SUMMARY THE BOOK OF WHY

JUDEA PEARL DANA MACKENZIE



Summary of "The Book of Why" by Judea Pearl and Dana MacKenzie

Written by Lea Schullery

The New Science of Cause and Effect

Introduction	5
The Beginning of the Causal Revolution	6
We Must Use Counterfactual Data to Truly Understand Results	8
The First Step of the Ladder of Causation Deals with Passive Observation 10	
Rung Two of the Ladder of Causation Is About Taking Action	12
The Third Rung of the Causation Ladder is About Identifying Countractuals 14	
Confounding Bias is a Lurking Third Variable that Scientists Must Tak Account	e into 16
Causal Relationships and Human Thought is the Key for Artificial Intelligence	18
Final Summary	20



Go to QuickRead.com/App now to download our app and get access to thousands of free book summaries as both text and audiobooks.

Get the key insights of non-fiction books in minutes instead of hours. Listen to our free audiobooks while you workout or on your commute to work.





Introduction

One simple question asked some tens of thousands of years ago, allowed humans the power to discover new ideas and create remarkable innovations. As a result of answering this question, humans have created organized societies, towns, cities, and eventually the science- and technology-based civilization we live in today. All of this came to be because humans asked a simple question: Why? You see, humans are able to recognize that certain things cause other certain things and that tinkering with one can change the other. No other species can understand this, not like humans. This new science is what Judea Pearl and Dana MacKenzie refer to as "causal inference," and it presumes that the human brain is the most advanced tool ever designed for managing causes and effects. Our brains can store a vast amount of causal knowledge, giving us the power to unleash this knowledge and answer some of the most pressing questions of all time. Questions like: "How effective is a given treatment in preventing disease? Did the new tax law cause our sales to go up, or was it our ad campaign? What is the health-care cost attributable to obesity? Can hiring records prove an employer is guilty of a policy of sex discrimination? And even, I'm about to quit my job. Should I?" These questions, while different, all concern themselves with cause-and-effect relationships. Today, science allows us to not just ask these questions, but answer them as well. This new science has created a simple mathematical language that can now be used to combine our knowledge with data and answer causal questions like the ones above.

Authors Judea Pearl and Dana MacKenzie hope that the new science of causal inference will help us better understand how humans understand cause-and-effect relationships better than computers and data. Furthermore, "in the age of computers, this new understanding can also bring the prospect of amplifying our innate abilities so that we can make better sense of data, both big and small."

The Beginning of the Causal Revolution

For the last few decades, the phrase "correlation does not imply causation" has been a mantra chanted by scientists. It's been largely accepted, partly due to the work and research of Karl Pearson, a twentieth-century English mathematician who zealously worked to prove that causation was nothing more than a special case of correlation. Simply put, Pearson believed that data was all there was to science. Nothing else. He believed this was true because causation could not be proven or represented by data. Therefore, causation is scientifically unacceptable.

Pearson belonged to a philosophical school called *positivism*, which believed that the universe is a product of human thought and that science is only a description of those thoughts. Therefore, causation is an objective process that happens outside the human brain, thus could have no scientific meaning. Prepared to discard causation completely, Pearson further proved his point by identifying correlations he believed were false or bogus. For instance, one such correlation is the one between a nation's per capita chocolate consumption and its number of Nobel Prize winners. This correlation seems ridiculous because we cannot fathom the idea that eating chocolate can cause a Nobel Prize!

What Pearson failed to point out, however, is that more people in wealthy, Western countries eat chocolate, and the Nobel Prize winners are typically from those countries. The ironic thing is that this is a causal explanation, which, for Pearson, is not necessary for scientific thinking. Furthermore, geneticist Sewall Wright later proved that causation could be represented mathematically, which he discovered while researching at Harvard University. While studying the coats and markings on guinea pigs, Wright sought to determine how hereditary their markings were.

Wright began to doubt that genetics alone determined the amount of white and suggested that *developmental factors* in the womb were causing the different markings and variations; therefore, he estimated the developmental factors by creating a mathematical formula. In Wright's case, the desired and unknown quantity is represented by *d*, which is the effect of "developmental factors" on white fur. Other quantities included in the equation included *h*, for "hereditary" factors, which are also unknown. Finally, Wright showed that if we know the causal quantities, then we can predict correlations in the data by a simple graphical rule. By creating a path diagram to represent these relationships, Wright demonstrated that developmental factors have an effect on the gestation period which have then have an effect on coat pattern, and so on.

By turning this path diagram into a mathematical equation, Wright was able to determine that 42 percent of the variation in coat pattern was due to heredity, and 58 percent was developmental. He later went on to publish a general paper called "Correlation and Causation" that explained how path analysis worked in settings other than guinea pig breeding. Of course, given the times, Wright was criticized and his findings were argued and debunked by his peers. Little did his peers know, however, that Wright's findings were just the beginning of the Causal Revolution.

We Must Use Counterfactual Data to Truly Understand Results

For scientists, data is everything. Data is essential when trying to determine the underlying cause of an effect or vice versa. Therefore, scientists rely on data for much of their work. Unfortunately, data can be skewed and can lead to erroneous misinterpretations. To show just how important and pivotal data can be, let's take a look at the public debate that erupted in Europe when the smallpox vaccine was introduced.

Unexpectedly, data showed that more people died from the smallpox vaccine than from smallpox itself. Naturally, some people used this information to argue that the vaccine should be banned, even though it was actually saving lives and eradicating the disease. Here's how that data became skewed: "Suppose that out of 1 million children, 99 percent are vaccinated and one percent are not. If a child is vaccinated. He or she has one chance in one hundred of developing a reaction, and the reaction has one chance in one hundred of being fatal." In other words, the numbers add up to 99 fatalities. "Meanwhile, 10,000 don't get vaccinated, 200 get smallpox, and 40 die from the disease. In summary, more children die from vaccination (99) than from the disease (40)."

Of course, parents were marching to the health department with signs saying, "Vaccines kill!" And can you blame them? The data seemed like it was on their side, but we must take a closer look at the numbers to truly understand the data. To do this, we must ask ourselves, "What if we had set the vaccination rate to zero?" If this were the case, we can conclude that out of 1 million children, 20,000 would have gotten smallpox, and 4,000 would have died. When we compare the counterfactual world with the real world, we see that *not* vaccinating would have cost the lives of 3,861 children.

Furthermore, data can be used to show the relationship between a child's shoe size and reading ability. "Children with larger shoes tend to read at a higher level. But the relationship is not one of cause and effect. Giving a

child larger shoes won't make him read better!" Instead, the variables can be explained by the child's age. The older the child, the better the reading ability. Pearson lacked similar common sense when making the Nobel Prize correlation with eating chocolate. For instance, if we look *only* at sevenyear-olds, then we can expect to see no relationship between shoe size and reading ability. It's this kind of junction that allows us to begin climbing the Ladder of Causation.

The First Step of the Ladder of Causation Deals with Passive Observation

It was the story of Adam and Even in the Garden of Eden that made author Judea Pearl recognize the creation of human knowledge and understanding. He recalls being concerned about the "notion that the emergence of human knowledge was not a joyful process but a painful one, accompanied by disobedience, guilt, and punishment." Is human knowledge worth giving up a carefree life in Eden? Surely, the agricultural and scientific revolutions that followed were worth all the economic hardships, wars, and social injustices of modern life, right? This philosophical question is what led Pearl to confront the Ladder of Causation.

He recognized that God asked Eve, "What is that you have done?" And Eve replied, "The serpent deceived me, and I ate." God was asking Eve "what," and Eve answered, "why." God asked for the facts and received explanations. In other words, humans have always been fascinated by the intricacies of cause-and-effect relationships.

There are three levels of causation or three rungs of the ladder. Most animals, as well as present-day learning machines, are on the first rung, learning by association. For instance, this is what the owl does as it observes its prey move and figures out where the rodent is going to be at the time the owl strikes. The owl doesn't concern itself with asking *why*. Similarly, a computer Go program does this when it studies a database of millions of Go games to figure out which moves are associated with a higher percentage of wins. This first rung relies on making predictions based on passive observations and is characterized by the question, "What if I see…"

For example, a marketing director might ask a director at a department store, "How likely is a customer who bought toothpaste to also buy dental floss?" Questions like these are the basis of statistics, and the first step in answering them is by collecting and analyzing data. To answer this question, we must first look at the data on the shopping behavior of customers who bought toothpaste. We must then compute the proportion of those customers who also bought dental floss. This proportion is called a *conditional probability*, and we can write it symbolically like this: *P(floss* | *toothpaste)* or "What is the probability of floss, given that you see toothpaste?"

Statisticians use methods like the one above to identify associations between variables; however, statistics alone cannot tell us about cause and effect. Is toothpaste or floss the cause? For the sales manager, it doesn't really matter. "Good predictions need not have good explanations. The owl can be a good hunter without understanding why the rat always goes from point *A* to point *B*." Systems that work at the first level of the Ladder of Causation lack flexibility and adaptability, but when we step up to the next level of causal queries, we begin to change the world. Questions like, "What will happen to our floss sales if we double the price of toothpaste?" requires a new kind of knowledge which we will find at rung two of the Ladder of Causation.

Rung Two of the Ladder of Causation Is About Taking Action

The defining query of the second rung of the Ladder of Causation is "What if we do…." Or, "if we *change* the environment, what will happen?" The "do" is important. The second rung is characterized by actively influencing outcomes, unlike the first rung which only relies on passive observation. How will doubling the price of toothpaste affect floss sales? We could write it symbolically like this: $P(floss \mid do(toothpaste))$, which asks about the probability of selling floss at a certain price, given that we set the price of toothpaste at another price.

Additionally, the manager might recognize that he has too much toothpaste in the warehouse. So, he asks the question, "How can we sell it?" or "What price should we sell it for?" These questions require an action, an intervention, which we perform mentally before deciding what kind of action to take. This requires a causal model. In fact, we perform interventions like this in our daily lives all the time. For example, if you have a headache, you might take an aspirin to cure it. You intervene on one variable to affect another one.

Unlike the first rung of the ladder, computers cannot perform on the second rung. They cannot answer these types of questions; after all, computers are not like humans who understand cause-and-effect relationships. Humans, on the other hand, have the ability to test the effect of something through controlled experiments, something humans have been doing since Biblical times. You see, when Ashpenaz, the overseer of King Nebuchadnezzar's court was tasked with identifying the best of the captured nobles to serve in the court, he was faced with a problem. As part of the education of these children, they would get to eat royal meat and drink royal wine. Of course, this is where the problem occurred.

One of his favorites, a boy named Daniel, refused to eat the meat for religious reasons. Daniel could not eat meat that wasn't prepared according

to Jewish laws, and so he asked that he and his friends be given a diet of vegetables instead. To prove that vegetables wouldn't affect their performance, Daniel asked Ashpenaz to conduct a controlled experiment. For ten days, four of them would only be fed vegetables while the rest feasted on the King's meat and wine. After ten days, they would compare the two groups. The results of the experiment proved that the vegetarian diet gave Daniel and his three companions more strength and mental stamina. In fact, the King was so impressed, that he gave each of them a favored place in his court.

You see, when Ashpenaz was faced with Daniel's problem, he asked a question about causation: Will a vegetarian diet cause my servants to lose weight? To answer this question, Daniel proposed a controlled experiment by setting up two groups of people who were similar in many ways and comparing the two after some time. While Daniel's experiment was strikingly modern for his time, he didn't think of one thing: confounding bias (which we will later explore in more depth). For instance, if Daniel's group was overall healthier than the control group, then their diet would have nothing to do with their healthy appearance! Thinking about all of the various factors that affect an experiment leads us to the third and final rung of the Causation Ladder.

The Third Rung of the Causation Ladder is About Identifying Countractuals

In the previous chapter, we discussed the causal relationship of taking aspirin to cure a headache. Once that headache is gone, you might begin to wonder *why* it is now gone. Was it the aspirin you took? The food you ate? The good news you heard? It's these kinds of queries that bring us to the top run of the Ladder of Causation, the level of counterfactuals. Counterfactuals require us to go back in time, change history, and ask "What would have happened if I had not taken the aspirin?"

Counterfactuals, unlike data, aren't always factual. But the human brain is constantly seeking explanations for such scenarios. For instance, Eve explained the reason for her actions was that the serpent deceived her. It is this ability that distinguishes humans from animal intelligence and machines. In fact, we see counterfactuals all the time in the courtroom as it is very old and known in the legal profession as "but-for-causation." For example, "if the defendant fired a gun and the bullet struck and killed the victim, the firing of the gun is a but-for, or necessary, cause of the death, since the victim would be alive if not for the firing." Similarly, if Joe blocks a fire exit with furniture and Judy dies in the fire because she could not escape, then Joe is legally responsible for her death even though he didn't light the fire itself.

Speaking of fire, a classic example that demonstrates necessary causation is that of a fire that broke out after someone struck a match. Many would argue that the fire wouldn't have happened if it weren't for the match being lit, but they forget to take the presence of oxygen into account. We ignore the causal relationship between oxygen and fire. Unlike humans, a computer cannot think in terms of causal relationships. According to a computer, both the match and oxygen would play an equal role in the fire since they are both *necessary causes*. As a result, the computer would determine that the oxygen is to blame for the fire. Furthermore, a computer would likely calculate the match as the *sufficient cause* of the fire. While both the presence of the match and oxygen are both necessary for the fire to break out, the computer can reason that the match was sufficiently responsible, making it the cause of the fire. All of this simply means that necessary and sufficient causes are crucial for answering causal questions and play an important role in the third rung of the Ladder of Causation. But now that you understand the three rungs, what's next?

Confounding Bias is a Lurking Third Variable that Scientists Must Take into Account

As we mentioned in a previous chapter, controlled experiments have been around for as long as humans. We also mentioned that Daniel's experiment was modern in all ways except one: he failed to take confounding bias into account. "Confounding bias occurs when a variable influences both who is selected for the treatment and the outcome of the experiment." Additionally, they are typically associated with the second rung of the Ladder of Causation, which means intervention is required to adjust the experiment.

For example, if you were testing a drug and gave it to patients who are younger on average than the people in the control group, then age becomes a confounder. Age becomes a "lurking third variable." If there is no data on the ages at all, then the results from the test can't necessarily be trusted or true. But if we know the confounding variable Z is age, then we can compare the treatment and control groups in every age group separately. In this scenario, we can take an average of the effects, weighting each age group according to its percentage in the target population. This method of compensation is called "adjusting for Z" or "controlling for Z."

Confounders, however, aren't always easy to compensate for. For instance, in the 1950s and 1960s debate about the link between smoking and lung cancer, confounders could be just about anything. Some suggested there could be a smoking gene that caused people to crave cigarettes and also made them more likely to develop lung cancer. Of course, there was an easy way for statisticians to test the effect of smoking through a randomized controlled trial (RCT) in which treatment, such as smoking, is randomly assigned to some individuals and not to others and the observed changes are then compared. But randomization, in this case, was certainly not ethical. Researchers couldn't ethically assign a random group of people to smoke for 30 years to test the link to cancer, the results could be deadly!

Of course, the debate between smoking and cancer wasn't about tobacco nor cancer. It was all about the innocuous word "caused," and it wasn't the first time physicians were confronted with perplexing causal questions. In fact, it was the mid-1700s when James Lind discovered that citrus fruits could prevent scurvy. But now we are met with the question, "What is it about citrus fruits that prevent scurvy?" At the time, vitamins weren't even invented yet, so scientists couldn't say *which* citrus fruits worked better than others at preventing the disease. So if we don't know *why* oranges work, we might be tempted to try another fruit if we run out of oranges.

Therefore, scientists use *mediation* to help answer these types of questions. The mediator for scurvy is vitamin C. Unfortunately, two expeditions of Robert Falcon Scott to Antarctica in 1903 and 1911 all suffered greatly from scurvy because of one thing: they didn't know the mediator. Sure, it was known that citrus fruits prevented scurvy but many believed it was a result of the fruit's acidity. The causal diagram was Citrus Fruits \rightarrow Acidity \rightarrow Scurvy. From this point of view, Coca-Cola would work! Upon hearing that despite taking lime juice, many explorers became ill, the medical community was shocked and confused. It wasn't until 1930 that Albert Szent-Gyorgi discovered that it was ascorbic acid or vitamin C that was the particular nutrient that prevented scurvy, changing the causal path to Citrus Fruits \rightarrow Vitamin C \rightarrow Scurvy.

Causal Relationships and Human Thought is the Key for Artificial Intelligence

Causal relationships are an important part of human discovery. If we never ask "Why?" then we never make new discoveries! But can we teach machines to understand causes and effects too? To answer this question, we'll need to take a look at causal models and big data. Today, we have more raw data than ever before as our world moves towards an online platform. "For example, in 2014 Facebook reportedly was warehousing 300 petabytes of data about its 2 billion active users. They had data about the games people play, the products they buy, the names of their Facebook friends, and of course, all their cat videos!"

But just as our online data is growing, so is our data in science. For example, the 1000 Genomes Project collected two hundred terabytes of information in what it calls "the largest public catalog of human variation and genotype data." But how do we extract meaning from all these numbers, bits, and pixels? Well, while the data seems immense, the questions are simple: "Is there a gene that causes lung cancer? What kinds of solar systems are likely to harbor Earth-like planets? What factors are causing the population of our favorite fish to decrease, and what can we do about it?" Many believe the answers to these questions will be found in data, but as you now know, these are all causal questions that cannot be answered on data alone.

Therefore, Big Data and causal inference must work together to determine the answers. The first step is to draw causal diagrams as we saw in the previous chapter. Once a diagram is drawn, it becomes possible to create a mathematical formula that demonstrates the relationship existing between correlation and causation. In a causal diagram, we can clearly see all the known factors in one place. These factors are then linked together with arrows, demonstrating how one directly affects another. Once they are linked, it becomes possible to see which are mediators and which are confounders. For example, let's begin with the assumption that blood pressure is known to be a possible cause of a heart attack, and Drug *B* is supposed to reduce blood pressure. Researchers might begin by drawing a diagram with arrows linking the drug and blood pressure, lifespan and blood pressure, and the drug and lifespan.

As you may know, age affects both blood pressure and lifespan, regardless of the drug, and is linked to both factors with an arrow that points in two directions, which allows us to see that age is a confounder. From here, the diagram can then be expressed in a formula. Because we have turned the causal relationship into a step-by-step logical process, we can then enter this cause-and-effect process into robots and computers. The formula could then be used to calculate both an answer and the statistical uncertainty in that answer. In other words, this would mean that computers would finally be able to ask *why?*

Final Summary

"Correlation is not causation" has long been accepted in the scientific community, and for good reason! Sure, the rooster crowing can be correlated with the sunrise, but it doesn't cause the sunrise. The problem, however, is that there is a large misunderstanding about what causation is. Today, we understand that causation is formed like a ladder with three rungs. As you climb the ladder, your questions become more complex and require a better understanding of causal relationships. Additionally, through the proper methodology, it is possible to determine when a correlation implies causation. Even more, this method could then be programmed into computers, allowing them to answer causal questions. Computers being able to answer why is the key to artificial intelligence and has the potential to open up a world of possibilities for scientific discoveries and advancements.



Go to QuickRead.com/App now to download our app and get access to thousands of free book summaries as both text and audiobooks.

Get the key insights of non-fiction books in minutes instead of hours. Listen to our free audiobooks while you workout or on your commute to work.



