

Chapter 1 · The Seen and the Unseen

This book builds models of economic life by hand, with a small number of agents whose behaviour you can recognise. Before we build the first one, we need to settle three things: what a model is, how a model should be judged, which depends on what the model is for, and why this book answers both questions differently from standard economics.

The chapter is light on philosophy and heavy on examples. It ends by showing why agent-based models are the natural tool for this book: they solve a problem, the problem of historical specificity, that neither standard economics nor its traditional critics could solve. Several of the points the chapter raises are developed properly later, when we have models in front of us to build and compare.

1. Two kinds of thing

You are standing at a crossing. A car is coming, and it begins to slow. You step off the kerb.

What did you see? The car, its position, its falling speed. What did you act on? Something you could not see at all: the driver's intention to stop. You staked your safety on a judgement about what was going on inside another person.

Or take an exam. The exam produces a number, but the number is not what your teacher cares about. She cares about what you have learned, and the grade is only a trace of it. Two students can earn the same grade, one because she understood the material and the other because he memorised it the night before.

Every science works with these two kinds of thing: what can be seen, and the hidden things that produce it.

- In physics, there are tracks in a detector, and the particles said to have made them.
- In medicine, there are symptoms, and the disease behind them.
- In economics, there are prices, purchases and choices, and whatever goes on inside people to produce them.

The second kind of thing is not exotic. We deal in it constantly and mostly without difficulty, as the person at the crossing does every day.

2. What a model is

A model is a simplified version of reality, made for a purpose. Everything in this book depends on taking each part of that sentence seriously.

2.1 Models simplify reality

Think of a map and the ground it represents. The map leaves out almost everything: the slope of each field, the trees, the stones underfoot, the colour of the houses.

That is exactly why it is useful. With the map you can find a route, estimate a distance, and judge how long the journey will take and how hard it will be. With the ground alone, in all its detail, you could do none of this. What a model leaves out is as deliberate as what it keeps.

2.2 The simplification serves a purpose

Suppose you need the area of a plot of land with an irregular boundary. There is no formula for such a shape. So you divide it, approximately, into triangles and rectangles, whose areas you know how to calculate, and add them up.

Now suppose you want to build on the same plot. You need a different simplification: straight lines along which walls can stand, and a drawing an architect can work from. The same reality yields different models for different purposes.

Two purposes run through this book. One is **calculation**, which is why so many models take mathematical form. The other is **understanding**: grasping something too complex to take in all at once. The models of the standard economics textbook were simplified so that they could be calculated with pencil and paper, and their defenders say they serve prediction; section 4 asks whether they do. The models in this book are simplified so that they can be understood. Later chapters show what that difference does in practice.

2.3 Models are metaphors

“Asad is a lion.” Taken literally, the sentence is false. Yet it says something true: that Asad shares a particular quality with a lion, his courage. A metaphor compares two quite different things in a few respects, usually without saying which respects it means. “An uphill struggle” works the same way.

Every model is a metaphor of this kind. Reality is so complex that every moment in it is unique; it has never happened before and will never happen again. If we treated each moment as unique, nothing we had learned before would apply to it.

To learn from experience at all, we must ignore what is unique about this moment and attend to what it shares with others we have known. That is what a model does. And it carries a risk: in setting aside the details, we may throw out the baby with the bathwater, and lose the one thing that mattered.

2.4 Always wrong, always necessary

The earth is roughly a sphere, but only roughly. That is a first simplification. To draw a map, we then flatten the sphere onto a page, which distorts it a second time. Every map we use rests on both.

So no model captures reality’s full complexity, and every model is, taken literally, false. Why use things we know to be false? Because there is no alternative. Our minds cannot hold reality whole; to understand it, or to act on it, we must simplify.

It follows that the right question about a model is never “is it true?” It is “has it kept what matters for its purpose?” Section 3 turns that question into a criterion.

2.5 The dangers of models

A model highlights some features of reality and makes the rest invisible. Over time we come to replace the reality with the model in our minds, and forget what was left out. Nassim Taleb calls this mistaking the map for the territory. It brings a false sense of security: the feeling that every risk has been calculated.

Long-Term Capital Management shows how costly that can be. The fund was founded in 1994, and two economists on its board, Myron Scholes and Robert Merton, shared the 1997 Nobel Prize for their method of valuing derivatives. Its models treated certain risks as very small. In 1998, after crises in Asia and Russia, those risks turned out to be very large. The fund lost about \$4.6 billion in less than four months and needed a rescue arranged by the Federal Reserve.

2.6 What a model contains

Like the reality it simplifies, a model has two kinds of part. Some parts can be checked against what we observe: prices, quantities, who traded with whom. Other parts stand for what we cannot see: the mechanisms, and in economics the motives, that produce what we observe.

A model with no hidden part is not a model at all, only a summary of data. A model with no observable consequences can never be brought into contact with the world. Every model worth having has both.

Exercise, before reading on. You are on the metro. A young man gets up and offers his seat to an older passenger. Write down three genuinely different accounts of why he did it, not three versions of the same account. Each should be an account that, if true, would have produced exactly what you saw. Keep what you write; the rest of the chapter uses it.

3. Two ways of judging a model

Here is a well-known kind of experiment. Each student sits alone in a room with a short quiz. The answers are printed on the back of the sheet. The instructions are to answer the questions first, then turn the sheet over, mark your own work, and hand in your score. Nobody is watching, and nothing prevents a student from turning the sheet over first.

Suppose that, across the class, one answer in five that was marked correct had in fact been copied from the back. Now build a model of the class. Here are two:

- **Model A.** One student in five is a cheater and copies every answer. The other four are completely honest.
- **Model B.** Every student cheats a little, copying about one answer in five and doing the rest honestly.

If all we observe is the overall rate, both models reproduce it exactly. Yet they describe two very different classes. In one, a few people have no scruples. In the other, everyone has a few.

The hidden part can differ in another way too. Take two students who did not cheat at all. One was honest because of the integrity in her heart. The other was honest because he suspected the room might be watched. Their score sheets are identical; their reasons are not.

Look back at your metro exercise. You have three accounts, and every one of them fits what you saw. Nothing you observed tells you which one is right.

So given a model, how do we judge it? The first answer is that we cannot, until we know what the model is for. A model is a simplification made for a purpose, and it can be judged only against that purpose. A map for walkers and a map for pilots are judged differently. Two purposes dominate discussion of economic models, and each brings its own criterion.

The first purpose is prediction: does the model reproduce what we observe, or will observe? Work out what the model implies we should see and compare it with what we do see. A close match counts in the model's favour. Call this the **nominalist** criterion. It is the purpose usually claimed for conventional economic models.

The second purpose is understanding: has the simplification kept the mechanisms that produce what we observe? On this view, a model is good if, having stripped away the complications of reality, it still keeps the motives and mechanisms that do the work, so that it helps us understand the phenomenon. Call this the **realist** criterion. It is the purpose of the models in this book.

These are not two kinds of model; any model can be judged for either purpose. And neither criterion makes a model true. Every model is a simplification, and taken literally false (section 2.4). Models A and B pass the nominalist test equally, yet they give very different understandings of the class. Fit alone cannot choose between them.

Behind the two criteria: two views of science

The two criteria come from two views of what science is.

One view, logical positivism, was highly influential in the early twentieth century. It held that science deals only with what can be observed. A scientific law is a pattern among observations, and it is tested by whether it predicts observations beyond the data in which the pattern was first found. Religion, on this view, rests on things unseen; science rests on things seen. But science also speaks of things nobody sees: atoms, forces, gravity. The positivists' answer was that any such unseen entity could be replaced by its observable consequences. "Gravity" is shorthand for the way bodies fall, and says nothing more. Behind this lay a hope: that knowledge confined to what anyone can observe would be certain, and everyone could agree on it.

The other view, realism, holds that a good model captures the real entities and causes that produce what we observe, even though those entities and causes cannot themselves be seen. The hidden part of a model is a claim about what is really there.

The nominalist criterion follows from the first view, and the realist criterion from the second. This book takes the realist side. Section 5 explains why confining attention to

observables does not make a model true, and the box in section 7 tells how the first view came to dominate.

4. The appeal of the nominalist criterion, and its limits

The positivist move is familiar from everyday life: replace what cannot be seen by its visible trace.

A scholarship committee has twenty applicants and three awards. It can rank the applicants by grade-point average. Everyone can check the ranking, and nobody can dispute it. Or it can try to judge which applicants are most capable. Then the members will argue, each convinced by different evidence, and there is no calculation that settles the argument.

Many institutions choose the first way, and for understandable reasons. Universities often promote professors by counting papers and citations rather than judging the quality of their work. A count commands agreement. A judgement of quality does not.

The nominalist criterion has the same strengths, and they should be taken seriously:

- **It can be checked.** Whether a model's output matches the data can be settled by calculation, without appeal to intuition, authority or taste.
- **It claims no more than it can test.** The realist puts unseen machinery into his model and asks us to take it on trust. The nominalist declines to, and refusing to claim what you cannot support is a virtue.
- **Prediction would be a demanding test.** Foreseeing what will happen is not easy.

Behind these strengths lies a single idea. Nominalists hold that by restricting attention to what can be observed and calculated, we reach knowledge on which everyone can agree. Whether that restriction also brings certainty is a question this book can leave aside.

The cost is visible in the scholarship example. The ranking by grades is beyond dispute, but it does not tell the committee who is most capable. It tells them who scored highest. Those can be very different things.

The criterion also has two limits that are easy to miss.

A match is not a confirmation. Checking whether a model fits is easy, but a fit does not show that the model is true, or even that it is better than its rivals. Models A and B fit equally well, and for any set of data many different models can be built that fit it.

The predictive claim is rarely tested. Conventional economic models are defended as tools of prediction, yet the claim that they predict well is seldom put to a serious test. When events did test it, as in the financial crisis of 2007–08, which standard models did not foresee, the claim failed.

5. No model can be checked for truth

It is often said that nominal models can be checked, because their output can be placed beside observations, while realist models cannot, because the hidden part of reality cannot

be seen. This contrast is mistaken. No model, of either kind, can be placed beside reality and checked for truth. Every model is a simplification, and taken literally false.

The history of science shows both halves of the problem:

- **The ether.** Physicists once held that light, like sound, needed something to travel through, and concluded that space must be filled with an invisible medium, the ether. For a long time the theory fitted what was observed and was treated as established. It was eventually abandoned. Fitting the observations had not made it true.
- **Dark matter.** Galaxies and clusters of galaxies behave as though they contain far more mass than we can see. Physicists propose that the missing mass is matter we cannot observe directly. It may be there. It may, like the ether, turn out not to be.

The difficulty reaches even into ourselves. If I say a man's nose is three centimetres long, you can measure it. If I say there is love in his heart, there is no tape to measure it with. Even about our own hearts we are usually right, but not always; people deceive themselves.

So claims about hidden mechanisms cannot be settled by checking; they need judgement, and judgements differ. Many philosophers of science reject realism for this reason. But the rejection rests on the mistaken contrast: fitting the observations does not establish a model either, as the ether shows.

We agree that claims about what is hidden can never be proven. But nothing else about a model can be proven either. The useful question about any model is not whether it is true but whether it serves its purpose, and section 9 sets out how to ask that question of a model built for understanding. Meanwhile we learn to live with uncertainty, and we build agreement in a way that does not depend on proof.

6. Living with uncertainty

Uncertainty is the condition of all human knowledge, not a special weakness of realism. The realist simply admits it openly.

Go back to the crossing. You read the driver's intention from the way he slowed, and you crossed. You had no proof. If he were your enemy, he could have slowed to deceive you and then accelerated. Yet you were right to cross, and you do the same every day. Ordinary life runs on knowledge that is never certain and is usually good enough.

Experience does not interpret itself

There is an old story about Robert the Bruce, who fought for Scotland's freedom against a far more powerful English king. He was defeated six times. Hiding in a cave after the sixth defeat, ready to give up and leave the country, he watched a spider trying to spin a thread from one side of the cave to the other. It failed six times and succeeded on the seventh. He went out, raised one more army, and won.

What should six defeats teach? The story says persistence. But a sensible adviser might have said the opposite: six times you have led your supporters to their deaths, and a

seventh attempt will do the same. The experience alone does not decide between these lessons. What we learn depends on the model we bring to it, which is the point of section 2.3. For a Muslim, the Qur'an is the first source of the models through which experience is read.

Degrees of certainty

We hold some beliefs at ninety-nine percent, some at ninety, and some barely at all. Real life requires keeping track of these degrees. A great deal of Western philosophy has instead treated knowledge as a binary: a claim is either true or false, known or not known, with no room for “perhaps”. Living with uncertainty means keeping the “perhaps”.

When to act and when to wait

If you must choose a school for your child this September, you cannot say “I don’t know.” You learn what you can, choose on the best evidence available, and stay ready to change course if better evidence arrives. The jurists of our tradition worked the same way: on the basis of the hadith that have reached me, this is my ruling; if you find something better, follow it.

If instead the question is what some distant government will do next, and nothing you do depends on the answer, you can suspend judgement. Nothing obliges you to have an opinion.

The scholar’s “I don’t know”

People find uncertainty uncomfortable, so they jump to conclusions. Every drawing-room conversation is full of confident opinions about matters nobody present knows much about. The Companions showed a different discipline. When the Prophet ﷺ asked them what day it was, a question whose answer they knew, they replied: Allah and His Messenger know best.

The scholar’s journey begins in the same place. First learn what others have said, and why: scholar X holds this, for these reasons; scholar Y holds that, for those. Understand all of them before judging any. A claim you cannot give up is a claim you cannot test; if you already know, there is nothing left for evidence to do. The willingness to say “I do not yet know” is the condition of learning anything.

And certainty of faith?

Someone will ask: do we not know with certainty that God exists? We do. But that certainty, *yaqin*, is not the conclusion of an argument. Al-Ghazali describes in *Deliverance from Error* a doubt that lasted about two months, in which he could not trust even the first principles of reason. He could not argue his way out, because any argument had to be built from the very principles in doubt. What restored him, he says, was a light that God casts into the heart.

Yaqin is a state of the heart, not a deduction, so it does not compete with the uncertainty of our models. Nothing in this chapter gives any economic model scriptural warrant. The models in this book are conjectures; some of them will be wrong.

7. The interior: economics' hidden layer

In economics, the hidden part of the picture is the inside of a person: wanting, knowing, hesitating, being pulled two ways, choosing.

Here economics differs from physics in a way that is easy to miss. A physicist has never been an electron; everything he believes about electrons is inferred from their traces. But each of us has direct access to exactly one human interior. You know what you wanted when you did that. You know the reluctance you felt and overrode. You know when the reason you gave was not the real reason.

From that one case we form conjectures about everyone else. The conjectures are never certain, and the whole of ordinary life is built on them. A teacher looks at a room and knows it has stopped following. The person at the crossing reads the driver's intention and steps off the kerb.

Why, then, does standard economics leave the interior out? Not out of carelessness. What I know of my own interior, I know about as surely as I know anything. But nobody else can check it. Knowledge that cannot be checked by others cannot command agreement, and the nominalist rule admits only what can. The exclusion is principled, given that rule.

Economics made this positivist move explicitly. Preferences cannot be observed, so Paul Samuelson's theory of revealed preference set out to build the theory of the consumer on observed choices alone, the same substitution as the scholarship committee's grades for capability. Whether the substitution succeeds is doubtful: choices follow the theory's rules only because there are preferences behind them. The further reading for this chapter takes up the question.

Our tradition has never treated the unseen as beyond knowledge. The Qur'an opens its description of the God-conscious with "those who believe in the unseen" (2:3). It also tells us where to look for the signs of what we cannot see: "in the horizons and within themselves" (41:53). The second of those sources is exactly the one the nominalist rule sets aside. And the teaching that actions are judged by their intentions is the same distinction applied to deeds: two people perform the same act, and what separates them is the hidden *niyyah*.

Why the unseen was set aside. Europe in the sixteenth and seventeenth centuries tore itself apart in wars along the lines of religious difference. To many thinkers the diagnosis seemed clear: people were killing each other over things nobody could see or check, matters of doctrine and the unseen, on which no agreement was possible. The remedy was to build knowledge on a footing everyone could share: restrict it to what anyone can observe and verify, and treat the rest as private opinion.

The Islamic experience suggests a different lesson. Difference of opinion need not lead to violence; it can be held with humility, and our tradition calls it a mercy. The two World Wars of the twentieth century, which were not wars of religion, showed that the deeper cause of such conflicts was the pursuit of power. But by then the

conception of knowledge as only what can be observed had been building for centuries. The diagnosis had failed. The remedy stayed.

8. Between universal laws and unique events

Agent-based models are not simply a convenient tool for realists. They solve a problem that neither standard economics nor its traditional critics could solve.

An experiment and its hidden motives

In an oral double auction, buyers and sellers gather in a room. Each seller has a good and a cost below which he will not sell; each buyer has a limit above which she will not pay. Sellers call out the prices they will accept, buyers call out the prices they will pay, and a trade happens whenever someone accepts an offer. When the experiment is well designed, prices usually settle close to the equilibrium of the textbook supply-and-demand diagram.

The experiment records what happened: which offers were made, which were accepted, at what prices. It does not record why. The reasons were inside the heads of the subjects. You could ask them afterwards, but their answers would be one more piece of evidence, not a window.

Now consider the two traditional ways of explaining the result.

Standard economics starts from universal laws: supply and demand determine an equilibrium price, in every market, at every time and place. That law covers this session, but it says nothing about these particular people and why they did what they did.

The historical and qualitative tradition starts from the other end. It studies a specific episode in depth: this session, these subjects, what each of them said and did. It can explain this session very well. But it has no way to carry the lessons to the next session, or to a real market. The economist Geoffrey Hodgson called this the problem of historical specificity, in *How Economics Forgot History*. One approach cannot see what is specific; the other cannot get beyond it.

What an agent-based model adds

An agent-based model cannot leave the motives unstated. To build one of this session, we must write down what each agent does, as a rule. For example: a seller accepts any offer above his cost plus a margin he has in mind, and lowers his asking price a little when no one has bitten for a while.

If agents following such rules reproduce what happened in the room, those rules are a hypothesis about what moved the real subjects. The model has stripped away the particulars of one afternoon and kept what we believe mattered. That is how an agent-based model closes the gap between the universal law and the unique event: it is a fully specified hypothesis about the factors at work, which can be carried to other cases and tested there.

The same logic applies to much larger episodes, such as why the global financial crisis happened, or what lay behind the rapid growth of East Asia and China. We will start with problems small enough to run by hand.

Notice also the difference between an experiment and an agent-based model. The experiment shows what people do. The model offers a conjecture about why. Experiments are costly, hard to control, and hard to repeat with the same people. A model can be run as many times as we like, with each rule changed one at a time.

Three kinds of agent-based model

Agent-based modelling did not begin in economics. It began in the study of complex systems. Watch a flock of birds turn: the whole flock banks together and holds its shape, as if it had a leader and a plan. It has neither. Each bird follows a few simple rules about the birds nearest to it: stay close, but not too close, and steer roughly as they steer. Put those rules into a few hundred artificial birds and the flock appears. Nobody designed the pattern; it emerges from local behaviour. Here the simplicity of the agents is essential. If each bird were complicated, the order might be hiding inside the birds, and the demonstration would prove nothing.

When economists took up agent-based modelling, they built something quite different. Their models have thousands or millions of agents, along with banks, firms, households and regulators, and they run on large computers. The agents may maximise profits or utility, as in standard theory. These models are usually judged by whether their output matches statistical patterns in real data, which is the nominalist criterion of section 3.

The models in this book are a third kind, built for understanding:

- **The agents behave in ways you recognise.** They want several things at once, the things may conflict, they know only what someone in their position would know, and they choose.
- **The institutions are written down.** Who may deal with whom, who names the price, whether a deal can be renegotiated, when an agreement becomes final.
- **The models are small.** Ten or twelve agents, run by hand on paper. When something interesting happens in a small model, we can say what produced it.

Approach	What is inside an agent	What surrounds the agents	Purpose, and how it is judged
Economic theory	One motive, complete information, perfect calculation	Prices, and little else	Prediction, claimed but seldom tested; agents optimise and the outcome is an equilibrium
Complexity models	A few simple rules; no motives	Space to move in	Showing how a pattern can emerge; output reproduces it

Approach	What is inside an agent	What surrounds the agents	Purpose, and how it is judged
Agent-based models in economics	Rules chosen by the modeller, often optimising	Banks, firms, households, regulators	Prediction; output matches statistical patterns in data
This book	Several motives that may conflict; partial knowledge; a real choice	Institutions written as explicit rules	Understanding; sufficiency, reproduction and plausibility (section 9)

9. Judging a model built for understanding

The models in this book are built for understanding, so they are judged by three questions. The nominalist expects agreement to follow from the facts; the realist has to build it, and agreement grows as a model answers these questions well.

Sufficiency

A model strips away most of reality. The first question is whether what remains is still enough: is this simple, crude model sufficient to give us understanding of the far more complex reality? This is the risk named in section 2.3, of throwing out the baby with the bathwater. A model of the metro that leaves out the other passengers can never tell us whether the young man stood up out of kindness or because he was being watched.

Reproduction

Does the model reproduce the phenomena we want to explain? A model that does not produce the phenomenon has nothing to say about it. But reproduction is necessary, not sufficient. Models A and B in section 3 both reproduce the overall rate perfectly, and they offer very different understandings of the class.

Plausibility

Do the mechanics inside the model look like a plausible explanation? The agents must behave in ways we recognise. The rule is: test it on yourself first, then against collective experience, meaning what other people, and the accumulated findings of research, tell us. Would you behave like this in that position? Do you know people who do? A model nobody recognises has probably got something wrong. So has a model everyone recognises that explains nothing.

When rival accounts are equally plausible, look for evidence about the mechanism itself, not only about the outcome. Rival motives often predict different behaviour in situations we have not yet looked at:

- **The metro seat.** Late at night, an older passenger boards an almost empty carriage. The account based on kindness predicts the young man still stands up. The account based on being watched predicts he stays seated.

- **The honest students.** Run the quiz again in a room where students are certain no one can see. Integrity predicts the same honesty. Fear of being watched predicts more cheating.
- **Models A and B.** Look at the individual score sheets rather than the overall rate. Model A predicts a few students with perfect scores and the rest honest; Model B predicts small inflation on nearly every sheet.

Accumulation: building agreement across cases

A hypothesis earns confidence when the same mechanisms reproduce more than one result. Suppose a set of rules for buyers and sellers reproduces the oral double auction. Now change the institution: sellers must post a fixed price and cannot bargain, as in most shops. Experiments with posted prices give quite different results. If the same rules, placed in the new institution, also reproduce those results, confidence in them grows. If they do not, we go back to the rules.

Because every rule in an agent-based model is written down, anyone can run it and anyone can challenge it. Disagreement stops being a clash of whole worldviews and becomes a dispute about a particular rule, which evidence can address.

Which group? Collective experience

One person's judgement about a hidden mechanism is a report. It becomes knowledge when it is set against the experience of others. But which others? The people in one room are a small group. The widest group available is collective experience: what many people, in many places and times, have found, including the accumulated findings of research.

This differs from a view of knowledge common in modern thought. Enlightenment thinkers treated every claim as simply true or false. When they saw that parts of their inherited knowledge were corrupt, they concluded that it should be thrown out and rebuilt from scratch by each reasoning individual. No one can do this. Each of us accepts nearly all of what we inherit and questions only a small part of it. Even our questions are asked in a shared language, whose categories already carry a vast set of assumptions about how the world is divided up.

The Islamic tradition does not depend on that binary of true and false. It treats inherited knowledge and the settled agreement of a community as reliable for the most part, though never certain. In this book, collective experience is the state of the art: where we start, and what any new model must answer to. Chapter 2 shows what this means for a single game played many thousands of times.

That is also why this book asks you to bring your work to others: your accounts, your models, your interviews, and especially the models that failed. A model that failed, with the point at which it failed, adds to everyone's experience.

Looking ahead

The next chapter puts these ideas to work on one very simple game, the ultimatum game, which has been played many thousands of times around the world. It shows what fit can

and cannot settle, how collective experience separates rival motives, and why intentions matter.

After that, we build our first model, of a market you probably know from the inside: students arriving at a university and looking for rooms to rent. Hal Varian opens his textbook *Intermediate Microeconomics* with this market. We will build it with ten students and ten owners, run it by hand, and ask whether the textbook's rent emerges, and what would have to be true of the market for it to do so.

Exercises

1. **The metro seat.** Take your three accounts from section 2 to a small group. Which account does the group find most plausible, and why? For each pair of accounts, describe a situation in which they would predict different behaviour.
2. **A third cheating model.** Design a third model of the quiz class that also produces an overall rate of one copied answer in five. What evidence would distinguish your model from Models A and B, and how would you gather it?
3. **Learning from experience.** Describe something you learned from experience. Then describe a different lesson the same experience could have taught. What led you to the lesson you drew?
4. **What the map leaves out.** Look up a route in a map app. List three things the route leaves out. For each, name a purpose for which leaving it out would be a serious mistake.
5. **From an account to a rule.** Take one of your accounts of the seat-giving and write it as a rule an agent could follow. Then describe a situation in which your rule and a rival rule would lead to different behaviour.

Further reading

Zaman, Asad, *Logical Positivism and Islamic Economics* (December 30, 2013). International Journal of Economics, Management and Accounting, Vol 21, No. 2, pp1-18.
<https://ssrn.com/abstract=2195043>

How logical positivism arose, how it shaped the methodology of economics, including the theory of revealed preference, and why it conflicts with an Islamic view of knowledge.