DON A. MOORE, PHD

PERFECTLY CONFIDENT

HOW TO CALIBRATE YOUR DECISIONS WISELY

FIRST EDITION

Library of Congress Cataloging-in-Publication Data has been applied for.

ISBN 978-0-06-288775-7

20 21 22 23 24 LSC 10 9 8 7 6 5 4 3 2 1

CHAPTER 1

What Is Confidence?

CONFIDENCE INTERVALS

To help you calibrate your confidence, I would like to invite you to play a little game with me. The table that follows lists ten quantities of which you are uncertain. For each one, please identify your answer, with a 90 percent confidence interval. A confidence interval consists of two numbers, one below your best guess and one above your best guess. That range should be wide enough that you are 90 percent sure the right answer is somewhere between. Obviously, the surer you are, the narrower you can make the confidence interval. If the question was about the date of your own birth, you could make the confidence interval very precise indeed. The less certain you become, the wider you should make the confidence interval. Your challenge is to calibrate your intervals so that you are 90 percent sure the right answer is inside it.

One way to think about it is like this: Set the lower bound so low that there is only a 5 percent chance the truth is below it (and a 95 percent chance the truth is above it). Then set the upper bound so high that there is only a 5 percent chance the truth is above it (and a 95 percent chance the truth is below it). With only a 10 percent chance the truth is outside the range, you should have a 90 percent confidence interval.

Without consulting any reference materials or other people, please estimate 90 percent confidence intervals for the ten quantities in the table.

	Lower Bound	Upper Bound
1. World population, according to the US Census Bureau on September 12, 2019		
2. Year in which Orville Wright took the world's first powered heavier-than-air flight		
3. Hourly wage that Steve Jobs paid designer Dean Hovey for creating Apple's mouse		
4. Maximum depth (below sea level) of the Mariana Trench in the Pacific Ocean		
5. Total revenues of the Tesla corporation in 2018 (according to its annual report)		
6. Year in which Daniel Kahneman won the Nobel Prize in Economics		
7. Amount of money that Google paid to buy YouTube in 2007		
8. LeBron James's average points per game in his NBA career, as of July 2019		
9. Year in which William James first taught a psychology class at Harvard University		
10. Year in which Maya Angelou was awarded the Presidential Medal of Freedom		

Did you really take the time to answer the questions? Please do. It gives you some skin in this game. It improves your ability to apply the insights offered here to yourself and to your own decisions.

For how many of these questions should the right answer have landed between your lower and upper bounds? Well, if you have calibrated your confidence correctly, then each one should have a 90 percent chance of hitting. Out of ten items, nine should be inside the confidence interval. Read on to discover the right answers. How many did you get?

	Correct answer
1. World population, according to the US Census Bureau on September 12, 2019	7.598 billion
2. Year in which Orville Wright took the world's first powered heavier-than-air flight	1903
3. Hourly wage that Steve Jobs paid designer Dean Hovey for creating Apple's mouse	$35
4. Maximum depth (below sea level) of the Mariana Trench in the Pacific Ocean	36,070 ft.
5. Total revenues of the Tesla corporation in 2018 (according to its annual report)	$21.46 billion
6. Year in which Daniel Kahneman won the Nobel Prize in Economics	2002
7. Amount that Google paid to buy YouTube in 2007	$1.65 billion
8. LeBron James's average points per game in his NBA career, as of July 2019	27.2
9. Year in which William James first taught a psychology class at Harvard University	1873
10. Year in which Maya Angelou was awarded the Presidential Medal of Freedom	2010

CHAPTER 2

How Might I Be Wrong?

How high a standard of rigor and logic do you hold yourself to? How can you avoid falling for false or misleading claims? By way of trying to assess the calibration of your own beliefs, I would like to ask you for a forecast about the end of the world, or at least some people's part of it. Of all those who died last year, what percentage of global deaths were due to each of the following causes? Please make these estimates without consulting any reference materials or other people.

I will provide the answers in short order. But first, let's think about another way to help you calibrate confidence in your beliefs. Chapter 1 asked you to assess your confidence in your beliefs by specifying 90 percent confidence intervals for your estimates of

Cause of Death	Percent of All Deaths
Road injuries, including vehicular collisions	
Accidental falls	
All other unintentional injuries (drowning, fire, poisoning, etc.)	
Self-harm, including suicide	
Interpersonal violence, including murders	
All other intentional injuries, including genocides and wars	

uncertain quantities. Now we're going to use another measure. Without consulting any reference materials or other people, please make your best guess for each of the ten quantities in the following chart.

	Best Guess	Confidence
1. Number of (full-powered) radio stations in the Christian Family Radio network as of July 2018		
2. Number of deaths worldwide due to motor vehicle accidents (in millions)		
3. Net worth of Jeff Bezos, as of July 2019, according to *Forbes* magazine		
4. Year in which the English king Charles I was beheaded		
5. Number of deaths in the September 11 attacks		
6. Total revenues of Amazon corporation in 2018		
7. Age at which Oliver Cromwell, Lord Protector of England, died		
8. Number of Jewish Zealots who died on Masada		
9. Assets under management by the hedge fund Bridgewater Associates		
10. Number of new saints canonized by Pope John Paul II		

Then indicate how confident you are in your answer by estimating the likelihood (0 to 100 percent) that your answer is close to (within 5 percent of) the right answer.

Did you really take the time to answer the questions above? Do it. C'mon! It'll be fun.

HYPOTHESIS TESTING

When you go into the world asking, "Is this hypothesis true?" you may be tempted to believe you are taking a neutral approach, but you are not. Simply the way you pose the question can influence the answer in subtle and surprising ways. It will be easier for you to think of evidence that allows you to answer yes. You will formulate questions that are more likely to generate affirmative answers. When you ask other people these questions, they will be more likely to respond in the affirmative or to provide you with evidence that supports your hypothesis. The question "Is this hypothesis false?" generates a different approach, a different line of thinking, different responses, and different conclusions. Ignorance of the way you bias your search for information leads you to be too confident in the biased conclusions that result.

Now let's compare how confident you said you were with your hit rate. How often were your guesses actually within 5 percent of the right answers?

To test your calibration, compare your average confidence across all ten questions with the frequency that your answers were actually within 5 percent of the truth. How did you do? If you're like most people, your confidence exceeded your accuracy. But you probably also did better than you did on the confidence interval task from chapter 1. As a rule, people show better calibration for probability estimates than for confidence intervals. There are a

	5% Below	Right Answer	5% Above
1. Number of (full-powered) radio stations in the Christian Family Radio network as of July 2019	47	49	51
2. Number of deaths worldwide due to motor vehicle accidents (in millions)	1.28	1.35	1.42
3. Net worth of Jeff Bezos, as of July 2019, according to *Forbes* magazine	$155 billion	$163 billion	$171 billion
4. Year in which the English king Charles I was beheaded	1566	1648	1730
5. Number of deaths in the September 11 attacks	2,846	2,996	3,146
6. Total revenue of Amazon corporation in 2018	$221 billion	$232.9 billion	$245 billion
7. Age at which Oliver Cromwell, Lord Protector of England, died	56	59	62
8. Number of Jewish Zealots who died on Masada	912	960	1,008
9. Assets under management by the hedge fund Bridgewater Associates	$152 billion	$160 billion	$168 billion
10. Number of new saints canonized by Pope John Paul II	459	483	507

couple of reasons for that. First, understanding the logic of confidence intervals depends on thinking about uncertainty as a probability distribution, which few people do naturally. Second, everyday life rarely requires us to specify confidence intervals, and so we get little practice or feedback using them.

Let me ask you for one more probability estimate, this one for your forecast of the percentage of all deaths this year worldwide that will be caused by injuries (as opposed to disease and

starvation). Please do not look back at your previous answers, but just estimate this one on its own:

Cause of Death	Percent of All Deaths
Injuries, both intentional and unintentional	

CHAPTER 3

What Is Possible?

FORECASTING

The superior alternative is asking about a probability distribution rather than a single point. In the case of weight guessing, that might look like this:

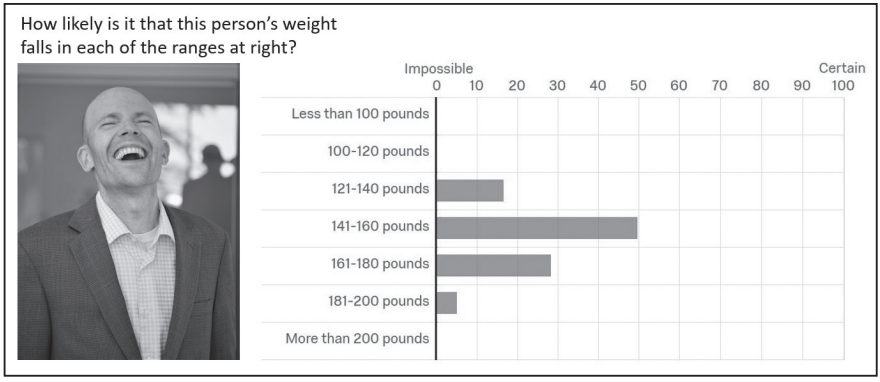

Histogram response scale

Asking people to complete a histogram like this forces them to broaden their thinking and consider the possibility that their best guess is wrong. When I ask the question this way, the one twenty-pound bin rated "most likely" to contain the right answer drops from 60 percent to 40 percent. In other words, when I ask them which twenty-pound bin is most likely, people tell me that, on average, they are 60 percent sure they picked the right one. But when I ask them to rate all the bins' probability of containing the right answer, no bin gets higher than 40 percent. Forty percent is still higher than it should be if hit rates average only 30 percent, but it's a whole lot closer to good calibration. Moreover, it produces a useful distribution that can help make decisions. For instance, if the question was about ibuprofen sales rather than body weight, knowing the distribution of possible sales figures would be quite helpful.

THINKING ABOUT UNCERTAINTY

The psychological weight accorded to a possible outcome is not the same as its objective probability. The relationship between the two is usefully illustrated in this subjective probability weighting function. In this graph, the x-axis shows true probability and the y-axis shows subjective psychologically weighted probability. The dotted line shows how a perfectly rational person would weight the probabilities. The solid curve illustrates the psychological reality:

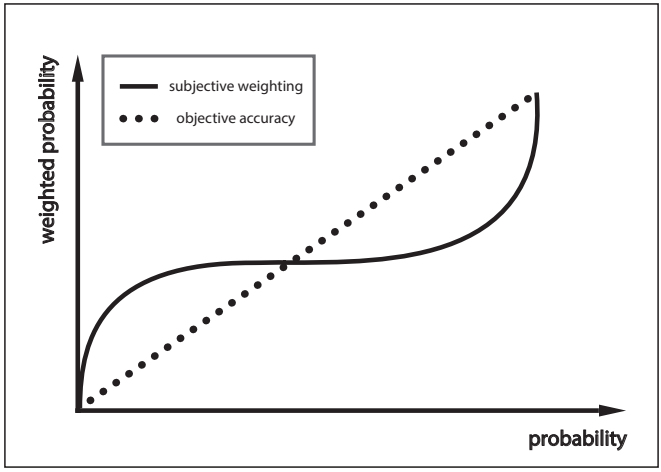

The probability weighting function

CALIBRATING YOUR CONFIDENCE

Plinko is a version of the Quincunx, a device useful for generating random distributions. The Quincunx, also known as the Galton Board, is a device invented by Sir Francis Galton to demonstrate the normal distribution generated by chance. Here is what it looks like:

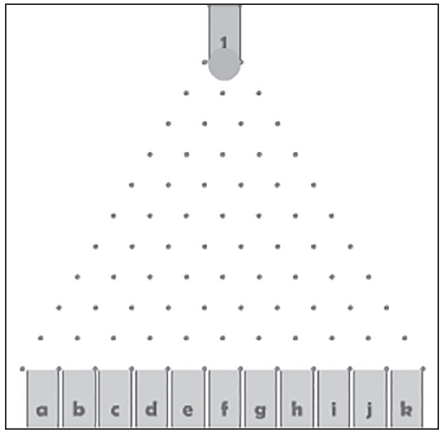

The Quincunx

Thinking in expected values can be useful in all sorts of other areas of life. Consider, for example, when you will complete a major project that you are working on right now. It could be the length of time it will take to finish an important report, finish a new building, or complete a new software product, or maybe the length of time until you get a major promotion. Really take a moment to identify a particular project. You may have already forecast when you will be done, and maybe you have told your forecast to other people, such as your boss or a customer. Now consider the possibility that it will take twice as long as you expect. How likely is that? It is also possible it will take half as long as you expect. If you estimate the probability of each of these possible completion times, you have just completed a histogram distribution. This simple process has created a histogram with four bins to which you can assign probabilities, as illustrated in the following table. *The very act of thinking through these possibilities will help you make a more accurate forecast.*

If, on January 1, you guessed that your project would probably be done in six months, then the procedure above would produce this binning arrangement:

Bin	Probability
Before April 1	
April 1 to June 30	
July 1 to December 31	
After December 31	

A histogram for forecasting the completion of a hypothetical major project.

However, there are probably more useful binning arrangements to consider than the one I threw out. Take some time and think through the possibilities. Maybe it makes more sense to have the bins represent months, or perhaps weeks. If the range of possibilities is narrow, then there may be few bins. Or maybe you want to consider as many as twenty or thirty bins. Once you have your bins, go through and specify the probabilities on each one. At first, don't worry if the probabilities don't sum exactly to 100 percent, just try to get them proportionately right. After you have that, you can go back and correct them so they sum to 100 percent. You may want to share your histogram distribution with others if it will help calibrate their expectations of you, providing them with more accurate and useful information. Naturally, you may want to share selectively. For instance, to an audience expecting a point prediction, you may want to share a later time point by which you are more than 90 percent confident you will be able to deliver.

But even more important than that is to share your forecast with your future self so that you can follow up and learn from it. Save your forecast in a place where you will be able to find it later. Then plan to hold yourself accountable for your forecast. Put a note on your calendar to review its accuracy down the line. Maybe you want to do that at the time of the first milestone, or maybe at the time of your best-guess forecast.

When my friend and colleague Julia Minson recently asked me to read and comment on a paper of hers, I told her that I would love to help but was unsure when I would be able to make the time. I could make a point prediction of a single date, but there were so many uncertainties about how long it would take that pretending to know with certainty was foolhardy. Instead, I gave her the following probability distribution:

Bin	Bin Probability	End Date	Cumulative Probability
1 day	0%	Within 1 day	0%
2–7 days	15%	Within 1 week	15%
8–14 days	25%	Within 2 weeks	40%
15–21 days	30%	Within 3 weeks	70%
22–28 days	20%	Within 4 weeks	90%
28–35 days	8%	Within 5 weeks	98%
Later	2%	Later	100%

A histogram for my forecast of when I would read my colleague's paper. The bin probability specifies the probability the outcome will land in that bin. The cumulative probability indicates the probability that it will happen before the end date.

By liberating myself from the false certainty of the point prediction, I enabled a more honest elaboration of the possibilities and became better able to communicate it to others, including my friend. That same information illustrated as a cumulative probability distribution looks like this:

Cumulative probability distribution

Thinking about uncertainty in terms of probability distributions can help you become better calibrated in your confidence. It demands that you think about the full range of possibilities and their likelihoods. And so, it places your confidence in the probability distribution, as opposed to a far more fallible best guess. If you had neglected to consider the chance that your invention would sell millions per year, pausing to reflect on that possibility would be worthwhile. Thinking about the future in terms of a probability distribution may not come naturally at first, but the effort pays off. When you grow more comfortable thinking this way, you can more clearly see the potential before you, improve your forecasts, and make better decisions. This provides a stronger foundation for well-founded expectations, whether you are trying to fly a plane, get your business off the ground, or play Plinko.

CHAPTER 8

Find the Middle Way

You may not be clinging to a narrow ledge thousands of feet up, but you undertake risky activities each and every day. Whether that is driving your car in heavy traffic, making high-stakes investment decisions, or navigating complex relationships, you need to be able to calibrate your confidence. Well-calibrated confidence is the map that can guide your life choices about what to undertake, how to direct your efforts, and what risks might get you killed.

This book has cataloged the situations in which each of us is likely to be either overconfident or underconfident. The following table summarizes some of those situations. It distinguishes among the three forms of confidence:

- Estimation quantifies how good you think you are, how likely you are to succeed, or how quickly you will get things done.
- Placement compares yourself with others.
- Precision assesses the accuracy of your beliefs or how sure you are that you are right.

Both overconfidence and underconfidence are common, each in different situations. The one notable exception is the distinct absence of underprecision. Research has failed to identify situations in which people are systematically underprecise—that is, insufficiently sure that their knowledge is correct. That is why I have so enthusiastically recommended that you consider why you might be wrong.

	Overconfidence	Underconfidence
Estimation	Wishful thinking (ch.4) Planning fallacy (ch.6)	Rumination and worry (ch.1) Risk exaggeration (ch.4)
Placement	Easy tasks (ch.1) Common events (ch.1) Better-than-average effects (ch.5) Moral superiority (ch.7)	Difficult tasks (ch.1) Rare events (ch.1) Impostor syndrome (ch.1)
Precision	90 percent confidence intervals (ch.1) Ideological certainty (ch.2) Religious zealotry (ch.2) Histogram analyses (ch.3) Equity trading (ch.7)	