WWW.RATIONALMINDSHOW.COM

guide to critical thinking

BY PIETRO BOSELLI, PHD

The modern world is the result not of individual intelligence, but of collective intelligence, thousands of years of accumulated knowledge, science, and complex institutions. As a result, the world around us is far too broad and nuanced for our intuition alone to grasp it, and so we must rely on tenets of reason, not only for understanding, but also for finding solutions to the challenges we all face.

Thinking rationally and statistically rather than emotionally and anecdotally is the more humanistic approach because it ensures real progress. Using reason and statistics does not mean reducing people to numbers. It means caring about the truth and about solutions.

The following concepts and principles will help construct a more critical approach to observing the world, to shaping our knowledge and opinions, and to making more accurate predictions.

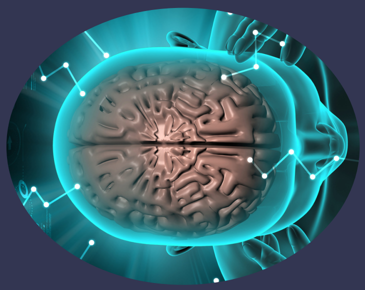

intuitive vs rational thinking

Intuitive thinking is the type of thinking that comes to us naturally and effortlessly. For example, we can recognize instantly from someone's face if they're happy or angry. Or if we see a marble rolling toward the edge of a table, we can predict it is going to fall off just by looking at it, without doing any calculations. If you are reading these words, you can do it effortlessly, because you are well practiced at reading. Rational thinking, on the other hand is the type of thinking that requires effort and concentration.

We are engaged in rational thinking when are performing difficult calculations in our mind, validating a logical argument, or even learning something new and complicated. For instance, when we learned how to read, it took practice and focus.

Intuitive thinking is fast, rational thinking is slow. Intuitive thinking is also known as common sense, and it is not just for humans - all animals have it as well. Intuitive thinking evolved to allow us to act quickly in nature, such as in fight or flight situations.

Intuitive thinking works via heuristics, which are subconscious mental shortcuts that our brain has evolved to give us quick answers to almost everything. And this works well in simple everyday situations, and for animals in the wild. But when trying to answer more complex questions, relying solely on intuition can - and often will - lead to logical fallacies and biases.

Unfortunately, these are often difficult to detect in ourselves and others because we all identify with our rational self. We think that all our conclusions have been reached through reason. But very often that is not the case. Quite often your intuition provides an answer and then you rationalize to justify it.

This sounds like a big claim, and I will make it even bigger by adding that not only are we all prone to logical fallacies and biases without realizing it, but also that the further away we are from the truth, the more confident we can be in our knowledge. But all this is best explained by explaining the concepts behind it with a few examples of how we can be inadvertently biased.

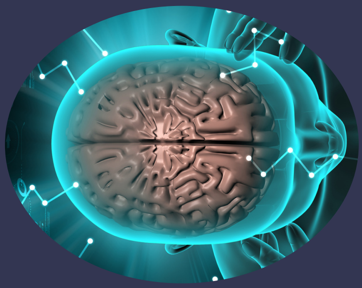

intuitive vs rational thinking

Consider the description of Linda.

Linda is a bright woman, 30 years old, single, with a major in philosophy.
As a student she cared about human rights and social justice.

Which of the following statements is the most probable?

Linda is a bank teller.
Linda is a bank teller and an outspoken feminist.

Most people will judge the second statement as the most probable, because it makes for a nice coherent story. Linda does not seem to fit the job of a bank teller but adding that she is a feminist seems to be more representative of her character. This is an example of thinking by representativeness. However, this choice contradicts logic.

The first statement is the most probable one, because the group of feminist bank tellers is wholly contained in the group of all bank tellers. If you bet on the first option, you would win if Linda was a bank teller, either feminist or not. With the second option you would win only if she was feminist. A more restrictive scenario hence a less probable one.

availability heuristics

Are there more atoms on Earth, or combinations of a deck of 52 cards?

What do you think?

Most people will answer the former. And this is because in our mind, we think of atoms as being incredibly small and numerous, so even if there are thousands of combinations possible in a deck of 52, surely there are more seconds since the beginning of the universe. It feels “right” to say this. The fact is that there are actually many times more combinations of a deck of cards than atoms on Earth. In fact, there are billions more combinations of a deck of cards than there are seconds since the Big Bang. But because your mind more easily summons the enormity of the number of atoms on earth, that information is more available, and so that becomes your answer.

The availability heuristic is a shortcut our mind uses to tell us the likelihood of something. If something comes easily to mind, if it’s more “available” to our mind, we intuitively expect it to be more likely. And this mechanism made perfect sense when we lived in the wild and were only aware of our immediate surroundings. If you saw many predators in a particular forest, you would easily remember that, as your life directly depended on it. The problem with the availability heuristic nowadays is that when we are asked about the frequency of something, we do not report the actual frequency. Instead, we report the ease with which we can think of examples of that class.

In the modern and complex world of today, our perceptions are easily distorted, chiefly by media, advertising, and organized religion, working together with various political factions to assert control by oversimplifying information in a way designed to leverage your mind to their desired outcome. For instance, the news reports only a non-random sample of what is happening in the world, focusing on the negative and the tragic. Those stories are carefully selected for maximum impact. Therefore, we think that violence and disasters are way more likely than they actually are. Frequent, dramatic, and very recent events are more available in our mind, as the memory is still fresh. Because the plane crash we have seen on the news is dramatic and easy to picture, we are more scared of flying than driving, even if driving is statistically more dangerous.

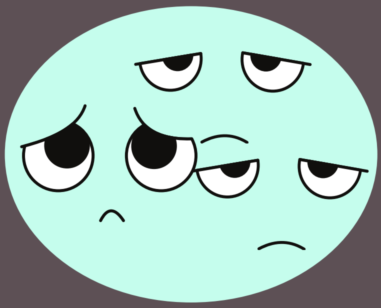

affect heuristics

We have seen how the availability heuristic is a shortcut our brain uses to judge something based on how easily it comes to mind. The affect heuristic is slightly different: it is a shortcut we subconsciously use to judge something based on our current emotions or the emotions evoked by that particular subject, without further analysis. When we are asked what we think about something, we tend to actually report what we subconsciously feel about it, and we think that is our rational answer.

For example, people who smile are judged more favorably independently of what we know about them. Experiments show that if I describe someone to you and I show you a smiling picture of them, you will like them more than if they weren't smiling. This is because we subconsciously feel better about a smile, even if we know smiles can easily be fake.

The most important consequence of the affect heuristic in our understanding of the world is that it influences how we judge an activity, or a technology, based on how we feel about it. This is especially true when the judgement is made in a hurry without actually investigating or evaluating in depth. In particular, if we like an activity or a technology, or if it is associated with a positive feeling, we tend to overestimate its benefits and underestimate its risks.

For instance, if you like smoking you will underestimate its risks. If you dislike social media, you will overestimate its risks. Our like or dislike of something is usually exacerbated by another bias known as confirmation bias. When we search for evidence, we tend to accept what confirms our presumed knowledge and tend to dismiss information which challenges it.

In this way, our likes and dislikes determine our beliefs about the world. Our political preference determines the arguments we find compelling. If we like a public policy, we likely believe its benefits are larger than they actually are, we neglect its risks, and we only search for evidence which confirms our belief.

It is important to be aware of these cognitive patterns in order to be more self-critical and objective in our analyses. A side effect of working from facts rather than feelings is that it can also cause us to be less politically polarized.

overconfidence

A rational explanation of anything needs complete and reliable information. The problem is that our intuitive thinking does not need all the information to create a story, as long as the story is coherent.

A coherent story is immediately and subconsciously accepted as true, and the more easily the story is formed, the more confident we are in it. Our intuitive system automatically ignores not only the fact that we might be missing parts of the story, but also the fact that some of the information might be inaccurate. The most common example of this is when we form an opinion about someone we just met: we are missing a big part of the picture and yet we cannot help forming an opinion about them.

Paradoxically, the less information we have, the easier it is to form a coherent story, because there are fewer pieces to the puzzle. This means that the less we know, the more confident we are in our knowledge.

In a famous series of experiments, legal cases were presented to different groups of people. One group was presented only the version of the plaintiff, one group only the version of the defendant, and one group heard both versions, like a jury would.

Everybody knew the setup of the experiment, and yet people who heard only one version of the story were more confident in their opinion that either the plaintiff or the defendant were right. People who heard both versions were not as confident, as reaching a definitive conclusion became harder.

"the less we know, the more confident we are in our knowledge."

overconfidence

Another feature of overconfidence is that people who just acquired some information about a subject suddenly feel happy with a simple story that is coherent in their mind. These people will be more confident than experts, who actually know that the story is way more complex than what it seems. The more skilled one becomes at something, the more they can recognize the limitations of their ability in that field.

This was presented in a famous paper by Kruger and Dunning, where they showed experimentally that people who scored in the bottom 20% on skills such as grammar and logical thinking consistently rate themselves above average. This is because they lack the very skills required to be self-aware of their lack of skills.

Put another way, incompetence cannot recognize incompetence. People with this kind of self-assured incompetence are easy pickings for those wishing to gain power through conspiracy theories, or prey upon prejudices, biases, and misinformation.

Meanwhile, people who were actually skilled rated themselves lower than they actually were, for two reasons: First, because they assume other people are as competent as they are, and second, because they are knowledgeable enough to recognize the difficulty in knowing all there is to know.

One thing to take away from this is that it's generally wiser to avoid the urge to jump to conclusions and resist the temptation of immediately reacting to new information. You should leave some time for more information to surface. Then deploy attention and rational thinking before expressing a judgement or opinion. A good example of this problem is breaking news: reporters tend to focus on the incomplete information currently available and purposely jump to a conclusion. We must avoid falling into the temptation of assimilating new information by immediately forming a strong opinion. Often only one tiny missing piece of information can lead to a very different conclusion.

correlation vs causation

A correlation is a measure of the relationship between two different variable quantities, which we simply call variables.

Height and weight of people are correlated. This means there is some sort of relationship between the two variables. On average, people who are 6'1" are heavier than people who are 6'. But not everyone who is taller is also heavier. You could also phrase it by saying that taller people are heavier on average, but that height is not a perfect predictor of weight.

Some things are not correlated at all. For example, your preference in music tells us nothing about your height. The correlation is zero. If we were to do a survey asking adult people their music preference and also their height, we would observe no relationship between these variables and could make no inference from the data.

Now that we understand correlation, a breakthrough in our understanding of the world will come by understanding and fully grasping the difference between causation and correlation.

When I was younger, my Grandma observed that almost all basketball players are tall. She insisted that I play basketball from a young age, to ensure that I would grow up to be tall.

Where was she wrong?

She committed one of the most common logical fallacies: inferring causation from correlation. She thought basketball "caused" people to be tall. She observed a correct correlation between being a basketball player and being tall but attributed a wrong causation.

It is true that nearly all basketball players are tall, but that does not mean that basketball causes you to be tall. Being tall gives you higher probability of becoming successful at basketball, and this is why basketball players are tall.

So WHY did my Grandma get it wrong?

correlation vs causation

It is because of our intuitive thinking. As we have seen, our mind needs to create a coherent story, and it needs to find a cause for things. It is an irresistible and automatic procedure.

In a well-known and now infamous study, it was shown that women who undertook hormone replacement therapy were on average less prone to heart disease. The conclusion made in the study upon observing this correlation was that hormone replacement therapy was the cause for lower risk of heart disease.

A later control study revealed that it was not the hormone replacement therapy that reduced heart disease. Women who could afford the expensive hormone replacement therapy were from a higher socioeconomic class and had on average better diets and exercise regimens, and therefore were less prone to heart disease. In the original study, a wrong causation was inferred from the correlation.

Unfortunately, nowadays we are bombarded with such studies with attention-grabbing but unsubstantiated headlines. They exploit the fact that the intuitive mind thinks causally. Thus, we are presented with many correlations which we wrongly conclude must be causal, because we are missing a part of the picture or a control study.

"we are bombarded with attention-grabbing but unsubstantiated headlines."

sample size

Statistics is the science of collection and interpretation of data.

Suppose we want to know how many people in the US currently support the president. We cannot ask everybody individually, so we choose a random sample of the population. The sample not only must be random, but also must be large enough in order to be representative.

It could be intuitive that larger samples are more accurate. What is less intuitive is the same fact, presented in a different way: small samples yield extreme results more often than large samples.

Consider a box containing green and red marbles in equal amounts.
Our volunteers cannot see the marbles and draw them “blindly”.

Jack draws 4 marbles at a time, observes how many red and how many green he selected, and then puts the marbles back in the box. He only writes down when he gets an extreme result of either all green or all red marbles. Jill does the same but draws 7 marbles at a time instead of 4, and also writes down every time she observes an extreme result, all green or all red marbles.

Who is going to see an extreme result more often, Jack or Jill?

I think we all intuitively understand that Jack is more likely to see extreme results with his small samples, which is to say seeing 4 marbles all of the same color is easier than seeing 7 all of the same color. This seems intuitive enough.

But we do not intuitively know how much more likely it is for Jack to see extreme results. If Jack draws 4 marbles at a time and Jill draws 7 at a time, Jack is actually 8 times more likely to see an extreme result. This sounds surprising to our intuitive mind, but extremely useful to know.

sample size

A good and famous example is provided by two researchers, Wainer and Zwerling. They considered a study in which the incidence of kidney cancer, a rare form of cancer, was measured in all of the over 3,000 counties of the US. It was observed that the counties with the lowest incidence of cancer are mostly rural and sparsely populated in the Midwest and South.

The first thing one might think in order to explain this observation is that perhaps people in the countryside have less pollution, the air is cleaner, and they eat fresher produce, so they are less prone to cancer.

By further inspection of this study, they then went on to observe that the counties with the highest incidence of kidney cancer are mostly rural and sparsely populated in the Midwest and South. So exactly the same regions!

Tongue in cheek, one might try to explain this by saying that the poorer counties are where people eat unhealthily, they smoke, drink, and have less access to healthcare. And this is why people in rural counties are more prone to cancer.

So, which one is true then?

The only true explanation is neither of them. The only valid explanation is that rural counties have small populations, which are more likely to show extreme results of both high and low incidence of cancer. This is the law of small samples. And that is all there is to know. All other causal explanations - pollution, fresh food, smoking, clean air, etc. - are just created intuitively and erroneously because our mind cannot help finding a cause.

"Small samples yield extreme results more often than large samples."

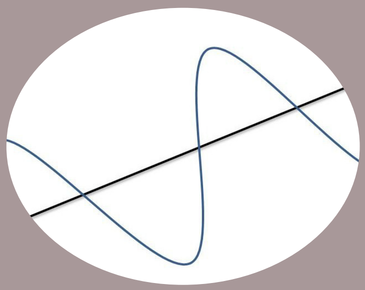

regression to the mean

If a random variable is extreme at its first measurement, it will be closer to the average (or mean) on successive measurements. This effect is called regression to the mean.

If today is the hottest summer day ever recorded, it is likely that in the following week the temperature will be lower, closer to the average summer day. Here the random variable being measured is “temperature in summer”, and the extreme is the highest temperature on record, which then regresses in the direction of the average temperature or mean temperature. A similar thing could be said about the coldest winter day, which is likely to be followed by a slightly milder day, or closer to the mean.

Another example: If you are having the best meal you have ever had, it is likely that tonight for dinner your meal will not match that - it will be just your regular meal.

An extreme is likely to be followed by something less extreme and more average.

In the 70’s cognitive psychologist Daniel Kahneman was called to help the Israeli army with Air Force pilot training techniques. The main military instructor reported that praising a pilot’s good performance was counterproductive because he would do worse the following time. The instructor then said that punishing bad performance always resulted in an improvement the following time. What the instructor had observed was regression to the mean, but he gave it a causal interpretation: according to him, improvement was caused by the punishment, and deterioration was probably due to complacency caused by the praise.

"An extreme is likely to be followed by something less extreme and more average."

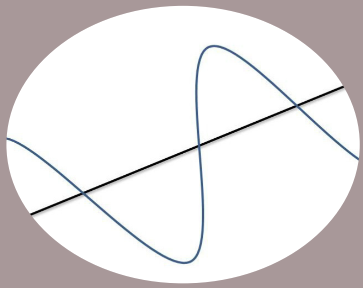

regression to the mean

In reality the pilot doing exceptionally well on one occasion was largely due to luck on that day (maybe favorable conditions, good mood, etc.). Because that was the best he had ever done, any change in circumstances meant he was more likely do worse after that. A pilot doing his worst on one day was also due to luck (maybe unfavorable conditions, tiredness on that day, etc.). From there he could hardly do worse, so the next time, he was bound to do better whether he received the punishment or not.

Our mind cannot help giving a causal explanation, even if the effect is merely an artifact of the statistics. This is detrimental in many studies that are published and ignore correct statistical inference.

Suppose someone is claiming that taking a particular pill for a month can cure unhappiness. So, they take the 1% unhappiest people, give them a pill for a month, and then ask them again how happy they are now. Many people will report they are happier, and the average happiness of this small group will now be closer to the mean happiness of the general population. But we know that very unhappy people are an extreme group and likely to move closer to the mean with time anyway. This is because everybody's circumstances change and after a few months, the 1% unhappiest will mostly be a completely different group of people. So how do we know if the pill had any effect at all?

A proper clinical trial requires a control group. We take the very unhappy people and split them into two groups. We give one group the actual pill and the other group a placebo pill, which looks exactly the same but has no active ingredient. After a month, we test the happiness of the two groups. We know that on average both groups will be happier due to regression. But if the group taking the pill is significantly happier on average than the placebo group, then we know the pill works. Otherwise, we discard the claim.

Any trial without a control group is most likely to not be invalid. This is because we can observe false results due not only to regression, but also to sample size effects, or wrong attributions of causation once a correlation is observed.

base rates

What is a base rate?

Let's start with a simple, textbook definition.

A base rate is the general probability of something being true, or the percentage of a population with certain characteristics.

For example, if 5 in 100 cars are red, the base rate for red cars is 5%. If we were to pick at random a car from all cars in the street, we know that the probability of it being red is 5%.

It is also worth offering a simple definition of probability: the likelihood of something happening where 0% means no chance and 100% means certainty.

Next example of base rate: If 1 in 1,000 people show side effects for a new drug, the base rate for the drug's side effects is 0.1%. If we give the drug to a person at random, the probability of them showing side effects is 0.1%.

Why do we need to know about base rates? Because ignoring them can lead to seriously wrong judgements of probability.

Assume that 1 in 1,000 drivers are currently drunk, so the base rate for drunk drivers is 0.1%. Consider a breathalyzer which never fails to detect a drunk person but displays a false drunkenness in 5% of the cases when the driver is sober. These are called "false positives".

Assume that the police are out on the street and stop one driver at random and perform the breathalyzer test. The test shows that the driver is drunk. What is the probability that the driver is actually drunk? Most people will intuitively say a high percentage up to 95%, thinking about the 5% false positives. But, the actual probability of the driver being drunk is only 2%!

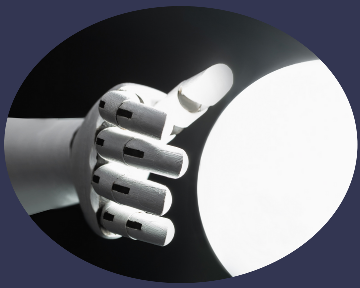

making better predictions

First, when we talk about probability, this could be interpreted as the frequency of an outcome when an action is repeated many times. For example, if I toss a coin only 3 times, I will not be too surprised if I get tails all three times. But if I toss it an infinite number of times, I will get tails 50% of the times, which means the probability of getting tails is 50%. So here, probability is the frequency of an outcome when repeating an action many times.

There is also another interpretation of probability, which is the degree of credence we assign to something. For example, if I am about to toss a coin, I can say the likelihood of it coming out tails is 50%. In other words, 50% is the level of credence I assign to the outcome “tails”.

Why is it useful to have this different definition?

Under this interpretation of probability, we can assign credence or degrees of belief to future events which will happen only once. For example, I could say that the Los Angeles Lakers have a 70% chance of winning the next game. Clearly, we cannot repeat the game an infinite number of times (like we do with the coin toss) to see if it is true that 70% of the time the Lakers will win.

Instead, the probability is assigned based on prior knowledge (how often the Lakers have been winning this season) and then updated as new evidence arises (for example, is the team in good shape today). This is called a Bayesian approach, as opposed to a frequentist approach, in which probability can only be measured for events that can be repeated many times.

If a hypothetical girl named Anna was able to read and write at the precocious age of 4, what is her high school grade in English most likely to be?

It is tempting to say she most probably will have an A in English. However, there are so many other factors that influence how well one does in school, and her precociousness probably has a smaller influence.

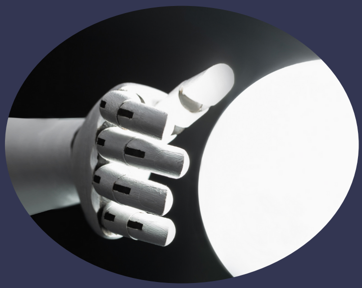

making better predictions

A better estimate would be made by taking a base rate: what is the average English grade at her school, for example a B, and then integrate the situational evidence (she was precocious so probably academically talented) by adjusting from the base rate going up. So perhaps a B+ is a better estimate of her grade.

This is a bit counterintuitive because our mind likes to think by representativeness. But we have seen how that can fail when thinking about probabilities.

As individuals we are all prone to biases and logical fallacies, but it does not mean we are irrational. Humans have discovered these logical fallacies precisely because we are capable of rational thinking, and we can define a baseline of rationality against which we can measure irrationality.

Humanity has amassed so much knowledge, created science and institutions which are tested against reality, and they work, and help us successfully navigate an increasingly complex world. By looking through the lens of reason, we can see the progress of humanity.

It is enough to look at the data and long-term positive trends, without falling prey to irrational fears and pessimism, to believe in the future of innovation while fostering constructive debate, tolerance, and pluralism.

episodes

www.rationalmind.show
www.youtube.com/rationalmind

ENERGY

June 28

Solar, wind, oil, nuclear - what is the real future of clean safe energy?

EMERGENCE

July 26

What can we learn from the swarm behavior that emerges from individual actions?

AGRICULTURE

July 5

Organic, conventional, GMO - what is the answer to feed a growing population?

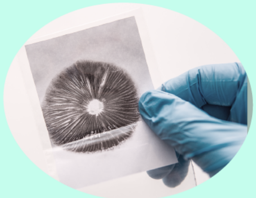

PSYCHEDELICS

August 2

Do psychedelic-derived treatments hold the key to treating debilitating medical conditions?

DECENTRALISATION

July 12

What are the impacts of decentralised finance, government, and communication?

SPACE

August 9

What is actually happening in space, and what does it mean for humanity?

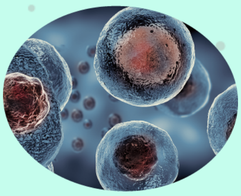

LONGEVITY

July 19

Are we on the cusp of revolutionary life extension, and what does it mean if we are?

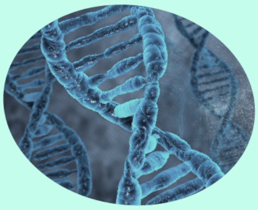

GENE EDITING

August 16

How close are we to personally curated cures for diseases?