1

Reason on Trial

In the cold autumn of 1619, René Descartes, then aged twenty-three and a volunteer in the armies of the Duke of Bavaria, found himself in what is now southern Germany with time to spend and nobody around he deemed worth talking to. There, in a stove-heated room, as he recounts in his *Discourse on Method,*[1] he formed the stunningly ambitious project of ridding himself of all opinions, all ideas learned from others, and of rebuilding his knowledge from scratch, step by step. Reason would be his sole guide. He would accept as true only what he could not doubt.

Descartes justified his rejection of everything he had learned from others by expressing a general disdain for collective achievements. The best work, he maintained, is made by a single master. What one may learn from books, he considered, "is not as close to the truth, composed as it is of the opinions of many different people, as the simple reasoning that any man of good sense can produce about things in his purview."[2]

Descartes would have scorned today's fashionable idea of the "wisdom of crowds." The only wisdom he recognized, at least in the sciences, was that of individual reason: "As long as one stops oneself taking anything to be true that is not true and sticks to the right order so as to deduce one thing from another, there can be nothing so remote that one cannot eventually reach it, nor so hidden that one cannot discover it."[3]

Why did Descartes decide to trust only his own mind? Did he believe himself to be endowed with unique reasoning capacities? On the contrary, he maintained that "the power of judging correctly and of distinguishing the true from the false (which is properly what is called good sense or reason) is

naturally equal in all men."[4] But if we humans are all endowed with this power of distinguishing truth from falsity, how is it that we disagree so much on what is true?

"The Greatest Minds Are Capable of the Greatest Vices as Well as the Greatest Virtues"

Most of us think of ourselves as rational. Moreover, we expect others to be rational too. We are annoyed, sometimes even angry, when we see others defending opinions we think are deeply flawed. Hardly ever do we assume that those who disagree with us altogether lack reason. What aggravates us is the sense that these people do not make a proper use of the reason we assume they have. How can they fail to understand what seems so obvious (to us)?

If reason is this highly desirable power to discover the truth, why don't people endowed with it use it to the best of their capacities all the time? After all, we expect all sighted people to see what others see. Show several people a tree or a sunset, and you expect them all to see a tree or a sunset. Ask, on the other hand, several people to reason about a variety of questions, from logical problems to social issues, and what might surprise you is their coming to the same conclusions. If reason, like perception, worked to provide us with an adequate grasp of the way things really are, this should be deeply puzzling.

Descartes had an explanation: "The diversity of our opinions arises not from the fact that some of us are more reasonable than others, but solely that we have different ways of directing our thoughts, and do not take into account the same things. . . . The greatest minds are capable of the greatest vices as well as the greatest virtues."[5]

This, however, is hardly more than a restatement of the enigma, for shouldn't the way we direct our thoughts itself be guided by reason? Shouldn't reason, in the first place, protect us from intellectual vices?

Descartes was the most forceful of reason's many advocates. Reason has also had many, often passionate, detractors. Its efficacy has been questioned. Its arrogance has been denounced. The religious reformer Martin Luther was particularly scathing: "Reason is by nature a harmful whore. But she shall not harm me, if only I resist her. Ah, but she is so comely and glittering. . . . See

to it that you hold reason in check and do not follow her beautiful cogitations. Throw dirt in her face and make her ugly."[6]

To be fair, Descartes's and Luther's views on reason were much richer and subtler than these isolated quotes suggest, and hence less diametrically opposed. Luther's invectives were aimed at the claims of reason in matters of faith. In a different context, the same Luther described reason, much more conventionally, as "the inventor and mentor of all the arts, medicines, laws, and of whatever wisdom, power, virtue, and glory men possess in this life" and as "the essential difference by which man is distinguished from the animals and other things."[7] Descartes for his part abstained, out of conviction or out of prudence, from critically examining faith in the light of reason.

Still, if reason were put on trial, both the prosecution and the defense could make an extraordinary case. The defense would argue, citing Descartes, Aristotle, Kant, or Popper, that humans err by not reasoning enough. The prosecution would argue, citing Luther, Hume, Kierkegaard, or Foucault, that they err by reasoning too much.

The defense and the prosecution could also produce compelling narratives to bolster their case.

Eratosthenes and the Unabomber

Do you doubt the power of reason? Just look at the sciences, the defense would exclaim. Through insightful reasoning, scientists have discovered hidden facts and deep explanations that would have been completely inaccessible otherwise. Modern science provides countless examples of the power of reason, but nothing beats, as a simple and compelling illustration, the measurement of the circumference of the earth twenty-two centuries ago, by Eratosthenes (276–195 BCE), the head librarian of the greatest library of the ancient world at Alexandria in Egypt.[8]

Already at the time, it was commonly accepted that the earth was spherical rather than flat. This best explained the curvature of the horizon at sea and the apparent movement of the sun and the stars. Still, it was, as the phrase goes, "just a theory." No one had traveled around the earth, let alone seen it from a distance as astronauts now have. How, then, could its circumference be measured?

Eratosthenes had heard that every year, on a single day, at noon, the sun shone directly to the bottom of wells in the distant town of Syene (now Aswan). This, he understood, meant that, there and then, the sun was at the zenith, vertically above the town. Syene therefore had to be on the Tropic of Cancer and that single day had to be the summer solstice (our June 21). Syene, he assumed, was due south on the same meridian as Alexandria. He knew how long it took caravans to travel from Alexandria to Syene and, on that basis, estimated the distance between the two cities to be 5,014 stades (an ancient unit of measure).

When, on the summer solstice at noon, the sun was vertically above Syene, by how many degrees was it south of the vertical in the more northern city of Alexandria? Eratosthenes measured the length of the shadow cast at that very moment by an obelisk located in front of his library (or so the story goes). He determined that the sun's rays were hitting the obelisk at an angle of 7.2 degrees south of the vertical. He understood that the sun was far enough to treat all rays that reach the earth as parallel, and that therefore the angle between the rays of the sun and the vertical at Alexandria was equal to the angle between the vertical at Alexandria and that at Syene, two lines that cross at the center of the earth (see Figure 2). In other words, that very angle of 7.2 degrees also measured the difference in degrees of latitude between Alexandria and Syene. He now had all the information he needed. Since 7.2 degrees is one-fiftieth of 360 degrees, Eratosthenes could calculate the circumference of the earth by multiplying by fifty the distance between Alexandria and Syene. The result, 252,000 stades, is 1 percent shy of the modern measurement of 24,859 miles, or 40,008 kilometers.[9]

Eratosthenes grasped the mutual relevance of apparently unrelated pieces of evidence (the pace of caravans, the sun shining to the bottom of wells, the shadow of an obelisk), of assumptions (the rotundity of the earth, its distance from the sun), and of simple geometrical ideas about angles and parallel lines. He drew on all of them to measure a circumference that he could imagine but neither see nor survey. What made his measurement not just true but convincing is—isn't it?—that it was a pure product of human reason.

How telling, the prosecution would object, that the defense of reason should choose as evidence such an exceptional achievement! It is an exception, and this is why it is still remembered after more than two thousand years. Ordi-

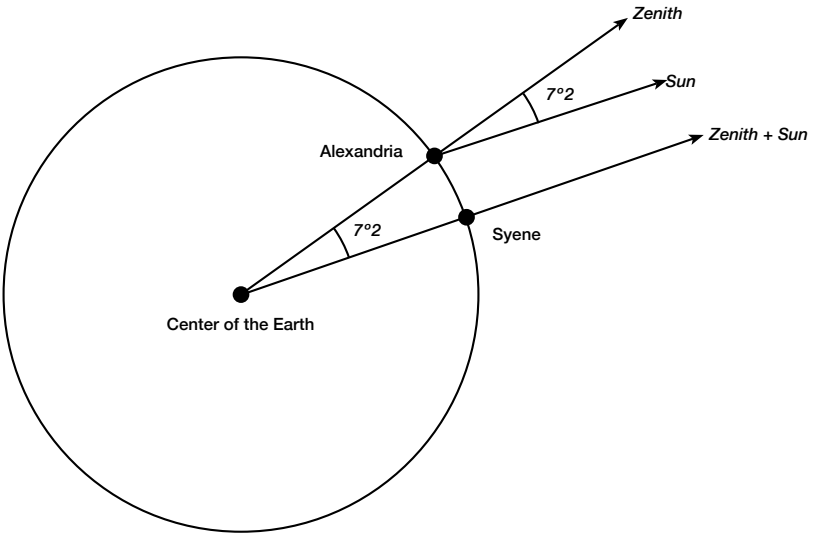

Figure 2. How Eratosthenes computed the circumference of the earth.

nary reasoning doesn't lead us far, and that is just as well, as often it leads in the wrong direction. Even extraordinary uses of reason, far from being all on the model of Eratosthenes, have led many thinkers badly astray. Publishers, newspapers, and scientific journals receive every day the thoroughly reasoned nonsense of would-be philosophers, scientists, or reformers who, failing to get their work published there, then try the World Wide Web. Some of them, however, reason not just to theoretical but also to practical absurdities, act on them, and achieve notoriety or even infamy. The prosecution might well at this juncture introduce the case of Ted Kaczynski.

As a young man, Kaczynski was unquestionably a brilliant reasoner. He had entered Harvard in 1958, at age sixteen. For his doctoral dissertation at the University of Michigan, he solved a mathematical problem that had eluded his professors for years, prompting the University of Berkeley to hire him. Two years later, however, he abandoned mathematics and academe to live in a shack in Montana, where he became an avid reader of social science and political work. Both his readings and his writings focused on what he saw as the destructive character of modern technology. Viewing technological progress as leading to disasters for the environment and for human dignity is not

uncommon in Western thought, but Kaczynski went further: for him, only a violent revolution causing the collapse of modern civilization could prevent these even greater disasters.

To help trigger this revolution, Kaczynski began in 1978 to send bombs to universities, businesses, and individuals, killing three people and injuring many others. He wrote a long manifesto and managed to have it published in the *New York Times* and in the *Washington Post* in 1995 by promising that he would then "desist from terrorism." The Unabomber, as the FBI had named him, was finally arrested in 1996 and now, as we write, serves a life sentence without the possibility of parole in a Colorado jail, where he goes on reading and writing.

What had happened to the brilliant young mathematician? Had Kaczynski's reason failed him, turning him into the "raving lunatic" described by the press? Kaczynski's family arranged for his defense to try to make him plead insanity. The defense of reason would no doubt concur: *un*reason had to be the culprit. It is unlikely, however, that Kaczynski suffered at the time of his arrest from any major mental disorder. He was still a smart, highly articulate, extremely well-read man. Defective reasoning, the prosecution of reason would insist, cannot be blamed for his actions. To see this, all you need do is read the Unabomber's manifesto:

> The Industrial Revolution and its consequences have been a disaster for the human race. . . . They have destabilized society, have made life unfulfilling, have subjected human beings to indignities, . . . and have inflicted severe damage on the natural world. The continued development of technology will worsen the situation. . . . The industrial-technological system may survive or it may break down. . . . If the system survives, the consequences will be inevitable: There is no way of reforming or modifying the system so as to prevent it from depriving people of dignity and autonomy. If the system breaks down the consequences will still be very painful. But the bigger the system grows the more disastrous the results of its breakdown will be, so if it is to break down it had best break down sooner rather than later. We therefore advocate a revolution against the industrial system.[10]

This, surely, is a well-constructed argument. Most of us would disagree with the premise that technological progress is a plain disaster, but actually, many well-respected philosophers and social theorists have defended similar views. What singles out Kaczynski, the prosecution of reason would claim, is that he pushed this radically pessimistic viewpoint to its logical consequences and acted accordingly. As one of his biographers put it: "Kaczynski, in short, had become a cold-blooded killer not despite of his intellect, but *because of it.*"[11]

So, the defense of reason would counter, the prosecution wants you to believe that the problem with Ted Kaczynski, the Unabomber, is that he was reasoning too much. His manifesto is indeed more tightly reasoned than much political discourse. What made him notorious, however, were not his ideas but his crimes. Nowhere in his writings is there even the beginning of a proper argument showing that sending bombs to a few powerless academics—his former colleagues—would kick-start a "revolution against the industrial system." When you are told that excessive reliance on reasoning led someone to absurd or abhorrent conclusions, look closely at the evidence, and you will find lapses of reason: some premises were not properly examined, and some crucial steps in the argument are simply missing. Remember: a logical demonstration can never be stronger than its weakest part.

Expert Witnesses for the Prosecution

Since historical illustrations, however arresting, are not sufficient to make their cases, defense and prosecution of reason would turn to expert witnesses. Neither side would have any difficulty in recruiting psychologists to support their cause. Specialists of reasoning do not agree among themselves. Actually, the polemics in which they are engaged are hot enough to have been described as "rationality wars." This very lack of agreement among specialists who, one hopes, are all good reasoners, is particularly ironic: sophisticated reasoning on reasoning does not come near providing a consensual understanding of reasoning itself.

The prosecution of reason might feel quite smug. Experimental psychology of reasoning has been fast developing since the 1960s, exploiting a variety of ingenious experiments. The most famous of these present people with

problems that, in principle, could easily be resolved with a modicum of simple reasoning. Yet most participants in these experiments confidently give mistaken answers, as if the participants were victims of some kind of "cognitive illusion." These results have been used in the rationality wars to argue that human reason is seriously defective. Reason's defenders protest that such experiments are artificial and misleading. It is as if the experiments were aimed at tricking sensible people and making them look foolish rather than aimed at understanding the ordinary workings of reason. Of course, psychologists who have devised these experiments insist that, just as visual illusions reveal important features of ordinary, accurate vision, cognitive illusions reveal important features of ordinary reasoning.[12] Philosophers, science writers, and journalists have, however, focused on the seemingly bleak implications of this research for the evaluation of human rationality and have, if anything, exaggerated their bleakness.

When you do arithmetic, it does not matter whether the numbers you add or subtract happen to be numbers of customers, trees, or stars, nor does it matter whether they are typical or surprising numbers for collections of such items. You just apply rules of arithmetic to numbers, and you ignore all the rest. Similarly, if you assume that reasoning should be just a matter of applying logic to a given set of premises in order to derive the conclusions that follow from these premises, then nothing else should interfere. Yet there is ample evidence that background knowledge and expectations do interfere in the process. This, many argue, is the main source of bad reasoning.

Here is a classic example.[13] In July 1980, Björn Borg, who was then hailed as one of the greatest tennis players of all time, won his fifth consecutive Wimbledon championship. In October of that year, Daniel Kahneman and Amos Tversky, two Israeli psychologists working in North America who would soon become world-famous, presented a group of University of Oregon students with the following problem:

> Suppose Björn Borg reaches the Wimbledon finals in 1981. Please rank order the following outcomes from most to least likely:
>
> 1. Borg will win the match.
> 2. Borg will lose the first set.

3. Borg will lose the first set but win the match.
4. Borg will win the first set but lose the match.

Seventy-two percent of the students assigned a higher probability to outcome 3 than to outcome 2. What is so remarkable about this? Well, if you have two propositions (for instance, "Borg will lose the first set" and "Borg will win the match"), then their conjunction ("Borg will lose the first set but win the match") cannot be more probable than either one of the two propositions taken separately. Borg could not both lose the first set and win the match without losing the first set, but he could lose the first set and not win the match. Failing to see this is an instance of what is known as the "conjunction fallacy." More abstractly, take two propositions that we may represent with the letters *P* and *Q*. Whenever the conjunction "*P* and *Q*" is true, so must be both *P* and *Q*, while *P* could be true or *Q* could be true and "*P* and *Q*" false. Hence, for any two propositions *P* and *Q*, claiming that their conjunction "*P* and *Q*" is more probable than either *P* or *Q* taken on its own is clearly fallacious.

Kahneman and Tversky devised many problems that caused people to commit the conjunction fallacy and other serious blunders. True, as they themselves showed, if you ask the same question not about Björn Borg at Wimbledon but rather about an unknown player at an ordinary game, then people do not commit the fallacy. They correctly rank a single event as more probable than the conjunction of that event and another event. But why on earth should people reason better about an anonymous tennis player than about a famous champion?

Here is another example from our own work illustrating how the way you frame a logical problem may dramatically affect people's performance.[14] We presented people with the following version of what, in logic, is known as a "pigeonhole problem":

> In the village of Denton, there are twenty-two farmers. All of the farmers have at least one cow. None of the farmers have more than seventeen cows. How likely is it that at least two farmers in Denton have the exact same number of cows?

Only 30 percent gave the correct answer, namely, that it is *certain*—not merely probable—that at least two farmers have the same number of cows. If you don't see this, perhaps the second version of the problem will help you.

To another group, we presented another version of the problem that, from a logical point of view, is strictly equivalent:

> In the village of Denton, there are twenty-two farmers. The farmers have all had a visit from the health inspector. The visits of the health inspector took place between the first and the seventeenth of February of this year. How likely is it that at least two farmers in Denton had the visit of the health inspector on the exact same day?

This time, 70 percent of people gave the correct answer: it is certain.

As the Borg and the farmers-cows problems illustrate, depending on how you contextualize or frame a logical problem—without touching the logic of it—most people may either fail or succeed. Isn't this, the prosecution would argue, clear evidence that human reason is seriously defective?

Expert Witnesses for the Defense

While many psychologists focused on experiments that seem to demonstrate human irrationality, other psychologists were pursuing a different agenda: to identify the mental mechanisms and procedures that allow humans to reason at all.

There is little doubt that some simple reasoning (in a wide sense of the term) occurs all the time, in particular when we talk to each other. Conjunctions such as "and," "or," and "if" and the adverb "not" elicit logical inferences of the most basic sort. Take a simple dialogue:

Jack (to Jill): I lent my umbrella to you or to Susan—I don't remember whom.
Jill: Well, you didn't lend it to me!
Jack: Oh, then I lent to Susan.
Jill: Right!

No need for Jack or Jill to have studied logic to come to the conclusion that Jack lent his umbrella to Susan.[15] But what is the psychological mechanism by means of which such inferences are being performed? According to one type of account, understanding the word "or" or the word "not" amounts to having in mind logical rules that somehow capture the meaning of such words. These rules govern deductions licensed by the presence of these "logical" words in a statement. Here is a rule for "or" (using again the letters *P* and *Q* to represent any two propositions):

"Or" rule: From two premises of the form "*P* or *Q*" and "not *P*," infer *Q*.

Several psychologists (Jean Piaget, Martin Braine, and Lance Rips, in particular[16]) have argued that we perform logical deduction by means of a "mental logic" consisting in a collection of such logical rules or schemas. When Jack and Jill infer that Jack lent his umbrella to Susan, what they do is apply the "or" rule.

According to an alternative explanation, "mental model theory" (developed by Philip Johnson-Laird and Ruth Byrne),[17] no, we don't have a mental logic in our head. What we have is a procedure to represent and integrate in our mind the content of premises by means of models comparable to schematic pictures of the situation. We then read the conclusions off these models. In one model, for instance, Jack lent his umbrella to Jill. In an alternative model, he lent it to Susan. If Jack's statement is true, then the two mental models can neither be both right nor be both wrong. When we learn that the "lent to Jill" model is wrong, then we are left with just the "lent to Susan" model, and we can conclude that Jack lent his umbrella to Susan.

Much work in the psychology of reasoning has been devoted to pitting against one another the "mental logic" and the "mental models" approaches. You might wonder: What is the difference between these two accounts? Aren't they both stating the same thing in different terms? Well, true, the two theories have a lot in common. They both assume that humans have mechanisms capable of producing genuine logical inferences. Both assume that humans have the wherewithal to think in a rational manner, and in this respect, they contrast with approaches that cast doubt on human rationality.

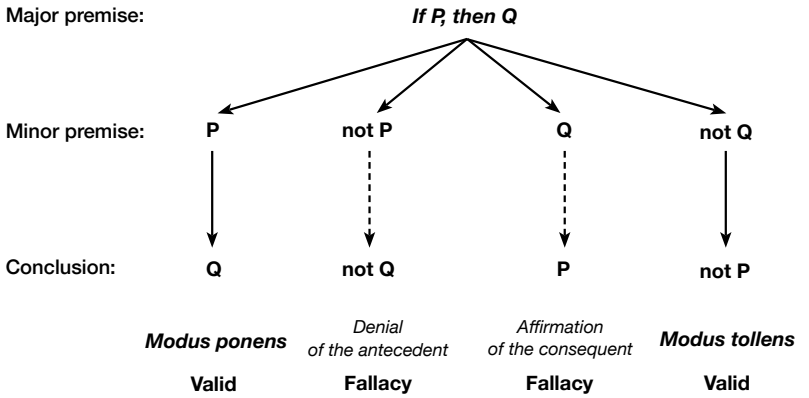

Figure 3. The four schemas of conditional inference.

The picture drawn by "mental logicians" and "mental modelers" is not quite rosy, however. Both approaches recognize that all except the simplest reasoning tasks can trip people and cause them to come to unwarranted conclusions. As they become more complex, reasoning tasks rapidly become forbiddingly difficult and performance collapses. But what makes a reasoning task complex? This is where the two theories differ. For mental logicians, it is the number of steps that must be taken and rules that must be followed. For mental modelers, it is the number of models that should be constructed and integrated to arrive at a certain conclusion.

The defense of reason would want these two schools to downplay their disagreements and to focus on a shared positive message: humans are equipped with general mechanisms for logical reasoning. Alas, the prosecution would find in the very work inspired by these two approaches much evidence to cast doubt on this positive message.

If there is one elementary pattern of reasoning that stands out as the most ubiquitous, the most important both in everyday and in scholarly reasoning, it is what is known as conditional reasoning—reasoning with "if . . . , then . . ." (see Figure 3). Such reasoning involves a *major premise* of the form "if *P*, then *Q*." For instance:

If you lost the key, then you owe us five dollars.
If pure silver is heated to 961°C, then it melts.

If there is a courthouse, then there is a police station.
If Mary has an essay to write, then she will study late in the library.

The first part of such statements, introduced by "if," is the *antecedent* of the conditional, and the second part, introduced by "then," is the *consequent.* To draw a useful inference from a conditional statement, you need a second premise, and this *minor premise* can consist either in the affirmation of the antecedent or in the denial of the consequent. For instance:

If there is a courthouse, then there is a police station. *(major premise: the conditional statement)*
There is a courthouse. *(minor premise: affirmation of the antecedent)*
———
There is a police station. *(conclusion)*

Or:

If there is a courthouse, then there is a police station. *(major premise: the conditional statement)*
There is no police station. *(minor premise: denial of the consequent)*
———
There is no courthouse. *(conclusion)*

These two inference patterns, the one based on the affirmation of the antecedent (known under its Latin name, *modus ponens*) and the one based on the denial of the consequent *(modus tollens),* are both logically valid: when the premises are true, the conclusion is necessarily true also.

But what about using as the minor premise the denial of the antecedent (rather than its affirmation) or the affirmation of the consequent (rather than its denial)? For instance:

If there is a courthouse, then there is a police station. *(major premise: the conditional statement)*

There is no courthouse. *(minor premise: denial of the antecedent)*

———

There is no police station. *(conclusion?)*

Or:

If there is a courthouse, then there is a police station. *(major premise: the conditional statement)*
There is a police station. *(minor premise: affirmation of the consequent)*

———

There is a courthouse. *(conclusion?)*

These two inference patterns (known by the name of their minor premise as "denial of the antecedent" and "affirmation of the consequent") are invalid; they are fallacies. Even if both premises are true, the conclusion does not necessarily follow—you may well, for instance, have a police station but no courthouse.

Surely, the prosecution would exclaim, all this is simple enough. Shouldn't people, if the defense were right, reliably perform the two valid inferences of conditional reasoning and never commit the two fallacies? Alas, the expert witnesses of the defense have demonstrated in countless experiments with very simple problems that such is not the case—far from it. True, nearly everybody draws the valid *modus ponens* inference from the affirmation of the antecedent. Good news for the defense? Well, the rest is good news for the prosecution: only two-thirds of the people, on average, draw the other valid inference, *modus tollens,* and about half of the people commit the two fallacies.[18] And there is worse . . .

Will She Study Late in the Library?

In a famous 1989 study, Ruth Byrne demonstrated that even the valid *modus ponens* inference, the only apparently safe bit of logicality in conditional reasoning, could all too easily be made to crumble.[19] Byrne presented participants with the following pair of premises:

Major premise:
If Mary has an essay to write, then she will study late in the library.
Minor premise:
She has an essay to write.

Participants had no difficulty deducing:

Conclusion: Mary will study late in the library.

So far, so good. To another group of people, however, Byrne presented the same problem, but this time with an additional major premise:

First major premise:
If Mary has an essay to write, then she will study late in the library.
Second major premise:
If the library stays open, then Mary will study late in the library.
Minor premise:
She has an essay to write.

From a strictly logical point of view, the second major premise is of no relevance whatsoever. So, if people were logical, they should draw the same valid *modus ponens* conclusion as before. Actually, only 38 percent of them did.

What Byrne was trying to prove was not that humans are irrational—mental modelers don't believe that—but that mental logicians have the wrong theory of human rationality. If, as mental logicians claim, people had a mental *modus ponens* rule of inference, then that inference should be automatic, whatever the context. Participants are instructed to take the premises as true, so, given the premises "If Mary has an essay to write, then she will study late in the library" and "Mary has an essay to write," they should without hesitation conclude that she will study late in the library. What about the possibility that the library might be closed? Well, what about it? After all, for all you know, Mary might have a pass to work in the library even when it is closed. A logician would tell you, just don't go there. This is irrelevant to this logic task, just as the possibility that a bubble might burst would be irrelevant to the arithmetic task of adding three bubbles to two bubbles.

Did mental logicians recognize, in the light of Byrne's findings, that their approach was erroneous? Well, no; they didn't have to. What they did instead was propose alternative explanations.[20] People might, for instance, consolidate the two major premises presented by Byrne into a single one: "If Mary has an essay to write *and* if the library stays open, then Mary will study late in the library." This, after all, is a realistic way of understanding the situation. If this is how people interpret the major premises, then the minor premise, "She has an essay to write," is not sufficient to trigger a valid *modus ponens* inference, and Byrne's findings, however intrinsically interesting, are no evidence against mental logic.

Is There a Defendant at This Trial?

The prosecution of reason might enjoy watching mental logicians and mental modelers, all expert witnesses for the defense, fight among themselves, but surely, at this point, the jury might grow impatient. Isn't there something amiss, not with the reasoning of people who participate in these experiments, but rather with the demands of psychologists?

Experimentalists expect participants to accept the premises as true whether those premises are plausible or not, to report only what necessarily follows from the premises, and to completely ignore what is merely likely to follow from them—to ignore the real world, that is. When people fail to identify the logical implications of the premises, many psychologists see this as proof that their reasoning abilities are wanting. There is an alternative explanation, namely, that the artificial instructions given to people are hard or even, in many cases, impossible to follow.

It is not that people are bad at making logical deductions; it is that they are bad at separating these deductions from probabilistic inferences that are suggested by the very same premises. Is this, however, evidence of people's irrationality? Couldn't it be seen rather as evidence that psychologists are making irrational demands?

A comparison with the psychology of vision will help. Look at Figure 4, a famous visual illusion devised by Edward Adelson. Which of the two squares, A or B, is of a lighter shade of gray? Surely, B is lighter than A—this couldn't be an illusion! But an illusion it is. However surprising, A and B are of exactly the same shade.

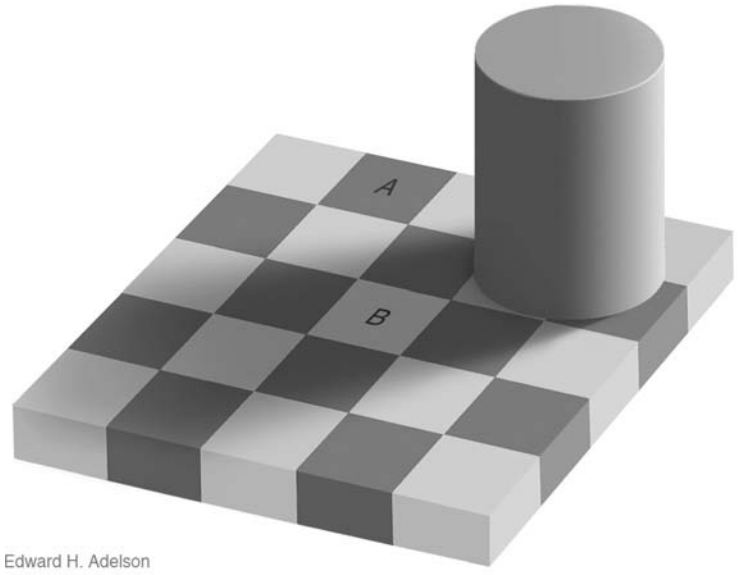

Figure 4. Adelson's checkerboard illusion.

In broad outline, what happens is not mysterious. Your perception of the degree to which a surface is light or dark tracks not the *amount* of light that is reflected to your eyes by that surface but the *proportion* of the light falling on that surface that is reflected by it. The higher this "reflectance" (as this proportion is called), the lighter the surface; the lower this reflectance, the darker the surface:

$$\text{reflectance} = \frac{\text{light reflected by the surface}}{\text{light falling on the surface}}$$

The same gray surface may receive and therefore reflect more or less light to your eyes, but if the reflectance remains the same, you will perceive the same shade of gray. Your eyes, however, get information on just one of the two quantities—the light reflected to your eyes. How, then, can your brain track reflectance, that is, the proportion between the two quantities, only one of which you can sense, and estimate the lightness or darkness of the surface? To do so, it has to use contextual information and background knowledge and infer the other relevant quantity, that is, the amount of light that falls on the surface.

When you look at Figure 4, what you see is a picture of a checkerboard, part of which is in the shadow of a cylinder.

Moreover, you expect checkerboards to have alternating light and dark squares. You have therefore several sound reasons to judge that square B—one of the light squares in the shade—is lighter than square A—one of the dark squares receiving direct lighting. Or rather you would have good reasons if you were really looking at a checkerboard partly in the shadow of a cylinder and not at a mere picture. The illusion comes from your inability to treat this picture just as a two-dimensional pattern of various gray surfaces and to ignore the tridimensional scene that is being depicted.

Painters and graphic designers may learn to overcome this natural tendency to integrate all potentially relevant information. The rest of us are prey to the illusion. When discovering this illusion, should we be taken aback and feel that our visual perception is not as good as we had thought it to be, that it is betraying us? Quite the opposite! The ability to take into account not just the stimulation of our retina but what we intuitively grasp of the physics of light and of the structure of objects allows us to recognize and understand what we perceive. Even when we look at a picture rather than at the real thing, we are generally interested in the properties of what is being represented rather than in the physical properties of the representation itself. While the picture of square A on the paper or on the screen is of the same shade of gray as that of square B, square A would be quite darker than square B on the checkerboard that this picture represents. The visual illusion is evidence of the fact that our perception is well adapted to the task of making sense of the three-dimensional environment in which we live and also, given our familiarity with images, to the task of interpreting two-dimensional pictures of three-dimensional scenes.

Now back to Mary, who might study late in the library. In general, we interpret statements on the assumption that they are intended to be relevant.[21] So when given the second major premise, "If the library stays open, then Mary will study late in the library," people sensibly assume that they are intended to take this premise as relevant. For it to be relevant, it must be the case that the library might close and that this would thwart Mary's intention to study late in the library. So, yes, participants have been instructed to accept as absolutely true that "if Mary has an essay to write, then she will study late in

the library," and they seem not to. However, being unable to follow such instructions is not at all the same thing as being unable to reason well. Treating information that has been intentionally given to you as relevant isn't irrational—quite the contrary.

It takes patience and training for a painter to see a color on the canvas as it is rather than as how it will be perceived by others in the context of the whole picture. Similarly, it takes patience and training for a student of logic to consider only the logical terms in a premise and to ignore contextual information and background knowledge that might at first blush be relevant. What painters do see and we don't is useful to them as painters. The inferences that logicians draw are useful to them as logicians. Are the visual skills of painters and the inferential skills of logicians of much use in ordinary life? Should those of us who do not aspire to become painters or logicians feel we are missing something important for not sharing their cognitive skills? Actually, no.

The exact manner in which people in Ruth Byrne's experiment are being reasonable is a matter for further research, but that they are being reasonable is reasonably obvious. That people fail to solve rudimentary logical problems does not show that they are unable to reason well when doing so is relevant to solving real-life problems. The relationship between logic on the one hand and reasoning on the other is far from being simple and straightforward.

At this point, the judge, the jury, and our readers may have become weary of the defense's and the prosecution's grandstanding. The trial conceit is ours, of course, but the controversy (of which we have given only a few snapshots) is a very real one, and it has been going on for a long time. While arguments on both sides have become ever sharper, the issue itself has become hazier and hazier. What is the debate really about? What is this capacity to reason that is both claimed to make humans superior to other animals and of such inferior quality? Do the experiments of Kahneman and Tversky on the one hand and those of "mental logicians" and "mental modelers" on the other hand address the same issue? For that matter, is the reasoning they talk about the same thing as the reason hailed by Descartes and despised by Luther? Is there, to use the conceit one last time, a defendant in this trial? And if there is, is it reason itself or some dummy mistaken for the real thing? Is reason really a thing?