Chapter 6

Coping With Uncertainty

At their next session, Mia began by saying "There are many beliefs and theories that do not have sufficiently high credibility to be labeled 'true.' When the credibility of a belief is below that needed for it to be called true, then we simply say it's uncertain. Sometimes, we can give it a credibility number, like there's a 70% chance of rain tomorrow, or a 60% chance that your cold will be over by Thursday."

> I can live with doubt and uncertainty and not knowing—I think it's much more interesting to live not knowing than to have answers that might be wrong. I have approximate answers and possible beliefs and different degrees of certainty about different things but I'm not absolutely sure of anything and there are many things I don't know anything about . . . —Richard Feynman, Physicist[1]

"A lot of medical knowledge is uncertain." Sam said. "For example, aren't there several theories about what causes cancer?"

"Right," Nick said, "And the diagnosis and treatment of many diseases also involve elements of uncertainty."

"And there are everyday beliefs that people have," Mia interjected. "What about a belief like 'the New England Patriots are going to win the

[1]From his interview with *Nova*, January 25, 1983.

Super Bowl'? Even ardent fans would admit they are at least a bit uncertain about it."

"Although there is debate about how uncertainty should be measured," Mia continued, "the most obvious method involves probabilities. As we mentioned earlier, Gio uses probabilities to indicate the strengths of his beliefs. There are a lot of technical details connected with the subject of probability, but we don't need to mention many of them."

"Thanks!" said Sam.

"But we do need to say enough to make the point that reasoning with probabilities requires some special modes of thinking," Mia said. "The first thing to know is that probability values are always numbers between 0 and 1."

"Isn't that the same as saying that they could be between 0% and 100%?" Sam asked.

"Expressing them as percentages is often done," Mia answered. "Anyway, if the probability of something is 1, then you are completely certain of it and would label it 'true.' If the probability of it is 0, then you would label it 'false.' And the probability of something and the probability of its negation must sum to 1. Either it is or it isn't. And, if there are several possibilities, the sum of all of their probabilities must equal 1."

"That's if they are all mutually exclusive and exhaustive," Nick amplified.

"What does that mean?" asked Sam.

"Mutually exclusive means that of all the possibilities only one of them can be the case," Nick replied. "It can't both rain and not rain, for example. Exhaustive means that the listed possibilities cover *all* possible cases. You can't leave any out."

"As an example," Nick continued, "let's suppose that there are only three people who might have written *Macbeth*, namely Shakespeare, Jonson, or Marlowe, and one, and only one, of them must have written it. Then, if you are certain that Shakespeare wrote it—that is, you think the probability that Shakespeare wrote it is 1—then you must also be certain that neither Jonson nor Marlowe did—that is, you must think that the probabilities that either one of them did are 0. Of course, you could be uncertain, and believe Shakespeare wrote it with probability 0.85, that Jonson did with probability 0.05, and that Marlowe did with probability 0.1."

"Ok," said Sam, "but how do we put probability numbers on beliefs? A professor in my geology class told us that there would probably be a strong earthquake in Los Angeles within the next thirty years, but what probability number should I attach to that belief?"

"There are two major ways to go about assigning probabilities," Mia replied. "The most obvious method involves collecting statistics about a large number of cases. Suppose, for example, that during the last 50 years, it rained three times in Pasadena on New Years Day. Then the probability of rain in Pasadena on New Years Day next year could be guessed to be about 3/50 or 0.06."

"That's if you think that the general pattern of the weather during the last fifty years will continue into next year," Nick elaborated.

"Yes," Mia agreed, "as statisticians might say, the statistics have to be *stationary*. Anyway, this version of assigning probabilities to events is called the *frequency method*. It's commonly used in many situations for which it's possible to collect a lot of data. People in the insurance business use it, for example, to estimate the probability that people will live to various ages. They look at the data."

"Statistical records and statistical data of all kinds are very useful for learning things," Mia continued. "For example, Gio learns to recognize what a person is saying by listening to many samples of what that person says. Those samples are data. Gio's internal mechanisms use this data to compute various probability values concerning relationships among the components of a person's speech, and he uses these values to make estimates about what words the person most likely is saying."

"But what about events for which there aren't enough statistics?" asked Sam. "What's the probability of a big earthquake in Los Angeles in the next thirty years?"

"People can make guesses about that," Mia replied, "but without a lot of data, the guesses would be pretty subjective. In fact, probabilities based on guesses are called *subjective* probabilities."

"Do you mean that people just make up some numbers?" Sam asked.

"No, not usually," Mia replied. "The guesses are based on informed opinion about the matter at hand. For example, geologists specializing in the science of earthquakes would base their earthquake probability estimates on many geo-technical considerations."

"People quote odds that a certain football team will win," Sam said. "Odds are given in horse races too. Are they related to subjective probabilities?"

"Yes, they can be," Mia replied. "Odds are just another way of expressing probabilities. If the odds on something happening are 4 to 1 against, then the probability of it happening is 0.2 or 20%, and the probability of it not happening is 0.8 or 80%. The odds are obtained by dividing 80 by 20. In football, the odds are sometimes subjectively estimated by experts who review past team performances and other possibly relevant factors such as whether or not rain is expected and whether certain key players are not going to play because of injuries. All of this information is combined into a judgment they make about the odds."

"Sometimes it's easier for an expert to use words instead of numbers to describe the likelihood of something," Mia continued. "For example, an expert medical diagnostician might say that it's 'extremely likely' that a certain patient has viral pneumonia. Or maybe she would say that it is 'somewhat likely.' So phrases like 'virtually certain,' 'extremely likely,' 'somewhat likely,' 'possible,' 'unlikely,' 'almost certainly not,' and so on might be preferred to numbers."

"I can translate words like those into numbers," Gio said. "Then I can use the numbers to do my calculations. For example, I translate 'virtually certain' into 0.999, and 'somewhat likely into 0.7."

"What if the so-called expert is wrong about the probabilities he or she guesses?" Sam asked. "Why should one trust just one expert?"

"Right," Mia agreed, "It's better to try to combine the judgments of several experts."

"How is that done?" asked Sam.

"One way is to use what is called a *pari-mutuel betting system*," Nick answered. "Several people place bets using initial odds provided by an expert. The results of all of these bets are used continuously, as they are made, to update the odds. For example, if there is a total of $3,000 so far bet on team A and \$1,000 bet on team B, then the odds in favor of team A can be set at 3 to 1."

"Another method is based on markets," Mia said. "People buy and sell 'contracts' about various contested beliefs, and the price of the contract is used to estimate probabilities."

"Do you mean like the stock market?" Sam asked.

"It's a bit more like futures markets for commodities," Mia answered. "Let's take football again as an example. 'Futures' are sold on teams A and B in a market set up for that purpose. The market maker and the contract purchaser have to agree on a price for the contract. Suppose the agreed upon price is \$0.75 for a contract on team A. That is, the contract purchaser pays \$0.75 for a one-dollar contract on team A. Then, if team A wins, the holder of the contract is paid one dollar. If team A doesn't win, the contract is worthless. Anytime before the game, people can buy and sell contracts to each other, and the price of the latest contract can be published for all to see—just like in a futures market. In this way, a price is established that reflects all of the information possessed by all of the people participating. This price will change over time as people take into account the latest information. If at some time the price for a one-dollar contract on team A is \$0.80, for example, the subjective probability that team A will win can be taken to be 0.8, and the odds in favor of team A would be 4 to 1."

"But," worried Sam, "if *everyone* can participate in the market, won't the odds be influenced by all the non-expert bidders? Shouldn't we just pay attention to how the 'smart money' invests?"

"Well, in principle the market ought to take care of that automatically," Mia guessed. "As long as there are enough experts around who judge the odds to be different from those so far established by the market and who have confidence in their opinions, they can bid heavily with a view to making a large profit at the expense of the more ill-informed participants. Their bidding will work to bring the odds more in accord with expert opinion."

"There's one problem I've thought of with those markets though," Nick said.

"What's that?" Sam asked.

"When a person buys a contract about some proposition, x, in a market," Nick replied, "he could be expressing confidence in x of course, but he may also simply believe that the market price of x will rise. For example, he may not really be very confident about x at all, but he may think that other people, foolishly in his opinion, will want to buy contracts at increasingly higher prices. He may be buying solely because he thinks he can sell his contract later at a higher price. And, if there are many people buying for that reason, the market for x may experience a 'bubble'."

"Yes," Mia agreed, "and a crash too, so watch out!"

"So, you're saying that the opinions of expert geologists, for example, would dominate the market about earthquakes in Los Angeles," Sam concluded. "But are there really markets of that sort?"

"Actually, yes," Mia said. "One such is the Iowa Electronic Markets (IEM). Others are the Foresight Exchange (FX) and the Hollywood Stock Exchange (HSX).[2] Traders in these markets can 'invest' in the outcomes of a wide variety of unresolved outcomes. For example, at FX, one can bet on whether physicists will discover the Higgs boson by certain dates in the future. At HSX, one can bet on who will win Oscar, Emmy, and Grammy awards. Prices in these markets are said to correlate well with actual award outcome frequencies."

"What happens for contracts where you might never know the outcome?" Sam asked.

"You could still have a market," Mia said. Market prices would fluctuate as new information comes in, and holders of contracts could always sell them to people who think they can still make a profit by buying a contract now and selling it to someone else at a higher price later. And, if someone thinks the price will drop, he can sell short."

"Is there a market for who wrote *Macbeth*?" Sam wondered.

"I wouldn't be surprised," Mia said. "At least there are scholars who debate the subject. But I imagine you would have to pay pretty close to a dollar for a dollar contract on Shakespeare."

[2]The websites are: http://www.biz.uiowa.edu/iem, http://www.ideosphere.com, and http://www.hsx.com. For an analysis, see "The Real Power of Artificial Markets," by David M. Pennock, Steve Lawrence, C. Lee Giles, and Finn Årup Nielsen, *Science*, 291(5506): 987, 9 February, 2001.

"So how do these market ideas relate to 'truth'?" asked Sam. "Let's see—if everyone participating in a market labeled some statement, A, true, then the odds in favor of it would be pretty high, and a one-dollar contract on A would cost a dollar, right?"

"That's an interesting way to define truth," Nick responded. "If you labeled something as true, you wouldn't want to bet against it, even if someone gave you very high odds. If a community believed a statement to be true, then a market involving that community wouldn't discount the contract on that statement. Of course, as we agreed earlier, just because you or a community label a belief true doesn't change reality."

"Here's what Justice Oliver Wendell Holmes, Jr. said," Gio offered.

> . . . the best test of truth is the power of the thought to get itself accepted in the competition of the market . . .[3]

"But Holmes probably wasn't talking about the kinds of markets Mia mentioned," Sam objected.

"But he was talking about betting," Gio said. "The author Lewis Menand thinks so anyway."

> Holmes, James, Peirce, and Dewey . . . said repeatedly [that beliefs] are just bets on the future.[4]

"Well, what about situations in which no one has any guess at all about how something will turn out?" Sam wondered. "How do you assign probabilities then?"

Mia replied "In a betting situation, the probabilities would then be distributed equally among all the possible different outcomes. So if there were no reason to favor team A or team B in football—assuming there could not be a tie—the probability that A would win would equal 0.5. That would also be the probability that B would win. Assigning probabilities that way is an instance of what is called *Laplace's Law of Indifference*.

[3]See *Abrams v. United States*, 250 U.S. 616,630 (1919).

[4]Louis Menand, *The Metaphysical Club: A Story of Ideas in America*, p. 440, New York: Farrar, Straus and Giroux, 2001.

Presumably market prices of contracts would reflect indifferent probabilities also if no one in the market had any opinions about outcomes."

"Here are some more examples using the law of indifference," Mia continued. "If you throw a single six-sided die—one of a pair of dice—and have no reason to suspect that it is 'loaded,' then the probability that it would come up 2, say, is 1/6. As are the probabilities that it would come up any one of the other five numbers. And if you flipped a fair coin, the probability that it would come up heads—or, alternatively, tails—is 1/2."

"Ok, suppose I assign probabilities to some of my beliefs. How do I use these uncertain beliefs to decide on actions?" asked Sam.

"You need to know more than just the probabilities," Mia answered. "You need to know how you value the consequences of the actions based on those beliefs."

"Here's how Gio does it," Mia went on. "Let's suppose he estimates that the probability that his favorite bridge across the Madison River is closed to be 0.4. Should he take a chance and try to cross the river using that bridge? He needs to know how he would value the outcomes. In this case there are four different possibilities:

1. Gio decides to use the bridge and the bridge is closed
2. Gio decides to use the bridge and the bridge is open
3. Gio decides not to use the bridge and the bridge is closed
4. Gio decides not to use the bridge and the bridge is open

"In order for Gio to decide what to do, he has to assign 'points,' called *utilities*, to each of these possibilities. What's the utility for Gio if he decides to use the bridge and it turns out the bridge is open?"

"I see," said Sam. "It's like a game in which Gio gets points depending on what he decides to do and what turns out to be the case. And Gio's decision about what to do should be the one that gets him the most points."

"Exactly," Mia replied. "Well, continuing with the example, let's make up a chart where we list the points that Gio assigns for each of the four cases that I mentioned. Here's a chart with the point scores and probabilities."

	Bridge is open 0.6	Bridge is closed 0.4
Decide to use bridge	10 (hooray!)	−8 (have to backtrack to take another bridge)
Decide not to use bridge	−4 (missed an opportunity)	2 (dodged a bullet)

"So, if Gio decides to use the bridge and the bridge is open," Mia continued, "he gets 10 points. But there is only a 60% chance that the bridge is open, so on the average, he would get only 6 points from that box. And, if he decides to use the bridge and the bridge is closed he would get –8 points; that is he would lose 8 points. So on the average he would get 0.4 x –8 = –3.2 points from that box. Taking both outcomes into account, if he decides to use the bridge, he could expect, on average, to get 6 – 3.2 = 2.8 points. A similar calculation for the case in which he decides not to use the bridge would give him –2.4 + 0.8 = –1.6 points. So, if the points in the boxes accurately reflect the consequences for Gio in each of those situations, he should go for it! Decide to use the bridge."

"But what if Gio does decide to use the bridge and the bridge is actually closed?" asked Sam. "It seems then that all those fancy calculations produced the wrong decision."

"Not necessarily!" Mia insisted. "One mustn't confuse a bad outcome with a bad decision! If the probabilities and the consequences are as we say they are, Gio made the right and *rational* decision. Of course, if finding out that the bridge was closed reminds Gio of something he should have taken into account *before* he found out that the bridge was closed, something that would have led him to estimate different probabilities or point values *before going to the bridge*, then perhaps his decision was based on the wrong numbers, and he could learn something."

"But how does one assign points to consequences?" asked Sam.

"Well, I assume you must know something about what the consequences mean to you," Mia replied. "Otherwise the decision couldn't have mattered very much to you anyway!"

"So, what if I have one of several choices to make and there are lots of different possible outcomes?" Sam asked.

"Same idea," Mia answered. "You list all your choices along the left hand side of a chart, and you list all the different outcomes with their probabilities along the top of the chart. Then you fill the boxes in with point scores for all the different combinations of decisions and outcomes and do the calculations. Select that decision that gives you the best average score."

"And what if it's hard to assign probability values to outcomes?" Sam asked.

"You can do the calculations several times with slightly different probability values and points to see if the best decision is very sensitive to the probability and point values, Mia suggested."

"And what if you have no clue about the probabilities?" Sam persisted.

"There is a field of study called 'decision-making under ambiguity'," Mia said. "If you had no reason to judge one outcome as any more probable than the others, you could use Laplace's law of indifference and divide probability values equally among all of the outcomes. Or, if you were unusually optimistic and not averse to risk, you could opt for that decision that had the highest score in its row. In our example, Gio would decide to use the bridge—optimistically assuming it would be open and getting a score of 10 if it was. If you were pessimistic and risk averse, you could opt for that decision whose lowest score in its row was best. Using that approach, Gio would decide not to use the bridge and settle for a worst-case score of –4."

"Well, I either have to decide to try the bridge or not," observed Gio, "too bad I can't do both—then either way I'd be right."

"Most decisions that we must make have this 'either-or' character," Nick replied, "and we have to do one or the other. But there are some where we can actually do a little bit of both."

"Really?" asked Sam. "What are those?"

"Suppose you own 100 shares of a stock and you think there is a 60% chance it will go up in the next week and a 40% chance it will go down," Nick suggested. "If you have to liquidate by the end of the week, you could sell 40 shares now and keep 60 shares to sell at the end of the week. That's what's called pursuing a 'mixed strategy.' Here's what can be said about it—if the same situation came up often, if the probabilities were accurate, and if you always used such a mixed strategy, in the long run you would maximize your return."

"Are there any situations where I could use a mixed strategy?" Gio asked.

"I suppose if we made you pay for battery re-charging and gave you information that would let you estimate future electric costs and their probabilities, you could use a mixed strategy to decide how much to charge your batteries today instead of waiting for a possibly lower price tomorrow," Mia replied. "Some actions can be modulated in this way—that is, they can be executed more-or-less forcefully. If you base one of these kinds of actions on an uncertain belief, the strength of the action ought not to be stronger than the strength of the belief. A half-hearted belief should permit no more than a half-hearted action."

> Conviction is something you need in order to act, ... But your action needs to be proportional to the depth of evidence that underlies your conviction. —Paul O'Neill, former Secretary of the Treasury[5]

"So, it sounds like dealing with probabilities is pretty important for sizing up beliefs and using them to decide on actions," Sam said.

"Yes," Mia replied, "and working with probabilities can sometimes be tricky."

"I do it effortlessly," Gio said.

"People have more trouble," Mia said. "There are many subtle matters to think about when dealing with probabilities. Some of these can be illustrated by card games. Take poker, for example. Getting dealt a hand with four aces is quite rare. But getting *any particular* hand is very rare also!

"There are lots of different hands," Gio said.

[5]Quoted by Ron Suskind in *The Price of Loyalty: George W. Bush, the White House, and the Education of Paul O'Neill*, p. 325, New York: Simon & Schuster, 2004.

“Will you calculate the number for us, Gio?” Mia asked. “How many different ways are there of dealing five cards out of a standard 52-card deck?”

“The number of different combinations is 2,598,960,” Gio answered.

“So the probability of getting any one of these is about one chance in two and a half million,” Mia said. “But you don’t say, ‘what a miracle! I got a hand that comes up only one time in two and half million!’ After all, you must get *some* hand!”

“When I play poker, I usually get a pretty worthless hand,” Sam complained.

“That’s because there are lots of hands that, unlike four aces, aren’t worth very much in poker,” Mia said.

“Roughly half of the possible hands can’t even beat a pair,” Gio said.

“Look at the odds against all of us being here having this conversation right now!” Nick said. “Lots of things had to fall into place just right. But, here we are!”

“Some of the reasoning needed when dealing with uncertain outcomes is counter-intuitive, and that’s why it’s important to know how to calculate with probabilities if you don’t want to be led astray on some important decisions,” Mia said. “Here’s another example. Suppose I’m giving a lecture in a large auditorium and announce that there is someone among the audience who has an uncanny ability to predict the outcomes of coin tosses. Then, I proceed with the following demonstration. First I ask all of those in the audience to make a prediction about the result of my next coin toss. I flip the coin, announce the result and ask all of those who predicted correctly to please stand up. We have to assume that everyone is being honest about this! In a large audience, about half of the people will stand up. Then I ask just those people who are standing to predict the outcome of the next coin toss. I toss the coin, announce the result, and ask those people who guessed incorrectly to please sit down. Now, somewhere around a quarter of the audience will be left standing. I continue this process until there is only one person standing. This might take eight or nine tosses, depending on the size of the audience and the way the coins actually fall. The one person who remains standing predicted every toss correctly! Does that person have special, magical powers of prediction? No, not necessarily.

Someone had to be left standing!"

"Wait a minute," said Sam. "Suppose that at a certain point in your process there are two persons left standing and that they both guess incorrectly on the next toss. After that, there would be no one left standing."

"Same idea," Mia replied. "Just *before* that last toss someone could make the claim that those *two* people had an uncanny ability to predict coin tosses."

"Something similar may go on among so-called experts at stock market forecasting or any other kind of forecasting, for that matter," Nick said. "If you have enough forecasters, *someone* among them is bound to have been right about the past. Would that same person be right the next time? Well, maybe—if there really were some basis on which accurate predictions could be made. But, otherwise, success is just a chance event."

"Sometimes probability calculations give surprising results," Mia said. "For example, did you know that if you have 23 people in a room, the probability that at least two of them will share the very same birthday—that is, day and month but not necessarily year—is over 50%?"

"I don't believe it," said Sam.

"Don't get caught betting against it," Mia cautioned. "Here's the calculation. Actually, it's easier to calculate first the probability that *no* two people will share the same birthday. Let's do it by letting people come into the room one at a time. When the second person comes in, the probability that he will *not* share a birthday with the first person already there is 364/365. When the third person comes in, the probability that he will not share a birthday with either of the first two is 363/365. And so on until the 23^{rd} person comes in. So we calculate the probability that no two of them share the same birthday to be equal to
$(364/365) \times (363/365) \times \ldots \times (343/365)$."

"That's equal to 0.492703," Gio volunteered.

"Then, the probability that at least two of them *will* share the same birthday is $1 - 0.492703$ which is 0.507297," Mia concluded.

"Well, as you said," Sam agreed, "reasoning with probabilities can be tricky."

"Many people believe a lot of silly things because they lack sophistication about probabilities," Mia said. "Knowing something about probability theory is very important for critical thinking."

"Here is what the economist John Maynard Keynes said about the subject," Gio said.

> In most branches of academic logic, such as the theory of the syllogism or the geometry of ideal space, all the arguments aim at demonstrative certainty. They claim to be *conclusive*. But many other arguments are rational and claim some weight without pretending to be certain. In Metaphysics, in Science, and in Conduct, most of the arguments upon which we habitually base our rational beliefs are admitted to be inconclusive in a greater or less degree. Thus for a philosophical treatment of these branches of knowledge, the study of probability is required.[6]

"I myself use 'Bayes networks' for reasoning with uncertain beliefs," Gio said.[7]

"Bayes networks? What are they?" asked Sam.

"They are used to link together my beliefs," Gio replied. "For example, if one of my beliefs, say B, affects another of my beliefs, say C, then my network would have a connection between those two beliefs."

"You can think of them as diagrams that show how the probabilities of various beliefs are related," Mia elaborated. "Several important reasoning strategies, even ones we humans use, can be illustrated using these diagrams."

"How about some examples?" Sam asked.

"Ok," Mia said, "Suppose a belief in B strongly supports a belief in C. The Bayes network fragment for that relationship might look like this:"

[6] John Maynard Keynes, *A Treatise on Probability*, p. 8, London: MacMillan and Co. Ltd., 1948. (First edition, 1921.)

[7] Kevin B. Korb and Ann E. Nicholson, *Bayesian Artificial Intelligence*, Boca Raton, FL: Chapman & Hall/CRC, 2004.

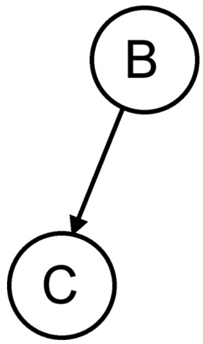

Gio then interrupted with a specific example: "I believe that if it's cold outside, the building heat will be on."

"Here is how Gio would represent that," Mia said.

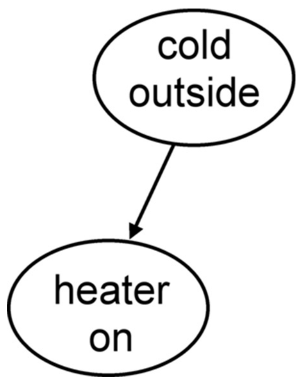

"The interesting thing," Mia continued, "is that we can often use the relationship backwards."

"Observing that the building heat is on increases the probability of it being cold outside," Gio volunteered. "I wouldn't need to go outside to find out."

"That's called *evidential* reasoning," Mia said. "We use positive evidence to increase our belief in a related proposition."

"Sometimes it's called *abductive* reasoning," Gio said.

"Evidential reasoning is what medical diagnosticians do when they try to decide which disease might be causing some observed symptoms," Mia said. "A disease very often has definite symptoms, and observing the symptoms is suggestive of the disease."

"Are you saying that doctors could use Bayes networks to help them decide what disease a patient might be suffering from?" Sam asked.

"There are diagnostic programs based on Bayes networks," Mia replied. "Some doctors use them for routine screening. At least, the modern ones and younger ones do. The programs seem to do a good job of capturing the skills of diagnosticians who have many years of experience and training."

"Evidential reasoning can be used in a negative way also," Mia continued. "For example, if a diagnostician *doesn't* observe a symptom associated with a particular disease, he or she would then tend to doubt that disease. That's what diagnosticians do when they attempt to rule out a disease. They check for symptoms whose absence rules out the disease."

"For example," Gio said, "if I notice that the building heat is *not* on, my belief that it is cold outside would weaken—maybe to the extent that I would conclude that it is not cold outside."

"Evidential reasoning can accumulate," Mia said. "Two or more pieces of *independent* evidence give extra support for a belief. We can elaborate Gio's network about it being cold outside to provide an example:"

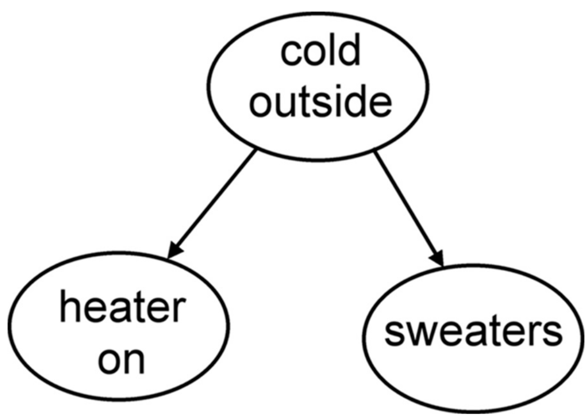

"If Gio sees people outside wearing sweaters, in addition to observing that the heater is on, that increases the weight of evidence for it being cold outside."

"We can use Gio's example to illustrate another kind of reasoning too," Mia continued. "It's called *explaining away*. Observing that the heat is on is evidence for it being cold outside *unless* the heater being on is explained by some other belief. For example, suppose Gio is told that the heating system is being tested. The network fragment in this case would look like this:"

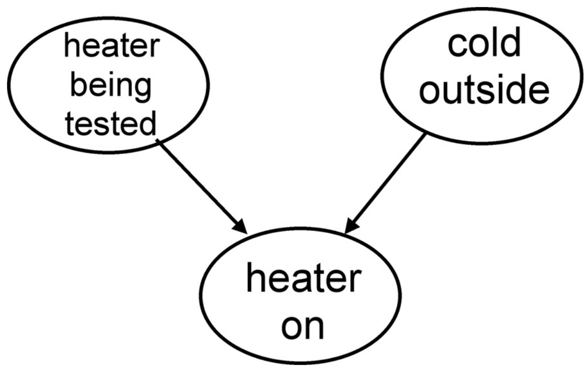

"Hearing that the heater is being tested would 'explain away' the evidence for it being cold outside."

"A belief in extra-sensory perception, ESP, can be weakened by explaining away anecdotal evidence for it by attributing such evidence to chance coincidence," Nick added.

"I've thought of another example," Sam said. "When you first see a magic trick, you might think there really is some kind of magic going on. If the magician showed you how he really did it, that would explain away the magic."

"Exactly right," said Mia.

"Interestingly," Nick commented, "the explaining-away strategy can be used to *strengthen* beliefs as well as weaken them. It all depends on whether we are talking about negatives or positives. Believers in ESP might attempt to explain away negative evidence by introducing other reasons for failure such as the subject wasn't trying hard enough and so on."

"People have worked out how to calculate the amounts by which probability values for beliefs should be increased or decreased given the probability values for other beliefs in the network," Mia said. "The mathematics is a bit complicated, so we won't go into it here."

"The calculations are easy for me," Gio said. "I use them constantly to keep the probabilities of all of my beliefs in tune with each other."

"Keeping everything 'in tune,' as Gio says, can be thought of as making his beliefs cohere with one another," Nick said. "We might say that Gio has been programmed with a version of the coherence theory of truth."

"Do you always use Bayes networks for your reasoning, Gio?" Sam asked.

"Only when I'm not quite certain about the statements I'm reasoning about," Gio answered. "When I'm confident about my beliefs, I use ordinary logic because the calculations are simpler."

"And what do you mean by that?" Sam asked.

"I have a lot of logical rules built into my software," Gio answered. "And ways to use those rules. For example, there's *modus ponens*."

"Modus what?" Sam asked.

"*Modus ponens*. It's a rule of inference. If I believe some statement, let's call it P, and I believe some other statement 'P implies Q', then, using *modus ponens*, I can infer Q," Gio answered.

"Oh, like 'All men are mortal,' and 'Socrates is a man,' therefore 'Socrates is mortal'?" guessed Sam.

"That sort of thing," Gio answered. "A lot of the time logical reasoning works just fine for me. I use Bayes nets only when I'm not certain about things."

"When the odds in favor of beliefs get very high, we don't need probabilities," Mia agreed. "But some 'certainties' occasionally slip back into the uncertain category."

"Well, this discussion of probabilities has been interesting, and I suppose I could consult a futures market to get probabilities for some of my beliefs," Sam said, "but what if I wanted to second-guess the market? How can I do my own evaluations?

"You could use something like the scientific method," Mia suggested.

"I know that scientists use it to help them with their theories," Sam said, "but can it be used for everyday beliefs too?"

"Mia and I think so," Nick replied.

"Well, let's hear more then," Sam said.

"I need to check on some programs I've been running," Mia said. "We can talk about science next time."