART I · HOW DO YOU KNOW?

CHAPTER 2

A Measurement with a Biography



A week later the part comes back from the laboratory. Attached is a report: diameter 9.974 mm, standard uncertainty 0.004 mm. Dana opens her logbook. Her five-reading average is 10.0 mm, with a spread-based uncertainty of about 0.07 mm. The laboratory number looks more convincing: more digits after the point, and a smaller uncertainty.

9.974

MM

Diameter per the laboratory report.

0.004

MM

Standard uncertainty of the result.

The reported estimate refers to a defined quantity and measurement conditions.

“Copy it over,” says Dana. “Theirs is more precise.”

“How do they know?” asks Timur.

The question sounds like nitpicking. The laboratory has better equipment, and its people measure every day. Dana sent the part there precisely to get a more reliable result. Now it lies before her.

Before reading on, decide: would you replace the logbook entry with the report’s number?



A number with more digits

Two more zeroes can be appended to the entry “10.0”. That gives “10.000”, though we learn nothing new about the part.

The digit count tells how a result is written.

Judging the measurement itself takes other information.

In the last chapter we evaluated the uncertainty of an average from the spread of five readings. The laboratory report states an uncertainty too, but it claims to cover every component that matters for this measurement. Setting it against a spread-only estimate and declaring victory in precision is premature.

Timur suggests reading the whole report. What exactly was measured? How did they arrive at 9.974? What stands behind the 0.004 estimate? These are the same questions about grounds we asked in the prologue. Only now the answer rests on the work of people far from us.

What exactly was measured

The first question begs a quick answer: the diameter, of course. But we already noticed that the word may not be enough. If the part’s cross-section is slightly oval, the size depends on the direction of measurement. If the part tapers toward the end, the place matters too. Two people can carefully measure one part and get different numbers, because they measured different sizes.

So the measurand needs a description detailed enough: which size of the part, in which section and direction, we are after. Then we can tell whether the report’s number refers to what Dana measured, and whether it will help settle the hole question.

Conditions matter too. When heated, a steel part expands. Geometric sizes are customarily stated at the standard reference temperature of 20 °C. Measuring may be done at another temperature, as long as its influence is accounted for. So the report must

distinguish the temperature they measured at from the one the result was referred to. How much the difference matters depends on the conditions and the uncertainty required. The bare entry “9.974 mm” alone cannot tell us that.

How a reading becomes a result

The next question is about the instrument. How is it known the way its readings relate to the part’s size? That takes calibration.

Imagine the instrument measuring the length of a reference standard. Its value is known with a stated uncertainty. It is compared against the instrument’s readings, allowing for the uncertainty of both sets of data. Then, from the established relation, they work out how to get a result from a reading — what correction to apply, for example. That is a fair picture of calibration in two stages.

The instrument itself may stay exactly as it was. Calibration reveals its characteristics so they can be allowed for in measurement. It does not by itself remove a bias of the readings, nor make the scale finer. Calibrate Dana’s caliper, and there will be data for corrections and uncertainty evaluation. The division value will not change from that.

And how is the standard’s value known, the one the instrument was compared against? It too came from a calibration, resting on another standard. From a laboratory measurement we can thus walk back to the original support. For length, such a chain ties a result to the definition of the metre through the speed of light in vacuum and the second. From that definition the unit is realised in practice. Through standards and calibrations its size passes on to the instruments that measure parts.

Every stage has its own uncertainty. Its contribution must be counted into the final result together with the influence of the measurement’s own conditions. All these quantities combine by the rules of uncertainty calculation; plainly adding up the numbers from the certificates will not do here.


A chain you can present

When a result is tied to a stated reference through a documented, unbroken chain of calibrations, that is called metrological traceability. In such a chain every link's contribution to uncertainty must be accounted for. Traceability belongs to a specific result: the whole path tying it to the reference must be shown.

So the phrase “the instrument is certified” is not enough. We need to know what the certificate attests, which standards were used, and how our result connects to them. The document itself may hold the needed information and references. Their content matters more than the paper's mere existence. We must also pin down what the result is traceable to. A tie to a particular specimen does not yet mean a tie to a unit of the International System of Units — the SI.

Now Timur's question sharpens. Which documents does the laboratory rest on? What do they confirm? How were the values obtained that its instrument's readings were compared against? We need not repeat every calibration in person, but we must understand what exactly we entrust to the specialists, and on what grounds.

A familiar image helps here: a family tree. In the Kazakh tradition of *jeti ata* seven generations of forefathers down the father's line are remembered by custom. Names are joined by kinship ties: through them a person understands his place in the family's history. In a measurement's history we are after ties too. A calibration certificate's number becomes useful once it is clear which standard and which comparison it links to. Behind every link of such a chain stand calibration records and calculations, by which the tie between measurements can be checked.



Three different questions

Dana wanted to simply copy the number over. Now she faces three questions to work through.

First — is the laboratory result traceable to the claimed reference? That depends on whether the calibration chain is confirmed and its links counted into the uncertainty evaluation. If yes, we know which measurements the result stands on. That does not yet rule out mistakes in handling the part itself.

Second — is this result enough for our task? The part must enter the hole. That takes the hole's dimensions, the part's allowed sizes, and the rule by which fitness is decided allowing for uncertainty. Even the small 0.004 mm will not settle the task if the wrong size was measured or the result sits too close to the allowed limit. Traceability helps justify a result; fitness shows whether we can use it for the intended purpose.

Third — do the two measurements agree? The gap between 10.0 and 9.974 is 0.026 mm. It is smaller than even the uncertainty Dana evaluated from the spread of readings. But for a verdict of compatibility that comparison is not enough. We must make sure one and the same quantity was measured, evaluate the uncertainty of the difference allowing for both results, and

pick a comparison criterion. Shared sources of uncertainty may move the answer too.

Dana's evaluation is still incomplete, so declaring the results compatible is premature. The bare difference of two numbers proves no one wrong either. We have yet to find out how large the disagreement is against what these measurements can be expected to show.

Dana keeps her readings in the logbook and attaches the report. Beside them she writes down what information still needs pinning down. That preserves the history of both measurements, and later it will show why a particular result was picked for the job.

A question for the laboratory

It helps to start with one question: which information in the report refers to our part and the conditions of its measurement?

The instrument's calibration certificate may have been issued a year ago. It holds information obtained back then. For today's result we also need to know how the instrument was used, which corrections were applied, which part of the detail was measured, and whether the instrument's needed characteristics still held. All of that ties the old calibration to today's number.

We started with two entries and wanted to pick the one that looks more precise. Now the laboratory result grows a biography: what was measured, how it was done, and which earlier measurements it rests on. From it we see more clearly what can be trusted and what still needs clarifying. That is how someone else's work becomes a ground of knowledge we can share.

In the next chapter steel gives way to peas on the table, and the notebook holds a table of possible trait combinations. Each combination has its share. The peas, however, never read the table.

PART I · HOW DO YOU KNOW?

CHAPTER 3

Peas Owe the Table Nothing



In Dana's notebook there is a square of four cells — the combination table familiar from biology lessons. Beside it stands written: "Three to one". Cross two plants, each carrying the hereditary factors for both smooth and wrinkled seeds, and the teaching scheme gives exactly this ratio: three smooth to one wrinkled.

On the table lie twenty seeds from a garden plot. By the description, they come from such hybrid plants. Dana sorts the peas into two piles: thirteen smooth seeds in one, seven wrinkled in the other.

"There should be fifteen and five," says Timur. "Unless the peas never read the textbook."

Before reading on, decide: is the table wrong, is something off with the peas, or are we asking too much of them?



What the table lists

First, let us sort out the four cells. Call the two variants of the hereditary factor A and a. In our scheme each parent carries both variants and passes one of them to the offspring. One variant arrives from the second parent too. Here are the pairs that can result:

	A from the second parent	a from the second parent
A from the first parent	AA — smooth seed	Aa — smooth seed
a from the first parent	aA — smooth seed	aa — wrinkled seed

The pairs Aa and aA hold the same variants. They stand written apart here to show which parent each came from. Where A is present, the scheme makes the seed smooth; wrinkled it is under the aa match. So smooth seeds own three cells out of four. Their share is exactly three quarters.

So far we only counted cells. Counting them as equally likely takes conditions: each parent passes A and a with equal probabilities, and one parent’s pick does not depend on which variant came from the other. The trait must also show itself as the table describes. Under these conditions the probability of a smooth seed is $3/4$.

Each cell shows a possible combination. The table sets no order in which the combinations must appear. It does not demand that three smooth seeds be followed by a wrinkled one.

What twenty seeds say

Thirteen out of twenty is 0.65. Three quarters is 0.75. Timur rightly noticed the shares did not match. But he read too exact a promise into the calculation.

If a smooth seed’s probability is $3/4$, the expected number of such seeds among twenty is $20 \times 3/4 = 15$. This computation needs neither hundreds nor thousands of seeds. Fifteen is the average number of smooth seeds we would expect from repeating such trials many times over. A single handful may hold more or fewer.

A coin makes this familiar. Even with heads and tails equally likely, twenty tosses owe us no exact ten heads. Twelve or eight may come up without exposing any error in what we assumed



about the coin. Inherited combinations owe a single handful no strict split by expected shares either.

With conditions unchanged and outcomes independent, in a large enough series a wide stray of the share from $3/4$ will be unlikely. That does not mean each next seed must drag us closer to the wanted proportion.

And an exact match even a large series does not promise.

We also need to know how our handful was got. Were the seeds picked with no preference for shape? Do they truly come from the plants the scheme describes? Were seeds of one type lost more often than of the other? A large count alone does not settle these questions.

What seven thousand seeds say

The pattern we are taking apart through the teaching table is known from Gregor Mendel's experiments. In a paper published

in 1866 he described several series of trials with peas. The first series reported concerns seed shape.

From 253 hybrid plants in the experiments' second year Mendel obtained 7,324 seeds: 5,474 round or rounded — the ones we call smooth here — and 1,850 angular wrinkled ones. The paper states the ratio: 2.96 to one.

5474

SEEDS

Round or rounded.

1850

SEEDS

Angular wrinkled.

7324 seeds in total in the series Mendel described; the ratio is 2.96 to one.

Close to three, but still not three.

The round seeds' share in this series is about 0.7474. Far closer to three quarters than our 0.65, yet no exact match again. The published data keeps this difference. We can repeat the ratio calculation and set it against the theoretical one.

Mendel also reports results for single plants. One yielded 43 round seeds and two wrinkled, another 14 and 15. Inside the large series, single groups differed markedly. The ratio of the whole series set no proportion for each plant.

There is one more detail without which the numbers are hard to check: what exactly counted as a round seed? Mendel specifies that this shape covers seeds with shallow dents too. He writes that he observed no transitional forms in the trials. Such details help another person sort seeds by the same traits. A count's result depends, among other things, on how clearly the groups we divide the seeds into are described.

Three different meanings of a share

Now three things can be told apart that are easy to take for one.

First — the share of cells in the table: three out of four. That is a counting result. If the table is drawn as we drew it, arithmetic

gives exactly $3/4$. What remains to check on plants is how well the scheme itself fits our case.

Second — the smooth seed's probability in the accepted model, also $3/4$. It ties the counting to assumptions about inheritance. From it we get the expected share and the expected number of smooth seeds. The expected share is $3/4$ for twenty seeds too. The series size moves the probability of the actual share straying from the expected one.

Third — the share got by counting seeds. Dana's is $13/20$, that is 0.65. Mendel's published series gives $5,474/7,324$, about 0.7474. Behind each of these numbers stand particular seeds and a way of picking and sorting them.

A gap between the expected and the obtained share must be judged allowing for the series size and the trial conditions. It may be an ordinary random wobble. Or it may signal that not all conditions hold: some seeds of one type failed to develop, say, or smooth ones were picked more often. A sorting mistake is possible too. The trait may also inherit in a more tangled way than our scheme assumes.

No gap may be declared random in advance. A small sample can also give reason to doubt the model, and a large one will not fix a picking error. Judging how unusual a result is under accepted assumptions takes statistical methods. We will come back to them. For now it is enough to see what must be checked before concluding.

What we know now

Timur expected fifteen smooth seeds and five wrinkled. His expected-count calculation was right. The mistake crept in when he decided that exactly so many must lie on the table.

Dana writes both results into the notebook: the model expects fifteen smooth seeds out of twenty; our handful holds thirteen. Beside them stay questions about the seeds' origin and the picking conditions. Now it is clear how an observation differs from an expectation, and what grounds each number has.


In the table we check the list of combinations and the count. In the model, the assumptions tying those combinations to inheritance. In the trial, the conditions, the picking of seeds, and the count. That shows where to look for the cause of a gap. One glance at two different shares is not enough for that.

In the first part we asked "how do you know?" of a measurement, a laboratory report, and a teaching table. Each time we had to find out how a number was got and on what conditions it can be trusted. But sometimes disagreement starts even earlier: we mean different things by what we count or name. The second part starts there. First in it comes Pluto.

PART II • WHAT EXACTLY ARE YOU CLAIMING?

CHAPTER 4

Pluto Went Nowhere



Timur has two school atlases at home. In the old one, a 2005 edition, there are nine planets. Last on the list stands Pluto. In the new one, published in 2007, there are eight planets, and Pluto has moved to a neighbouring page, among the dwarf planets.

9

PLANETS

In the 2005 atlas.

8

PLANETS

In the 2007 atlas; Pluto is classed as a dwarf.

The classification changed. Pluto itself stayed put.

“So where did it go?” Timur asks. “Did it blow up or something?”

Nothing of the sort happened to Pluto. It kept circling the Sun on its stretched orbit while astronomers gathered in Prague and discussed what to call it. Pluto itself never found out.

Before reading on, decide: what changed between the two atlases, if Pluto never went away?



Object, observations, word

To sort it out, separate three things: the heavenly body itself, our information about it, and the category we file it under.

Pluto exists regardless of which atlas page we put it on. No decision of astronomers made it smaller, stripped it of moons, or pushed it off its orbit.

Information about it changed as observations piled up. Astronomers refined its size, mass and motion, and studied the neighbouring bodies. In the 1990s and 2000s more and more objects turned up beyond Neptune's orbit. Among them was Eris, discovered in 2005 — a world comparable in size to Pluto. The old picture of the Solar System grew more tangled.

A question of classification arose. If Pluto is a planet, should bodies like it be called planets too? Which likenesses matter for that: size, shape, place among neighbours? The familiar list of nine names no longer showed where the boundary runs.

On 24 August 2006 the International Astronomical Union — the IAU — adopted in Prague a definition of a Solar System planet. Before that there had been no IAU-adopted definition. New observations helped raise the question, and the new definition changed the classification. That change is what our atlases reflect.

What the definition says

Here are the main conditions. Moons stand outside this division. Under the 2006 definition, a Solar System planet is a heavenly body that:

- orbits the Sun;
- has enough mass for its own gravity to give it a nearly round equilibrium shape;
- has cleared the neighbourhood of its orbit.

A dwarf planet meets the first two conditions but has not cleared its orbital neighbourhood, and is not a moon. In this classification planets and dwarf planets form different categories. All other bodies orbiting the Sun, moons excepted, count as small Solar System bodies.

The word “cleared” does not mean the surroundings must be perfectly empty. It means the body has come to dominate its orbital zone through mass and gravity. Planets may keep neigh-

bours. Pluto fails this condition, though it orbits the Sun and is nearly round in shape.

The IAU separately stated that Pluto belongs to the dwarf planets under the adopted definition. And a footnote to the definition listed the eight planets: Mercury, Venus, Earth, Mars, Jupiter, Saturn, Uranus and Neptune.

Here two ways of grounding an answer appear. To learn which category the IAU assigned to Pluto, we can turn to the resolution itself. To see why it meets that category's conditions, we need information about its properties. Citing the resolution settles the accepted classification but does not replace explaining its grounds.



A body not fully known

Now imagine another body, found beyond Neptune's orbit. It is known to orbit the Sun, to be no moon, and to have enough mass for its own gravity to have given it a nearly round equilibrium shape. About its orbital neighbourhood there is no information yet.

Timur wants to send it straight to Pluto's page. It looks alike, after all.

But under the stated conditions the answer turns on whether the body cleared its orbital neighbourhood.

Outward likeness does not settle that.

And no information about clearing does not mean the neighbourhood is uncleared. We do not yet know whether the needed condition holds.

So the answer runs: on the available information, it is not established whether the body is a planet or a dwarf planet. That is no new category of heavenly bodies. It describes what we know.

Such an answer beats a guess: it points at the information needed for the next step. The body's surroundings and its gravitational pull on neighbours remain to be studied. Hand a machine the written definition plus only the stated information, and the correct output is the same: too little data to pick a category.

We will keep coming back to this distinction. There can be grounds to confirm a claim, grounds to refute it, or too little information for either. The last must not be turned into a negative answer just to fill a row in the table.

As of which date

Back to the 2005 atlas. Did it become wrong retroactively? The 2006 resolution itself does not make its earlier entry a mistake. A list of nine planets matched the then-accepted classification. In the new atlas the editors used the IAU definition and got a different list.

Comparing these books, it helps to ask which classification each uses and which time its account refers to. If the 2007 atlas promises to follow the 2006 IAU definition yet counts Pluto among nine planets, that is an inconsistency. But a year on the cover alone cannot judge a whole book's contents: it may, say, narrate the history of astronomy.

Keep the old atlas, and furnish the new edition with an explanation of the change. Then the reader will see why the answers

differ, and will not mistake a classification change for a heavenly body's disappearance.

A definition has an adoption date, information has a time of arrival, a book has a publication year. These dates answer different questions. Between parts two and three we will take them apart in more detail. For now remember: comparing answers means pinning down both the time and the rules they were got by.

What we know now

Timur asked where Pluto went. The short answer: nowhere. Dana opens both atlases at the wanted pages. On one, Pluto joins the familiar list of nine planets; on the other, it counts as a dwarf under the 2006 definition. Now it is clear why the captions differ.

Picking a category rests on which properties of objects we count as essential for comparing them. Applying a definition needs information about those properties. In Pluto's story, observations prompted the classification's revision. The astronomers' decision fixed how the terms are used.

But work remains even after a definition is adopted. How to establish that gravity gave a body its equilibrium shape? What data judges the clearing of orbital neighbourhoods? Near Jupiter's orbit, say, asteroids travel — yet their motion is tied to its gravitational pull, and their very presence does not strike Jupiter off the planets. One literal reading of the word “cleared” is not enough here.


The resolution provides for a procedure to assign borderline objects to dwarf planets or another category. This footnote holds no ready procedure. The definition sets conditions, and checking them takes observations and scientific work.

So the question “where did Pluto go?” leads to another: what exactly do we claim when we call an object a planet? In the next chapter we go further and take apart a word with three definitions at once.

PART II · WHAT EXACTLY ARE YOU CLAIMING?

CHAPTER 5

One Word, Three Definitions



In chemistry class the teacher writes the formula BF_3 on the board — boron trifluoride.

“Is this an acid?”

“No,” says Timur. “There is not even hydrogen in the formula. What would it donate?”

“Yes,” says Dana. “The textbook gives it as an example.”

They share one textbook. The formula on the board is one too. Rechecking it is pointless: neither Timur nor Dana mixed up anything about the substance’s composition. But before deciding who is right, find out what each of them calls an acid.

Three questions under one word

Chemistry textbooks usually introduce three approaches to acids and bases. Each tells them apart by its own trait.

After Arrhenius, an acid in water solution forms hydrogen ions, and a base forms hydroxide ions. In water a hydrogen ion binds to a water molecule, forming the hydronium ion H_3O^+ . That refinement matters to a chemist; for us it suffices to notice that the definition describes a substance’s behaviour in water solution.

After Brønsted–Lowry, an acid donates a proton to another particle, and a base accepts one. The proton here is the H^+ ion. Water is not required for such a reaction. But whatever the substance reacts with must be taken into account: water itself, say, may accept a proton in one reaction and donate one in another.

After Lewis, an acid accepts an electron pair to form a chemical bond, and a base supplies that pair. No hydrogen is needed for this.

These definitions arose from different chemical tasks. Arrhenius cared about water solutions. The Brønsted–Lowry approach allowed describing proton transfer beyond them. Lewis proposed looking at bond formation through a donated electron pair.

So behind a short “is this an acid?” there may stand different questions: what happens in water, who donates a proton, or who accepts an electron pair. Until we pin down the question, a “yes” risks meeting a “no” that belongs to a wholly different property.



What happens with BF_3

Start with Timur's argument. The BF_3 molecule truly holds no hydrogen atoms, so it cannot donate an H^+ ion. By Brønsted–Lowry this molecule is no acid. Here we can answer “no”: the composition is known and rules out exactly the possibility we ask about.

Now look at Dana's example. The boron atom in BF_3 has its outer electron shell short of eight electrons. It can accept an electron pair to form one more bond. The nitrogen atom in the ammonia molecule NH_3 owns a lone pair — two electrons not yet engaged in bonding with another atom.

When BF_3 meets NH_3 , that pair forms a bond between boron and nitrogen. Both electrons for it come from nitrogen, and in the formed bond the pair becomes shared. By Lewis's definition boron trifluoride acts as the acid, ammonia as the base. Here too the “yes” has a perfectly definite ground: BF_3 's ability to accept an electron pair.

The molecule did not change from our opening another textbook chapter.

We examined another of its properties.

And what of Arrhenius's definition? It takes us back to water solution. The reaction with ammonia alone does not say how BF_3 behaves in water. That takes another process to consider.

BF_3 reacts with water, and other chemical particles appear in the solution. So no hydrogen in the starting formula does not yet mean the resulting medium cannot be acidic. Two claims must be kept apart: "the BF_3 molecule can donate a proton" and " BF_3 reacting with water makes an acidic medium". We have taken the first apart already. The second needs information about the reaction with water and its products.

Lacking that information, we should say what exactly is missing. But it would be wrong to declare the substance's behaviour unknown to chemistry only because we have not yet read the needed section.

Three definitions, three questions

Arrhenius: what happens in water? Brønsted-Lowry: who donates the proton? Lewis: who accepts the electron pair? Answering means picking the question that fits the task.

May we pick a convenient definition

So far everything looks suspiciously handy. Want to say "acid" — pick Lewis. Want to say "not an acid" — cite Brønsted and Lowry. Does three definitions mean winning any argument?

Only by swapping the question mid-conversation.

Suppose we discuss the bond forming between BF_3 and NH_3 . Lewis's definition helps describe what both molecules do in this reaction. The remark "but BF_3 donates no proton" is true, yet it does nothing to explain this bond: no proton transfer is at issue here at all.

But if we ask whether the BF_3 molecule itself can be a proton source, its ability to accept an electron pair changes nothing in the answer. It still owns no hydrogen atom to pass on as H^+ .

A definition sets the trait to check. After that we need information about the substance and the reaction. Preference alone — ours, the teacher's, or the textbook author's — is not enough.

Of course, ambiguity may surface after the answer already sounded. Then it is fine to stop and say: “I meant something else. Let us pin it down.” Conversations often advance exactly like that. Trouble starts when a person quietly shifts a word’s meaning while pretending to have proved one and the same thing all along.

Nor do the three approaches split substances into three disjoint lists. They overlap widely, but they describe different sides of reactions. Arrhenius’s approach suits plain water solutions; Brønsted–Lowry suits proton transfer; Lewis suits the bond in our BF_3 and NH_3 example. The choice depends on what we want to understand.

What Timur and Dana argued about

Back to the board. The teacher asks both to manage without the word “acid”.

“ BF_3 cannot donate a proton,” says Timur.

“But it can accept an electron pair to form a bond,” says Dana.

Now these claims can be seen to agree. Timur’s first “not an acid” was too broad, and Dana’s answer lacked a qualification. Once both named the property they meant, the reason to disagree was gone.

In the last chapter we had to pin down by which rule a heavenly body is called a planet. Here we had to ask which property hides behind a familiar chemical term. Sometimes that suffices to see: two answers do not contradict each other at all.

Pinning down, of course, settles no argument by itself. We may agree on every word’s meaning and still part over our estimate of the information — or find there is too little of it. But then at least it is clear what needs checking. While the question’s meaning shifts from speaker to speaker, even a right answer is easy to mistake for a wrong one.

In chemistry definitions have names. In everyday speech there are usually no such hints. In the next chapter we will try to work out what “speaking Kazakh correctly” means. There, one word will hide four different demands at once.

PART II · WHAT EXACTLY ARE YOU CLAIMING?

CHAPTER 6

It Can Be Said Shorter



Timur reads aloud Abai's quatrain, the one assigned for homework. The first line in the book runs: "Talai söz budan buryn köp aitqamyn." Reaching the last word, Timur says "aitqanmyn".

"It says *aitqamyn*," Dana notes.

"That is just the short form. The correct one is *aitqanmyn*."

"Are you reading the poem — or correcting it?"

Timur looks at the book. He only meant to prepare well for the lesson. Instead, Abai has ended up sitting a spelling test.



Two forms side by side

In Kazakh some verb forms live both in full and in shortened shape. Take *baramyn* and *baram* ("I am going"), *kelip qaldy* and *kep qaldy*. Such variants turn up in speech and in literature alike.

Shortening alone does not yet make a form wrong.

Rabigha Syzdyq takes apart these cases in the "Söz sazy" section of her Kazakh language reference. The section covers orthoepy — pronunciation norms. §33 lists full and shortened forms, then gives a direction for reading aloud: keep the variant the text uses.

If *baram* is written, it needs no expanding into *baramyn*. If *baramyn* is given, it should not be shortened to *baram* at will.

The direction works both ways. The full form gains no advantage just for holding more letters.

That is no demand to pronounce everything letter by letter. Spoken speech follows its own sound patterns. The narrower question here is whether one verb form may replace another in reading aloud.

Now Timur's mishap is clearer. He recognised the shortened form but decided it had to be restored. The exercise, though, asked him to read the given text aloud. The reference cites Abai's quatrain precisely for practising such reading.

What changes in the verse

The shortfall could be blamed on metre: the poet ran out of room, shortened the word, and Timur put the extra syllable back and spoiled the line. A handy explanation, but wrong for this pair. *Aitqamyn* and *aitqanmyn* hold three syllables each. A shorter word need not hold fewer syllables.

Keeping the form in reading aloud needs no guessing why the poet picked it. This text holds *aitqamyn*, and the rule says to render that variant. The artistic effect may be discussed apart. Its answer will take other arguments than "the full form is more correct".

But there are cases where shortening does tie to metre and rhyme. The reference's next paragraph takes apart forms like *tildesedi* and *esedi*. The third-person ending *-dy* or *-di* under the stated conditions may shorten to *-d*; a final *-t* pronunciation turns up too: *tildeset*, *eset*.

The example given is Asset's poem "Inju-marjan". There the shortened forms help hold eleven-syllable lines and rhyme with the name "Asset". Here the tie to the verse's workings is explained straight in the source.

And Syzdyq mentions such shortenings not only in verse but now and then in prose too. So a plain border of "allowed in verse, banned in prose" will not hold. The particular use and the reference's stated conditions must be allowed for.

One more distinction: a permission to shorten a form does not demand shortening it every time. From an author's power to use the device it does not follow that a reader may redo every ending he meets. An author composing a line and a person reading it aloud have different tasks.

Four questions instead of one

What did Timur mean by "correct"? Swap it for sharper questions.

Do people say so? That is a question of usage. A form may be heard in conversation, found in a speech record or a literary work. Such observations show it exists. But one fact of usage is not enough to count it fitting for every text.

Is it accepted to speak or write so in this situation? Here we turn to norm and communication conditions. The literary norm covers speech and writing both; a document, a conversation and a poem's performance may ask different things. It cannot be decided in advance that the full form is always standard and the shortened one always outside the norm. That must be checked against a fitting source.

How to read the given text? This question belongs straight to Timur's lesson. For the forms at hand §33 gives a definite direction: keep the full or shortened variant the text uses. That another situation might allow another form does not cancel the reading task.

What does the choice give the author? That is a question of expressiveness. Shortening may build rhythm, carry a rhyme, or render a character's voice. But the explanation must be sought in the work itself. No shortening's cause may be declared to be metre, as we nearly did with *aitqamyn*.

These questions connect, but their answers do not replace one another. A pronunciation reference will not explain every poetic design. A found example will not prove a form fitting for all situations. And a remark on official writing will not yet tell how to perform a poem's line.



The eager corrector

There is nothing unusual in Timur's intent. The better we memorise a rule, the readier we are to apply it everywhere. A short form looks incomplete, and the hand itself reaches out to add the missing letters.

But before correcting, we must understand whether there is a mistake. We may face an allowed variant. Or a character's voice. Or a mistake may truly be there — only taking the particular case apart will establish that, not distrust of everything short.

Timur reads the line once more, now with *aitqamyn*.

"So the full form is wrong?"

"We are talking about this line. It holds *aitqamyn*. That is what we read," Dana answers.

This small refinement keeps us from a new extreme. We do not swap the "expand everything" rule for an "shorten everything" one. We learn to notice which task each direction belongs to.

What "correct" means here

Usage, norm, reading of the given text and artistic choice need different grounds. The answer about the shortened form depends on which of these tasks we are solving.

Before correcting

In the last chapter the word “acid” hid different definitions. Here the word “correct” turned out to hold questions of usage, norm, reading aloud and artistic choice. We divided them and got an answer to one particular question: how Timur should read the line before him.


That took two supports. First we looked at which form the text gives. Then we found the rule for reading it aloud. One rule without the text would have been too little: we would not know which variant to keep. One text without the rule was not enough for Timur either — he saw the shortening and decided to fix it.

Refining need not end every dispute. But it shows which information may change the answer. If another edition holds another form, the editions will have to be taken apart. If we write our own letter, a rule for that task will be needed. A ready answer may not be carried over while the conditions it was got under go unattended.

In this part of the book we worked out what exactly we claim. Next we will ask why it may be believed. But before moving on, stop at one more detail: as of which date does our information hold.

INTERLUDE

Knowledge Has a Date



Picture a laboratory bench and a weight marked “1 kg”. Today is 19 May 2019. Under the definition in force, the kilogram equals the mass of the international prototype — the platinum-iridium cylinder kept near Paris.

The next day, 20 May, a new definition takes effect. Now the kilogram is set through the Planck constant: its exact numerical value in SI units is fixed. The cylinder stays in its vault. Nobody touched the weight. The definition change made it neither heavier nor lighter.

And yet the question “what defines the kilogram?” wants a different answer yesterday and today.



Already adopted, not yet in force

The decision was taken in advance, on 16 November 2018, at the General Conference on Weights and Measures. The resolution named both the new definition and the changeover date — 20 May 2019.

19 May

2019

The definition via the mass of the international prototype applies.

20 May

2019

The definition via the Planck constant takes effect.

The switchover decision was taken earlier, on 16 November 2018.

So months before the changeover one could know how the kilogram would be defined. But that knowledge did not cancel the old definition ahead of the appointed day.

Suppose that on 19 May someone reads the news of the Conference's decision and says: "The kilogram is now defined through the Planck constant." He found a genuine decision and grasped the change's substance. The mistake fits into one word — "now".

That day the exact answer would run: the new definition is adopted but takes effect tomorrow. Saying "no such definition exists" would be wrong too. It already existed as an adopted decision about an upcoming change.

A rule's text and its time of effect are different pieces of information.

Finding the first does not yet establish the second.

Three dates easily confused

To take such a case apart, three questions help.

Which time is at issue? We may ask what defined the kilogram on 19 May 2019, or on 20 May. To answer, the picked date must be matched against the new definition's effective date. In the first case the old one applied, in the second the new one.

Which information do we allow in? The past can be taken apart against everything available to us today. Or another task can be set: to recover what a person could have concluded from the information available to him then. In the second case later publications must not be quietly added. Nor is "was published" the same thing as "was known to this person".

When do we answer? A historical question can be asked a week or many years after the event. In June 2019, say, we may work out which definition applied on 19 May. The answer is given in June but concerns May. A new date on our wall clock does not carry the new definition into the past.

These distinctions do not require three calendar stamps under every phrase. They matter when the answer's meaning depends on time. Mistakes come easiest when a decision is already adopted but not yet in force, or when news of an event arrived much later than the event.

The definition changes, the weight stays

The changeover was arranged to keep the kilogram's old size within the precision the best measurements allowed. Scale users did not have to recalculate all results into another unit. What changed was the definition's ground: instead of one object's mass — a fixed value of a constant of nature.

A material sample's mass may slowly drift, from contamination or cleaning, say. A constant-based definition does not depend on one cylinder's safekeeping. That does not free measurements from uncertainty, but lets unit-realisation methods improve without changing its definition.

The second and the metre were revised differently. Their physical ground stayed: for the second, a fixed transition in the caesium-133 atom; for the metre, the speed of light in vacuum. The wordings were brought to the common form of definitions through fixed constant values. So to the narrow question “which physical ground does the unit rest on?” the answers match on both sides of 20 May. But the definition texts are not verbatim identical.

So a date helps see both change and continuity. It does not oblige us to hunt a difference where none exists.

Knowing then and knowing now

Back to the 2005 atlas from the Pluto chapter. It lists nine planets — under the then-accepted classification. The 2006 decision changed that classification. It did not turn the old atlas into evidence that its makers could not count.

But with claims about nature's workings it is sometimes otherwise. Suppose observations back a general rule, and later

an exception turns up. The old conclusion may well have been grounded in the available information. From this it does not follow that nature obeyed it with no exceptions before the discovery.

We return to this distinction in the genetic-code chapter. For now two thoughts suffice: a person may have had serious grounds for his conclusion, yet the conclusion itself may still have proved too wide. To assess grounds, recover the then-available information. To check the claim, allow in what was learned later too.

A date explains why a person may have reached a certain conclusion. It neither excuses a mistake nor makes the false true.



One more question for a claim

We already hold five questions: what exactly is claimed, which information backs it, which conditions are accepted, how the conclusion was reached, and what would make us revise it. Add one more: as of which date do you know it?

Sometimes it needs refining: are you speaking of how things stood then or now? Do you use the information available at that moment, or allow in later findings? When was the conclusion itself made?


An answer grows no weaker from these refinements. We see better what it stands on and where its scope ends.

On the bench lies the same weight. What changed is the answer to where we take the unit for measuring its mass from. And that change has a date.

PART III · WHY SHOULD THIS BE TRUE?

CHAPTER 7

A Formula with Conditions



In the college laboratory stands a cart with a load. Timur and Dana are preparing an experiment, and meanwhile going over a worked example. The cart's mass with the load is 19.0 kg . A string pulls it horizontally, and it moves in a straight line with an acceleration of 1.5 m/s^2 . What force must be applied?

"Mass times acceleration," says Timur. "Nineteen by one and a half. Twenty-eight and a half newtons."

"And if the dynamometer on the string reads thirty?"

Timur checks the multiplication. It still comes out the same.

"Then the dynamometer is lying."

His arithmetic is right. But before blaming the instrument, look at which force Timur computed and which Dana is about to measure.



What hides behind the letter F

In Newton's second law, force is the resultant of all external forces. They add up with directions allowed for. For our problem pick the system: the cart with the load. The string pulls it forward, and resistance forces point backward. Gravity and the support force cancel each other in this example: there is no vertical acceleration.

We care about motion along the horizontal track. In short: $F = ma$, where F is the resultant along the track, m is the mass of the cart with the load, and a is its acceleration. Substitution gives:

$$19.0 \text{ kg} \times 1.5 \text{ m/s}^2 = 28.5 \text{ N.}$$

The dynamometer on the string measures tension. To get the resultant from it, the remaining forces must be allowed for. If the string pulls with 30 N, say, and total resistance is 1.5 N, the resultant is 28.5 N. Both numbers agree with the calculation.

28.5

NEWTONS

Resultant along the track.

30

NEWTONS

String tension under the assumed 1.5 N resistance.

Worked example: mass 19.0 kg, acceleration 1.5 m/s². The resistance is yet to be determined.

The 1.5 N of resistance we have just assumed for the example. To explain a real instrument's reading, it will have to be determined.

Fitting the gap between two numbers does not yet establish the cause of the disagreement.

Timur answered the question about the resultant. The question about the string's tension stays open: with no information on resistance, one product of mass and acceleration is not enough for it.

Conditions beside the formula

Beyond the meaning of the letter F , two more clarifications are needed.

What happens to the mass? We treat the cart and its fastened load as one system of constant mass. Through the experiment it loses nothing and picks nothing up on the way. The familiar writing of the second law fits such a description.

But if sand pours out of the cart, or a rocket throws off gases, motion with departing matter must be taken apart, allowing for what crosses the system's borders and how it moves. Simply plugging a changing mass into the familiar formula, without asking what crosses the system's borders and how it moves, will not do.

Measured against what is the acceleration? In our example, against the laboratory. For this experiment we count its frame of reference as inertial: a body with zero total external force would keep its velocity there in size and direction. The earthbound laboratory meets this condition only approximately, but for an ordinary cart experiment the approximation is enough.

An accelerating railway car shows why this matters. Imagine the cart first resting in the car, able to roll with practically no resistance. The car starts accelerating forward. The cart lags behind it, and a passenger sees it move backward. That takes no force pulling it toward the rear wall: it is the car's motion that changes.

The passenger may describe what happens against the car. But then the equation of motion needs an inertia force, tied to the acceleration of the chosen frame. Another way is to reckon against a frame that can be counted as inertial. What matters is stating which way we picked.

The formula takes one line. The clarifications run longer, because they tie its letters to a particular cart, its forces, and a way of watching motion.



When numbers disagree

Now change a condition. Let 30 N be named the resultant itself, and take the 19.0 kg mass and 1.5 m/s^2 acceleration, for a start, as exact stated quantities.

These three values do not satisfy $F = ma$: the product is 28.5 N. We have found an inconsistency. Keeping the formula and its conditions of use, all three numbers cannot be accepted at once.

But the computation does not say what exactly to fix. Perhaps the mass is misstated. Perhaps the wrong acceleration got copied. Perhaps tension was again called the resultant. We must go back to where the data came from and what it means.

With real measurements one more question appears. Mass, acceleration and force are known with uncertainty. So the comparison must weigh not only two numbers — 28.5 and 30 — but the uncertainties of the independently measured force and of the computed result. The latter depends both on measuring mass and on measuring acceleration. Under a shared measurement method the ties between the input data have to be allowed for too.

Until that information is in, we cannot say whether the gap matters. First make sure the same quantity is compared, then estimate how exact the comparison is. We will come back to this in the experiment chapter.

Newton's second law cannot be declared refuted over two differing readings. But neither may the experiment's own de-

scription be placed beyond checking in advance. A gap is a reason to check the data, the counted forces, and the accepted approximations.

Units join the calculation too

Even before the experiment we can check whether the formula yields a quantity of the wanted kind. In our product kilograms multiply by metres per second squared. That is exactly what a newton is:

$$1 \text{ N} = 1 \text{ kg}\cdot\text{m}/\text{s}^2.$$

If a computed force comes out written in seconds, something is wrong already: in the formula, the substitution, or the answer's entry. No dynamometer is needed for that verdict.

Take an even simpler case. At a steady 2 m/s for 3 s a body covers 6 m. Speed times time gives length. The answer “6 seconds” fails on dimension. But “9 metres” is wrong too, though its dimension fits. Checking units sifts out some errors; by itself it proves neither the number right nor the formula applicable.

In a school problem the 28.5 N result is usually rounded to 29 N, keeping two significant figures — as many as the 1.5 acceleration entry holds. That is the rounding rule accepted for such a calculation. In a measurement report the rounding place is matched to the result's evaluated uncertainty.

And in both cases the unit stays beside the number. “29 N” names a force. A lone “29” does not yet say what we found.

What exactly the calculation grounded

Timur adds the word “resultant” to his answer.

“So we leave the dynamometer alone for now?”

“First let us find out what to compare its reading against,” says Dana.

Now the answer can be traced from start to finish. We picked the cart with the load, stated constant mass and a frame of reference, named which forces we add. Under these conditions

the second law gave 28.5 N. We checked the multiplication and the units.


Such a calculation shows what follows from accepted premises. Seeing how well it describes the real cart takes measurements and a check of the assumptions. And grounds to trust the physical law itself come from experiment: correct multiplication cannot prove a law of nature.

In the next chapter we move to a triangle on paper. There too are a formula and conditions, but another way of grounding appears — mathematical proof. And no ruler, however good, replaces it.

PART III · WHY SHOULD THIS BE TRUE?

CHAPTER 8

A Triangle on Paper




On a sheet of paper a triangle is drawn. One angle looks right. Dana holds a set square to it: it seems to fit. Timur takes a ruler and measures the sides. Three centimetres, four and five — as far as the markings tell.

“Right-angled,” he says. “Three, four, five. A classic.”

“And if the last side is a touch longer than five?”

“Then measure more precisely.”

Measuring more precisely is indeed possible. But even the best ruler cannot do one thing: turn a measurement result into an exact condition of a mathematical problem.



Look, measure, state

First Dana looked at the angle and held the set square to it. That gave her grounds to count the angle as close to right. How close depends on the instrument, the lines' thickness, and the care of comparison. One “seems to fit” is not enough for an exact estimate.

Then Timur measured the sides and checked the familiar equality:

$$3^2 + 4^2 = 5^2, \text{ that is } 9 + 16 = 25.$$

The computation is right. But it took in rounded measurement results. We do not yet know how far they stray from the lengths we try to determine. The agreement of these numbers' squares does not yet prove the angle on the sheet to be exactly 90° .

Now set another task. Consider a triangle on the Euclidean plane whose sides are exactly 3, 4 and 5 units of length. They may be stated in centimetres too: what matters is that these are exact conditions, not ruler readings.

For such a triangle the equality holds exactly. By the converse of Pythagoras's theorem the angle opposite side 5 is right. Neither pencil-line thickness nor set-square precision affects this conclusion. The drawing helps picture the triangle, but in this problem its properties are set by conditions.

Timur quietly moved from measured lengths to exact numbers. His conclusion about the mathematical 3–4–5 triangle was right. Judging the drawn angle takes returning to measurements and their uncertainty.

Which way the theorem works

Pythagoras's theorem says: in a right triangle the hypotenuse's square equals the sum of the legs' squares. The hypotenuse is the side opposite the right angle; the legs are the two sides forming it.

We already know the angle is right, and get a relation of lengths. Timur went the other way: the length relation is known, and the angle is yet to be determined. That takes the converse theorem. It states: if in a triangle one side's square equals the sum of the two others' squares, the angle opposite that side is right.

In the *Elements'* first book Euclid proves these statements apart: first proposition 47, then 48. The second uses the first but needs extra steps.

Simply swapping assumption and conclusion is not enough.

47

PROPOSITION

From the right angle to the equality of the squares on the sides.

The direct and converse theorems are proved separately; the triangle is on the Euclidean plane.

48

PROPOSITION

From the equality of the squares to the right angle.

In our example both directions are proved, so moving from equal squares to a right angle is allowed. Not every statement works that way. In the next chapter we meet a check built on a property of prime numbers. But it will also be passed by a number that is no prime.

But does the triangle exist?

Swap the lengths for 2, 3 and 4. Such a triangle exists, but it cannot be right-angled: $2^2 + 3^2 = 13$, while $4^2 = 16$. Were there a right angle in it, the longest side would lie opposite, and the equality would hold by Pythagoras's theorem. It does not hold — so there is no right angle.

Here we used the direct theorem: showed the equality a right triangle needs is broken. That is a different move of reasoning than in the 3–4–5 example.

Now take 1, 2 and 3. The squares need not even be reached. The two short sides together match the third in length. Join their ends, and the vertices land on one straight line. No ordinary triangle results.

In a triangle any two sides' sum must strictly exceed the third. Euclid proves this inequality in proposition 20. For checking stated lengths it weeds out cases where a triangle's angles are not yet open for discussion: there is no triangle with such sides at all.

The last triple is 0, 3 and 3. Now the arithmetic looks flawless: $0^2 + 3^2 = 3^2$.

But a triangle's side cannot be of zero length: its two ends would coincide. The strict inequality is broken too: $0 + 3$ is no more than 3. The squares agree, yet the converse theorem cannot apply. Its condition starts with the words "in a triangle".

The difference between the examples runs deep. Sides 2, 3 and 4 set a triangle that is not right-angled. Sides 1, 2 and 3, or 0, 3 and 3, set no non-degenerate triangle at all. One same short "no" would hide different reasons for refusal.

Apply and prove

So far we have used a ready theorem. Why trust the theorem itself?

Ten right triangles can be drawn, their sides measured, their squares checked. A thousand can be drawn. Such observations help spot a pattern, but prove nothing for all triangles. Measurements are approximate, and checked cases are ever only some of all possible.

Mathematical proof is built differently. We accept starting positions and derive the statement from them step by step. In the *Elements*, the proof of Pythagoras's theorem rests on constructions, properties of equal figures, and earlier proved area relations. It concerns an arbitrary right triangle in the geometry at hand.

Drawing a conclusion about the 3–4–5 triangle means applying a known result to a particular case. Proving the theorem means grounding the general move from condition to conclusion. These are different tasks.

The drawing is useful here: it helps see the construction, find the idea, or spot a missed condition. But the step "these segments are equal because they look alike in the picture" is no proof. Equality must follow from the construction or known properties.

If an exact counterexample meeting all its conditions turns up against a stated mathematical claim, hunt the error in the wording or the proof. A measured mismatch on paper needs another check: how exactly the drawing is made, what was measured, and how uncertainty was evaluated. It may refute our

description of the drawing, but by itself it does not refute the theorem about exact objects.

If drawn on a sphere

Until now we spoke of the Euclidean plane. For another geometry the conditions need revisiting.

Picture a perfect sphere with an equator and poles. Run a path from the North Pole to the equator down one meridian, then along the equator for a quarter circle, and back to the pole down another meridian. The result is a spherical triangle with three right angles.

Its sides are arcs of great circles — circles whose centres match the sphere's centre. All three picked arcs match in length. The sum of two such lengths' squares does not equal the third's square, though every angle is right.

That exemplifies another geometry; it is no exception to Pythagoras's theorem on the plane. Before carrying a familiar formula onto a new object, check whether its premises survived the move.



What stayed on the sheet

“So the drawn angle is not right after all?” Timur asks.

“We did not establish that either,” Dana answers.

Now they can put it more precisely. The set square showed no visible deviation. Rounded measurements gave the 3–4–5 triple. A triangle with such exact sides on the Euclidean plane is right-angled. The last statement is proved — but it cannot retroactively make the first two observations exact.


In the last chapter correct multiplication did not replace checking the experiment. Here a correct equality does not replace checking the theorem's premises. To feel a conclusion's force, keep the whole path to it: what is given, which statement is applied, and whether its conditions hold.


In the next chapter the arithmetic will add up again. This time hunt the error in what we decided to conclude from it.

PART III · WHY SHOULD THIS BE TRUE?

CHAPTER 9

A Number That Passed the Test





Timur is checking whether the number 341 is prime. In a book he found a way: raise two to the power one less than the number under check and look at the division remainder. He computes the remainder of 2^{340} divided by 341. It comes out as one.

“Prime,” says Timur.

Dana writes in the margin:

$$341 = 11 \times 31.$$

“And what do we do with that?”

Timur rechecks the remainder. All correct: one. Dana’s product is correct too. They will have to revisit how Timur reached his conclusion.

What the theorem promises

A prime is a natural number above one with only two positive divisors: one and itself. If a number above one has other positive divisors, it is composite.

Timur’s check rests on Fermat’s little theorem. First take it apart on small numbers. Take the prime 7 and base 2 — the number to raise. The exponent is one less than seven, that is six:

$$2^6 = 64 = 9 \times 7 + 1.$$

Sixty-four holds nine sevens, with one left over. That is the division remainder.

The theorem states: if p is prime, and an integer a is not divisible by p , then a to the power $p - 1$ leaves remainder 1

on division by p . Our example fits it: seven is prime, two is not divisible by seven, the remainder is one.

In the check we pick a base above one but below the number under check. If that number is prime, the theorem's condition holds and the remainder must equal 1.

The direction here is fixed: knowing the number is prime, we can predict the remainder. Concluding from the remainder that the number is prime takes a separate ground.

A right remainder, a wrong conclusion

For 341 Timur's result can be checked without writing out the huge number 2^{340} . It suffices to notice:

$$2^{10} = 1024 = 3 \times 341 + 1.$$

So 2^{10} leaves remainder 1 on division by 341. When numbers, each one more than some multiple of 341, are multiplied, the product leaves remainder 1 too. And 2^{340} is a product of thirty-four factors of 2^{10} . The remainder truly is one.

At this point Timur took an extra step: "Primes pass the check. My number passed it. So it is prime."

The theorem allows no such move. It resembles the reasoning "if it rained, the road is wet; the road is wet — so it rained". But the road may have been watered. One same observation fits different explanations.

Dana's note shows this with no analogies. Eleven exceeds one, is less than 341, and divides it with no remainder. So 341 is composite. That conclusion needs only the equality $11 \times 31 = 341$. No need to separately prove both factors prime.

341

COMPOSITE NUMBER

$341 = 11 \times 31$. The passed base-2 test does not undo this factorisation.

The base-2 check is passed, and the number is composite.

These facts do not contradict each other. They refute the attempt to reverse the theorem, not the theorem itself: it never claimed a composite number could never yield remainder 1.



What another base changes

“Maybe two is unlucky?” Timur asks.

Try three. The remainder of 3^{340} divided by 341 is 56. No one this time.

Had 341 been prime, with the picked base the theorem would have had to yield remainder 1. A different number came out. So 341 is not prime. That is one more proof of compositeness, needing no knowledge of the factors 11 and 31.

Now the difference between the check’s two outcomes shows. With its conditions met, a remainder other than 1 lets us reject primality. Remainder 1 establishes a narrower fact: the number passed the test for the given base. That is not enough to decide whether it is prime or composite.

It would be wrong to say a passed check reported nothing at all. We learned a property of the number for sure. The mistake starts when that property is taken for another, stronger one.

And switching bases does not always help. For example, $561 = 3 \times 11 \times 17$ passes the Fermat test for every base coprime to

561. “Coprime” means the two numbers share no positive divisor but one.

So trying only such bases will never expose 561’s compositeness. The hedge cannot be dropped: three, for instance, itself divides 561, and with base 3 no one results. There are other primality checks besides. The Fermat test’s limit does not become a limit of all arithmetic.

Checking the check itself

So far we took the remainders to be computed right. But that too needs establishing. For large powers it is handy to count step by step, keeping only the remainder after each multiplication. On division by 341 such remainders lie between 0 and 340, though intermediate products may run larger.

A long entry invites mistakes. Write, say, that the remainder for 2^5 is 33, though $2^5 = 32$. Later steps may look neat and still rest on a wrong link. One plausible answer at the end is not enough to count the whole calculation correct.

But checking arithmetic is only one task. There is also checking the reasoning’s completeness. We need to make sure it starts from true initial data and reaches exactly the power whose remainder we want to know. Several correct lines by themselves do not guarantee that.

A skipped line need not mean a mistake: the author may have cut an obvious move. What matters is whether that move can be grounded. Our short computation through 2^{10} is complete, though it does not write out every intermediate power. And a long table of remainders may hold a gap that keeps the conclusion from being checked.

Such work can be handed to a program, if it checks both the moves’ rightness and that the whole chain leads to the stated result. The mere absence of errors found does not yet mean the check is complete.

Finally, even a flawless remainder computation will not fix Timur’s logical mistake. Checking the calculation and establishing what follows from it are two separate tasks.

What we learned about the number

Timur crosses out the word “prime” and writes: “Passed the Fermat test for base 2.”

“Now that is right,” says Dana.

Beside it stays her entry: $341 = 11 \times 31$. It answers the question Timur asked at the start: the number is composite.

We got several exact results. The base-2 remainder is 1. The base-3 remainder is 56. The number has the divisor 11. Each result has its ground, and we need to understand which claim it lets us make.


This chapter ends the third part. The formula needed conditions, the geometrical theorem the right conclusion direction, and the numerical check an understanding of what its result means. In all three cases a pretty answer was too little: we had to trace why it follows from what is given.

Next we turn to claims that must be revised as new information arrives. What exactly changes then: the object itself, the conditions, or our understanding?

PART IV · WHAT WOULD MAKE YOU CHANGE YOUR MIND?

CHAPTER 10

The Same Three Letters



On a card is written UGA. Dana knows it is a codon — a triplet of nucleotides in RNA, denoted by three letters. Timur remembers this triplet from the school table of the genetic code.

“Stop,” he says. “Protein assembly ends here.”

Dana turns the card over. On the back it says: “Read in human mitochondria.”

“And now?”

Timur looks at UGA once more. The letters stayed the same. But the note on the back is part of the problem, and it must be allowed for.

How three letters are read

In protein synthesis the ribosome reads information off messenger RNA. Its record uses four letters: A, U, G and C. They are read in triplets, starting from a fixed place. In our problem one whole codon is already picked out — UGA.

From four letters 64 triplets can be made. In the standard table of the genetic code, 61 codons match amino acids of twenty types, and three codons mark termination. So several different codons may match one amino acid.

61

CODONS

Correspond to amino acids.

3

STOP CODONS

Mark termination.

64 triplets in the standard genetic-code table. Other systems need their own table.


UGA in this table is one of the stop codons. Usually it signals the synthesis of the polypeptide chain, from which a protein forms, to finish. The assembled chain does not break in the process: its assembly ends.

But the standard table does not cover every variant of the genetic code. Mitochondria are intracellular structures with their own DNA and their own protein-synthesis machinery. Some

codons in protein synthesis off mitochondrial genes read differently.

In the vertebrate mitochondrial code, humans included, UGA matches tryptophan — one of the amino acids. At that place the next link joins the chain. In the NCBI tables the standard code carries number 1, the vertebrate mitochondrial code number 2.

Now it is clear why the card's back was needed. Timur recalled the right row of the standard table but had not yet established whether it fits this entry. For the case stated on the card the answer is tryptophan.



If the back is blank

Take the note away. Before us again is only UGA. We know it is a codon, but we do not know in which system it is read.

We may answer conditionally: in the standard table it is marked as a stop codon, in the vertebrate mitochondrial table as a tryptophan codon. But picking one assignment for our card from these three letters alone is impossible.

“Not established” here does not mean “UGA cannot be a stop”. Nor does it mean “so it is tryptophan”. We know of different matches but lack the information to apply one of them to the particular case.

Two situations must be told apart. In the first, the biological context is unknown. In the second, it is established and points to another assignment of the codon. The first needs more information. The second needs the answer changed.

Mixing these situations up, it is easy either to invent the missing part or to miss the meaning of information already received. Timur might keep insisting on “stop”, citing the textbook. The textbook’s row would be no wronger for that, but another table would still apply to the card.

The table describes what happens

In the Pluto chapter we also picked the rule the answer depended on. But here the ground for picking is different.

The International Astronomical Union set the criteria of the “planet” category. Observations help determine whether a heavenly body meets them. The genetic code table describes matches that researchers uncover while studying protein synthesis. No decision can make a mitochondrion read a codon differently.

Of course, people invented the letters, the names and the table numbers. But behind the entry UGA → tryptophan stands a checkable claim about a biological process. It does not become true because such an answer suits us.

So for the card, picking a favoured table is not enough. We need to find out where and how the RNA is read, and use the description matching that system. If a table’s prediction diverges from observations, check both the information about the system and whether the description itself suffices.

What became known in 1979

On 8 November 1979 the journal *Nature* published the paper “A different genetic code in human mitochondria”. Its authors compared the sequence of a human mitochondrial gene against the amino-acid sequence of the matching protein from bovine heart. This analysis showed UGA used as a tryptophan codon, not as a synthesis-termination signal.

The publication did not change how mitochondria work.

It reported a difference that had been found. The paper's date belongs to the history of our knowledge, not to the start of a new biological rule.

Here recall the interlude's question: what information was available when the conclusion was drawn?

Observations in studied systems might give grounds to expect familiar matches to hold elsewhere too. Such an expectation may be a sensible working hypothesis. But the claim "UGA always means stop, in any system" is far stronger. Carrying a result over to cases not yet studied needs grounding.

So the date alone cannot judge whether earlier conclusions were well grounded. We need to know what exactly the researcher claimed, what data he rested on, and how widely he spread the conclusion. Nor does one publication date mean everyone knew of it the next day.

New information lets a generalisation be revised. That needs no giving up of correctly described cases. It suffices to admit that the old wording covered more than its grounds allowed. And the refined claim has a narrower meaning: the word "always" is gone from it.

That is how grounding and truth differ. From the available information a person might have serious grounds for an assumption that later failed. New data helps see its boundary but does not make the old mistaken forecast right retroactively.

What changed in the answer

Timur adds to the word "stop": "in the standard table". Below he writes the card's answer: "tryptophan".

"Two answers now?" he asks.

"Two matches. The card states one case," Dana answers.

The UGA letters did not change. What changed is our understanding of which process they belong to. Without the note we

could name possible matches. With the note we gained grounds to pick the needed one.

The answer grew precise thanks to particular information. But from this it does not follow that any claim can be saved by adding hedges. Sometimes its scope of application must change; sometimes it must be admitted false. The condition must reflect the studied system's workings and be backed by data.

In the next chapter three signs will converge on one conclusion. We will work out whether they give grounds to trust it and what is missing for checking.

PART IV · WHAT WOULD MAKE YOU CHANGE YOUR MIND?

CHAPTER 11

Three Signs Agree



Back from a forest trip, a basket stands on the kitchen table. Timur takes out a mushroom: a pale cap, gills beneath it, a ring on the stem. The stem's lower part stayed in the ground.

"The silver spoon did not darken," he says. "Worms ate it. And it smells good. Three signs, all in favour."

Dana opens the field guide. The description she reads states the colour of the spore print — the trace left by spores fallen from the cap. They have no such print.

"And what was at the stem's base?" she asks.

Timur looks at the cut.

Before reading on, decide: three reassuring signs are in, several pieces of information missing. Is that enough to send the mushroom to the frying pan?



What the signs check

Timur reasons in order. He remembers three rules: a poisonous mushroom must give itself away by darkened silver, worms will not eat it, its smell is unpleasant. He found nothing of the sort and concludes: the mushroom is edible.

But before counting passed checks, ask whether the rules themselves can be trusted.

A silver spoon is no universal indicator of mushroom poisons. Feeding traces show another creature ate the mushroom but do

not establish its safety for humans. A pleasant smell guarantees no edibility either. These signs may give no warning of a dangerous mushroom at all.

Timur truly noticed something. The mistake creeps in at the move from observations to conclusion. Between “worms ate” and “humans may eat” stands a rule that cannot be trusted. Two more unreliable signs do not fix that move.

The field guide serves another task: helping establish which mushroom we face, by characteristic traits. That may take information about the cap, gills, stem, spores and where it grew. What matters is a set of traits telling similar species apart.

We will not identify the basket’s mushroom: this chapter replaces neither a forager’s reference nor a specialist’s help. We need to understand what Timur’s confidence rests on and which grounds Dana still lacks.

Not seen does not mean absent

Dana reads about the stem’s base. In some mushrooms there is a volva there — a veil remnant that may look like a sac or a cup. But in their specimen the stem’s lower part is missing. So “no volva” cannot be written down. They never examined the needed part of the mushroom.

A similar story with spores. Gill colour and spore-print colour are different traits; one cannot simply stand in for the other. The needed observation has not yet been made.

That does not mean getting the print would necessarily pin down the species. It would add one trait to weigh with the rest. Matching on one count does not replace identification.

In the Pluto chapter we already told lack of information apart from mismatch with a criterion. Here the distinction is the same: “we did not see the stem’s base” does not mean “there was nothing characteristic at the base”. A gap in observations cannot be turned into a handy answer.

Timur suggests counting: three signs for, two checks undone. But this is no vote. What matters is whether each trait helps tell a safe option from a dangerous lookalike.

Report and decision

"We have not identified the mushroom" is a report about knowledge. "We will not eat it" is a decision under scarce information and a high cost of error. It does not require declaring the species established.

Unknown does not mean inedible

"So you think it is poisonous?" Timur asks.

"I do not know what mushroom this is. So we will not eat it."

These two sentences say different things. The first speaks of Dana's knowledge. The second, of her decision.

From our failure to identify the species it does not follow that the mushroom is poisonous or inedible. But deciding to eat it takes reliable grounds. A mistake here may cost too much, and an unknown mushroom can be declined.

So the rule "do not eat until established edible" needs no prior proof of poisonousness. It lets us act short of information without passing caution off as knowledge of the mushroom's properties.

Identifying the species and finding out whether the mushroom suits food are not one and the same question either. The second needs reliable information on edibility and conditions of use. One name on a found page is not enough.

Dana may take a definite decision while keeping the answer uncertain. There is no contradiction here.

How many independent grounds here

With the signs the first problem showed: they are unreliable. But behind the words "three confirmations" another mistake may hide — counting one and the same ground twice.

Recall the measurement from chapter two. Its result went into the report, from the report into the summary, from the summary into the presentation. All three documents hold one number. With no new measurements made, we face three records of one result.

Their agreement shows the number travelled with no divergence. But if the original measurement held an error, three identical records will not remove it.

A new measurement may add information. And then it is worth finding out whether it repeats the old error's cause. Another person with another instrument may still use the same wrong standard or the same mistaken correction. Independence cannot be established by merely recounting participants.

With the mushroom signs the case is different. Smell, feeding traces and the spoon's state are different observations. Calling them one observation just because all three signs are familiar from childhood would be wrong. The rules' shared origin does not yet prove the results dependent. The signs may well diverge.

But even independence does not turn a bad check into a good one. We need to find out both where the information comes from and how it helps answer the asked question.

So Dana needs the spore print not for a fourth tick. It may help tell species apart in identification. Its value ties to that task, not to a place in the queue of evidence.



An hour with no symptoms

“And if a person ate such mushrooms and feels fine after an hour?” Timur asks. “Is that a check already?”

There is an observation: after an hour the person notices no symptoms. But what follows from it?

Some mushroom poisonings show only after many hours. So it is with death-cap poisoning. Feeling well soon after eating is therefore compatible with dangerous poisoning and does not confirm the mushroom's safety.

Here the chapter on the number 341 comes in handy again. It passed the check while staying composite. We learned a definite property of the number but would err in ascribing it more. So here: "no symptoms yet" is a narrower claim than "no poisoning".

By symptom-onset time alone neither the poisoning's severity nor which poison caused it can be judged. Symptom-onset time alone cannot support such a conclusion.

From this reasoning no advice follows to wait a little longer. If an unknown mushroom is already eaten, seek medical help at once, without waiting for worse. Its safety cannot be tested on oneself.

What will change the answer

Timur sets the spoon aside.

"So three is too few?"

"First we must understand what they confirm," Dana answers.

Evidence count matters when it is clear how each piece was got and what each is worth. Several reliable observations may strengthen a conclusion.

Several retellings may only multiply the original error.

And several bad checks may leave unnoticed what needed noticing.

In this scene Dana is right to refuse counting the mushroom edible by signs. And she neither identified its species nor proved it poisonous. Her decision rests on the available information's limits and the possible mistake's cost.


Revising the answer about the species takes suitable identification traits and their competent assessment. Revising the eating decision takes still more: grounds to count the mushroom safe under particular conditions.

In the next chapter we return to the cart and plan a check in advance: we will state which results help pick between two explanations, and which leave the question open.

PART IV · WHAT WOULD MAKE YOU CHANGE YOUR MIND?

CHAPTER 12

An Experiment That Could Change Our Minds




Back to the cart from chapter seven. Timur and Dana are not pulling the string yet: they are drawing up an experiment plan. On the sheet are the same numbers — mass 19.0 kg , acceleration 1.5 m/s^2 , and the assumed dynamometer reading of 30 N .

Mass times acceleration is 28.5 N . The gap to the instrument's reading is one and a half newtons.

"I would still start with the dynamometer," says Timur.

"And if it is fine?" asks Dana. "Maybe we missed resistance. Or wrote down the mass wrong."

Before reading on, decide: can these three numbers pick one explanation?



Two explanations of one reading

First recall what exactly we compare. A dynamometer on a horizontal string measures its tension. Mass times acceleration gives the resultant. These forces need not match: something may resist motion. The bare difference of numbers proves no instrument fault.

Consider two possible explanations. For now they are guesses fitted to a conditional example, not measurement results. In both calculations we assume the force and acceleration readings carry no substantial systematic shift.

First: the mass is right, but there is resistance. The cart with the load weighs 19.0 kg. Suppose a 1.5 N force opposes its motion. Then from the 30 N tension, subtracting resistance leaves 28.5 N — exactly what mass times acceleration gives.

For further calculation suppose resistance stays roughly constant over the picked stretch and in the regime at hand. That assumption remains to be checked.

Second: the mass is written down wrong, and resistance may be neglected. Suppose the cart with all its fastened equipment weighs 20.0 kg. Then at 1.5 m/s^2 acceleration 30 N are needed. Perhaps some part never made it into the mass entry.

Both explanations fit the starting numbers. But these are no two independent confirmations: in each case we picked the unknown so the calculation would give the wanted result. Now check what these explanations predict beyond the starting case.

Change the acceleration

“And if we pull harder?” Timur suggests.

Plan motion at about 3.0 m/s^2 acceleration — twice as much, with the same cart and load. What tension does each explanation predict?

In the first case mass times acceleration makes 57 N. Add the assumed 1.5 N resistance: the expected tension is 58.5 N.

In the second case the mass is 20.0 kg, and resistance we neglect. The expected tension is 60 N.

Acceleration	19.0 kg and 1.5 N resistance	20.0 kg with no notable resistance
1.5 m/s^2	30 N	30 N
3.0 m/s^2	58.5 N	60 N

In the first row the predictions match. In the second they differ by 1.5 N. Now we have a chance to tell the explanations apart, if measurement precision suffices and the accepted conditions hold.

58.5

NEWTONS

19.0 kg and 1.5 N resistance.

60

NEWTONS

20.0 kg with no notable resistance.

Predictions at 3.0 m/s^2 acceleration. The difference is visible only with sufficient comparison precision.

Repeating motion at the old acceleration is useful too: the readings' spread can be estimated and the result checked for reproducibility. But if we get about 30 N again, preferring one of these two explanations over the other on such a match will not do. Repetition answers its own question.

There is a more direct move — weighing the cart with all its equipment anew, checking the weighing method and the result's entry. For choosing between 19 and 20 kg masses that may suffice. Raising acceleration after that is optional. It is needed if we also want to check how tension depends on acceleration.

Which conditions we check with the explanation

For the calculation we accepted that the cart moves horizontally, the string pulls along motion, mass stays unchanged through the experiment, and acceleration is measured against the laboratory. As in chapter seven, we count its frame of reference as inertial with precision enough for this experiment.

Force and acceleration must be compared over one and the same stretch of motion. A reading at the first jerk cannot be set against an acceleration measured later. Nor may the number 3.0 be counted exact just because we wrote it in the plan: the calculation takes the actually measured acceleration.

A separate question is whether resistance will hold. A stronger pull may change the motion's conditions. So the string's direction, the track's state and the speed range we compare results in must be watched. The same instruments and the same cart by themselves do not secure the same conditions.

We need not change exactly one quantity in every experiment. But we must understand which changes the calculation allows for and which may spoil the comparison. If resistance shifts markedly, the 58.5 N prediction no longer follows from the first explanation as stated.

Will precision suffice

One and a half newtons between predictions looks convincing on paper. In the laboratory the whole comparison's uncertainty must be estimated.

It depends on measuring tension and on the calculation's input quantities: mass, acceleration, assumed resistance. If the calculations use shared readings or corrections, that too must be allowed for. From one quivering needle no one can decide whether the predictions can be told apart.

For the example suppose such an estimate is done. A rule picked in advance: a result counts as compatible with a prediction if their difference in magnitude is under 0.5 N. Here 0.5 N is a conditional threshold for the whole comparison, grounded in the uncertainty estimate. It is not the dynamometer's division value, nor a promise that error will never exceed half a newton.

At 3.0 m/s^2 acceleration a 58.6 N result satisfies this rule for the first prediction only. A 59.9 N result — for the second only.

And 59.25 N fits neither: it differs from both predictions by 0.75 N. It cannot be entered for a favoured explanation just because it landed somewhere in the middle.

Now picture a less precise comparison, with a 1 N threshold set by the same principle. The same 59.25 N is compatible with both predictions. The number stayed put, but the chance to conclude changed.

So we provide for three cases in advance:

- **Result compatible with one prediction only.** Grounds to prefer this explanation over the second under the check's accepted conditions.
- **Result compatible with both.** This check failed to pick. Uncertainty may need shrinking, or predictions spreading wider.

- **Result compatible with neither.** Measurements, their uncertainty estimate and the explanations themselves need checking.

One of the two need not turn out right.

Even the first outcome will not prove the only possible explanation found. A constant overread of the dynamometer, say, may look in calculation like constant resistance. Telling even these options apart takes a separate check of the instrument.

Write down before you know

Timur reaches for the string, but Dana turns the sheet to him.

“First let us write down what we will count as a result.”

The sheet must keep both explanations, their assumptions, expected values, the measurement method and the comparison rule. The number of repeats and how to process their results must also be decided in advance, so the handiest reading is not picked afterwards.

A plan does not make a researcher infallible. It shows which decisions were taken before new data arrived, and which appeared after.

If a fault or an unallowed-for condition shows, the plan may change. But then the old entry must be kept, the edit explained, and the refined explanation checked apart. Otherwise an assumption invented after the experiment is easily passed off as a prediction the experiment confirmed.

Here the interlude’s dates matter again: what information was available when we drew the conclusion? For an experiment one particular question joins it: was the comparison criterion picked before the new result or fitted to it?

What could change our minds

Timur and Dana have no new reading yet. But now they hold a plan by which it can move their opinion. They know what they compare, which differences they expect, and under which conditions the comparison makes sense.

The answer “I am ready to change my mind” grew more definite. Name a result weakening the first explanation, a result weakening the second, and a result after which both must be revised. The chance that precision will fall short of a pick can be admitted in advance too.


Such an experiment promises no certain end to the dispute. It lets the next conclusion rest on new grounds and show how they differ from the old ones.

In the last chapter we move to music. There a check establishes the rule kept, while the question of whether it turned out well stays open all the same.

PART IV · WHAT WOULD MAKE YOU CHANGE YOUR MIND?

CHAPTER 13

A Dombyra Owes Fux No Exam





In music-theory class Dana lays a two-voice exercise on the stand. The lower melody is set by the teacher; the upper she composed herself. Each note of the lower voice matches one note of the upper — this way of building is called first-species counterpoint, or “note against note”.

The teacher looked the work over: the assignment’s demands are met. Dana plays the exercise on the piano.

“So it turned out well,” she says.

Timur, who plays the dombyra, listens to the end.

“All correct. But I don’t like it.”

Before reading on, decide: does his answer contradict the teacher’s words?

What can be checked against the notes

Dana studies strict-style counterpoint — joining melodic lines under fixed demands. This teaching tradition ties to Johann Joseph Fux’s treatise *Gradus ad Parnassum*, published in 1725.

Take one demand apart: in such a two-voice exercise the voices may not move in parallel perfect fifths. To understand what that means, two consecutive harmonies suffice. This fragment takes one rule apart; it is no finished exercise.

Let the lower voice go from C to D, the upper from G to B. All four pitches sit in the first octave. First C and G sound together, then D and B.

Voice	First harmony	Second harmony
Upper, original	G	B
Upper, altered	G	A
Lower, in both	C	D

In the original, C to G is a perfect fifth. A fifth spans five steps: C, D, E, F, G, counting both ends. D to B is a sixth. No two fifths in a row here.

Now replace the upper B with A, one step down. That gives C–G and D–A: two perfect fifths in a row. Both voices rise by a tone — the lower from C to D, the upper from G to A. These are the parallel fifths banned in the picked assignment.

We can point at the exact breach: it arises in the move from the first harmony to the second. Reading the notes right and applying the rule suffices for that. A listener’s personal impression will not change the intervals.

But we checked only one move under one demand. From this it does not follow that the original meets all strict-style conditions. Nor does it follow that changing the note leaves no trace elsewhere: in the whole exercise its ties to neighbouring pitches need checking.

What “correct” means

For our scene accept that the teacher checked the whole assignment and truly found no breaches. Dana may rightly say: “The conditions are met.” That is an important result — she coped with what she learned.

But the phrase “it turned out well” may mean more. Is the exercise interesting to hear? Does the melody develop convinc-

ingly? Would you want to hear it again? Checking the set list of demands does not automatically answer these questions.

Counterpoint rules help learn to hear and build voice combinations. Following them is no empty formality. But between mastering a device and applying the device well lies the musician's work.

So it was with the number 341. The check established a definite property, and Timur wanted a wider conclusion out of it. Now Dana makes a similar move: from meeting the assignment's conditions she concludes the musical result is bound to be good.

Timur rejects that conclusion. He never claimed the exercise was written wrong. Nor has he said yet that it is bad: only that he does not like it.

Music can be argued with reasons

"What exactly don't you like?" Dana asks.

Timur asks to play once more and stops her at one place.

"I keep waiting for the melody to move on. And it comes back here again."

Now the conversation has a ground. They can look at which turn repeats, hear how it sounds in different places, and discuss what Timur expected.

Dana may answer that the return is designed: it binds the phrase's parts. The teacher may suggest changing one turn and comparing both versions. Then it grows clearer what the repeat gives — wholeness or a sense of sameness.

The repeat's bare fact can be checked against the notes. Its expressive meaning takes other reasoning: the place in the phrase, the performance and the design must be allowed for.

From "the turn is repeated" it does not follow that "the music is dull".

Some musical demands can be stated exactly and put into a check. Others allow several convincing solutions. That does not

make discussion pointless: the speakers may present examples, compare performances and revise estimates.

A teacher's experience helps notice what a pupil still misses. But citing the teacher does not replace his reasons. Hearing what exactly he counts as good, and why, beats getting one more unexplained "fine".

Three different judgments

The assignment conditions are met. The musical solution works. I like it. These judgments may coincide, but each needs its own grounds.

Taste can be explained too

In Dana and Timur's dispute three questions tangled: whether the work meets the assignment, how well it is made, and whether a particular listener likes it. They connect, but an answer to one does not replace answers to the rest.

A personal impression need not be unaccountable. Timur can explain what he likes in music, which motion he awaits, and what tires him. These words will help Dana understand his reaction, though they will not oblige her to feel the same.

He may also simply say "don't like it", offering no analysis. That reports his present impression. Moving from it to "the exercise failed" takes more grounds and an understanding of which signs mark success.

Knowing the design or another performance may change the impression. It may not. A person can appreciate mastery and still not love a work. There is no contradiction here for one more check to remove.

What the dombyra has to do with it

After the exercise Timur takes the dombyra and plays a familiar kui. Dana listens, no longer checking against the demand list.

The dombyra is a two-stringed instrument, but two strings alone do not turn any piece into a "note against note" exercise. A kui owes no answer to the assignment Dana received.

Its rhythm, the melody's motion, the sound combinations, the playing manner can be studied. Performance decisions can be compared and discussed with a musician who knows this tradition. Such a conversation needs grounds belonging to the music itself and the asked question.

Judging a kui only by whether it suits a strict-style teaching assignment teaches its fit for that assignment. That answer does not replace an artistic assessment of the kui.

The point is not that Fux's rules are banned from the dombyra. A special exercise or arrangement can be composed for it, with counterpoint demands picked in advance. Then checking them is fitting. What matters is the work's aim and conditions, not the instrument's name.

Before staging such an exam, understand which skill is to be checked. Otherwise the mark will concern an assignment the musician never received.



Which conclusion we may draw

Dana returns to her score page.

"Good as a finished assignment," she refines. "And the music can still be worked on."

"Now it's clear," says Timur.

He needed no breach hunt to keep his impression. She needed no belittling of the done work to admit the check's limits.

Grounds are needed not only where an answer comes by computation. They live in a musician's judgment too, when he points


at a particular phrase and explains what he hears in it. Such reasons can be discussed, even if they never bring all listeners to one estimate.

What matters is understanding which conclusion they support. Meeting a rule, a good solution to a musical task, and listening pleasure may coincide. But none of them guarantees the other two.


In the epilogue we return to the claim you wrote down in the prologue. Now more questions can be put to it — and what your answer stands on can be said more precisely.

INTERLUDE

The Word “Knowledge”



We went through thirteen chapters but never defined knowledge. We spoke of measurements, sources, rules, proofs — of why we count our answers right. Now, before returning to your prologue phrase, pin down what ties these conversations together.



Confidence and its grounds

“I know the house was built in 1968” can be said in full confidence. But the person across the table may ask: where did that date come from?

If the answer is “I remember well”, we can pin down what exactly the person remembers. If he shows a registration certificate, we can read the entry and find out which year it states: of construction, of commissioning, or of something else. The conversation continues thanks to information that helps check the claim.

Confidence and groundedness are different things. One may never doubt and be wrong. One may speak cautiously of a conclusion with serious grounds behind it.

**Tone of voice alone decides nothing
about who is right.**

In a conversation about knowledge two questions help tell apart: how sure a person is, and what his conviction stands on. In this book we cared above all about the second.

Three conditions and one difficulty

One known philosophical approach describes knowledge as justified true belief. A person accepts some claim; the claim is true; the person holds sufficient grounds to accept it. These conditions help tell belief, truth and justification apart.


Already in Plato's *Theaetetus* a close thought is discussed: true opinion must come with an account. In the dialogue, though, that answer meets criticism. Ascribing to Plato a ready definition, simply handed down through the ages after, would be inexact.

There is another difficulty, pointed out in 1963 by the philosopher Edmund Gettier. Picture two job applicants — Smith and Jones. The manager assured Smith that Jones would be picked. Smith also counted the coins in Jones's pocket: ten of them. He concludes that the man who gets the job holds ten coins in his pocket.

But the job goes to Smith himself. And he, unbeknownst to himself, holds ten coins too. His conclusion proved true, though it rested on information about another person.

Smith held grounds for his belief, and it turned out right by coincidence. This example shows the difficulty: no lucky outcome replaces the belief's tie to truth. Exactly this reasoning targets the three conditions' sufficiency.

For this book's tasks the whole dispute over knowledge's definition needs no settling. But it cannot be bypassed by simply renaming any explained claim as knowledge.



Confidence, justification, truth

Confidence describes a person's attitude to a belief. Justification is the reasons for holding it. Truth is whether the claim itself is right. One question does not replace the others.

Checking and its limits

We never dropped the question “is it right?”. For its sake we checked measurements, read sources, computed and planned experiments. Factoring 341 established the number is composite. New information on the genetic code showed the old generalisation's boundary.

The catch is that not every justification guarantees. A well-backed supposition may prove too wide. Observations in studied systems do not yet establish that UGA in every possible case means synthesis halt.

A date helps understand which grounds were available to a researcher. It does not make a mistaken claim true for its time. But the prospect of a future review does not mean nothing can be found out now either.

So we need both questions: what lets a claim be counted right, and how reliable is this way of checking? Describing a ground in detail does not yet show it can be trusted.

What the book’s questions give

We learned to take a claim apart so another person could grasp and check the reasoning. For that we pinned down what exactly is said, which information is given, which conditions are accepted, how the conclusion was reached, and what could change it. When it mattered, we added questions about time.

These questions make no definition of knowledge. They help discuss grounds: find the source document, repeat the calculation, notice a missed condition, or suggest a check whose results could change our minds.

A filled card by itself certifies nothing. It may hold wrong information, an unreliable rule, or a mistaken conclusion. The record’s use is that it is now clearer where to hunt the problem.

No promise comes with it that every mistake will show at once and touch only one line. Sometimes a refinement suffices, sometimes the conclusion must go, and sometimes several tied assumptions need review. Readiness to revise helps in this work but does not replace it.

We aim at a right answer and a clear path to it. That is the task of this book about grounds of knowledge. For it, no need to declare every discussed conviction already accomplished knowledge.

More than can be written down

The word “know” has other uses too. Knowing how to ride a bicycle or tune a dombyra means holding a skill. A description of moves may help learning but by itself does not replace practical knack. And not every master can explain in detail everything he does right.

Knowing a person or knowing a city well is another case. From acquaintance and experience many claims can be drawn, but a short record about a person does not replace knowing him.

Perception and memory joined all our stories too. Someone saw an instrument’s reading, heard a melody, recalled a road.

The book turned to these abilities but attempted no full theory of their work and reliability.

Focusing on claims, we picked one side of the conversation about knowledge. The book's questions may help in discussing skills or personal experience too, but they neither exhaust them nor turn all knowledge into a set of records.


Before the epilogue

Return in mind to the words "I know it". Behind them may stand more than can be explained at once. And yet one can start with a concrete question: what gives me grounds to think so?

With it we begin taking your prologue phrase apart. Not to replace confidence with doubt, but to understand better what in this answer can be trusted.

EPILOGUE

An Answer Together with Its Grounds



In the prologue you wrote down one phrase — something ordinary you were sure of. Take it out and read it once more.

Your opinion may not have changed while reading. Or you may now want to add a hedge, or check what once seemed obvious. First three prologue examples, then yours.

What changed in familiar phrases

“The house was built in 1968.” Where does that date come from? It may stand in a registration certificate, a listing, or a plaque by the entrance. But first we must read what exactly is named: the year construction finished, of commissioning, or of the document’s issue.

Then we can find out how the entry reached that source. Several identical dates do not yet mean several independent confirmations: the information may have been copied from one document. It also matters whether the entry concerns the wanted building.

After checking, the source phrase may stay. But until the document is checked, it is more exact to say: “This source states 1968; what the date means exactly and where it comes from I have yet to find out.” We already know where to continue the search.

“In the mornings another road gets you to work faster.” Which trips do we compare? How many were there, on which days and at what time did we leave? Did we measure door to door, or only time in motion?

A few lucky trips may hint the new route is worth trying. For a surer conclusion we must compare it against the old one under matching conditions, allow for time spread, and decide in advance how to judge results. One smaller average is not enough to promise a faster way every morning.

The conclusion may have to narrow: the other route is faster on certain days or with an early start. Or data may so far fall short of a pick. The needed hedges will come from observations, not from a wish to save the phrase.

Saying “baram” is wrong. In chapter six we told usage, norm, reading and artistic choice apart. A shortening’s bare existence establishes no mistake. But neither does a form met in a text become fitting for every situation thereby.

For reading the considered verb forms the reference gives a concrete direction: keep the full or shortened variant the text uses. That is no general rule to pronounce everything literally letter by letter. When we write our own text, that task’s conditions must be taking apart.

So the source blanket condemnation of the shortened form fails checking. We replace it with more exact judgments. That is a substantive change of answer, not a way to declare the old claim right.

In all three cases one question helps find out: what do the available grounds let us say now? Sometimes the answer stays, sometimes it changes. Sometimes a check’s honest result is admitting the question is still open.

Now your phrase

Write beside it answers to the familiar questions. No need to fill gaps with guesses. Where you do not know something, mark what exactly remains to find out.

What exactly do I claim? Pin down the phrase’s meaning enough to see what to check. If a road is “faster”, compared to which, at what time, and for which way of getting about?

Which information backs it? Name a document, observations, a calculation, or a person’s evidence. Where possible, state where they can be seen.

Which conditions do I accept? Say, that the document concerns the same house, the instrument readings suit comparison, and the rule applies to our case. Keep checked things apart from what is so far only assumed.

How did I reach the conclusion? Show the move from information to answer. Read the date yourself? Compared trip times? Applied a rule? Relied on another’s estimate? One reasoning may join several such steps.

What could change my answer? Name a concrete piece of information or a check's result. Thinking this over in advance helps, while no new data exists yet. If the criterion later has to change, explain why.

Does time matter here? If so, pin down which time you speak of, which information you allow in, and when the conclusion was made. These dates may not match.

After such taking apart the phrase may stay as it was. But now the person across the table finds it easier to grasp why you accept it, and you — to see which questions to it stay unanswered.


What stands behind a short answer

Around the prologue table familiar lines sounded: "I saw it myself", "I was told", "It says so", "I worked it out". Now each of them can be continued.

After "I saw" — pin down what could be observed and under which conditions. After "I was told" — ask where the teller gets it and on which point he can be trusted. After "it says so" — find out the entry's meaning and source. After "I worked it out" — show the source data, assumptions and the calculation's course.

No single line completes a check by itself. But each may start one. We turned to different grounds because the questions differed: to set a part's size, to read a line, to prove a number's property, to discuss a musical phrase.

Not every dispute ends in one answer. Still, we can pin down more exactly where the speakers part: in information, in conditions, in reasoning, or in estimate criteria. That alone lets the conversation continue on substance.



What such taking apart gives

In the knowledge interlude we told confidence, groundedness and truth apart.

Written-down grounds guarantee no right answer.

They too need checking.

A card helps make reasoning open to discussion. Another person may find a gap, suggest new data, or show a mistake. He may also agree, seeing sufficient grounds.

Not everything we know and can do finds words at once. An unexplained impression may start a taking apart; a skill may help notice an important detail. But when we offer others a conclusion, we should explain where possible what it stands on and why it deserves trust.

Returning to the source

The name Arxo recalls the Greek *archē* — beginning, source, first principle. For this book the tied question of grounds of knowledge matters above all.

With time an answer grows familiar. We remember a date, a route, a formula or a rule, while the circumstances in which we learned them fade. It seems we always knew it.

Returning to the source means seeing again behind a ready phrase an observation, evidence, an accepted principle, someone's work. One memory of where we first heard the answer is not enough. We must understand why we trust it now.

Arxo helps trace how a conclusion depends on stated rules and information. But checking the information itself, picking fitting conditions and judging whether the grounds suffice still falls to us. This work's meaning is keeping the tie between what we claim and what lets us claim it.

LAST PAGE

Your prologue phrase still lies before you. Now beside it may stand a document reference, a calculation, an important hedge, or a question for the next check.

You may leave it unchanged. You may rewrite it. What matters is that grounds show behind the answer — and that we can tell whether they suffice.

What do we know, and what does it stand on?

Grounds. How We Know and Why We Can Be Wrong

First English edition · Edition 2026.10

© 2026 Rifat Jumagulov

Published by Arxo

This book is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0), <https://creativecommons.org/licenses/by/4.0/>. You may copy, share, translate and adapt it, including for commercial purposes, provided you give appropriate credit, link to the licence and indicate whether changes were made.

The licence covers the text of this edition and, to the extent they are protected, its illustrations. It grants no rights to the Arxo name or logo. Quoted passages and other third-party material remain under their own terms; their sources are named where they appear.

Translations and adaptations made by others are not official editions. Please mark them as unofficial and link to the original.

The illustrations were generated with Google's Nano Banana Pro image model and selected and edited for this edition.

Set in Charis SIL, IBM Plex Sans and IBM Plex Mono, all under the SIL Open Font License 1.1.

The current edition, free in PDF, EPUB and on the web: books.arxo.io