6

Decisions, Decisions

IF YOU COULD KNOW HOW your decisions would turn out, decision making would be so easy! It is hard because you have to make decisions with imperfect information.

Suppose you are thinking of making a career move. You have a variety of next steps to consider:

- You could look for the same job you're doing now, though with some better attributes (compensation, location, mission of organization, etc.).
- You could try to move up the professional ladder at your current job.
- You could move to a similar organization at a higher position.
- You could switch careers altogether, starting by going back to school.

There are certainly more options. When you dig into them all, the array of choices seems endless. And you won't be able to try any of them out completely before you commit to one. Such is life.

How do you make sense of it all? The go-to framework for most people in situations like this is the **pro-con list**, where you *list* all the positive things that could happen if the decision was made (the *pros*), weighing them against the negative things that could happen (the *cons*).

"I don't believe in pressuring my children. When the time is right, they'll arrive at the default choice and go to law school."

While useful in some simple cases, this basic pro-con methodology has significant shortcomings. First, the list presumes there are only two options, when as you just saw there are usually many more. Second, it presents all pros and cons as if they had equal weight. Third, a pro-con list treats each item independently, whereas these factors are often interrelated. A fourth problem is that since the pros are often more obvious than the cons, this disparity can lead to a **grass-is-greener mentality**, causing you *mentally* to accentuate the positives (e.g., *greener grass*) and overlook the negatives.

As an example, in 2000, Gabriel finished school and began a career as an entrepreneur. Early on, at times, he considered switching to a career in venture capital, where he would fund and support companies instead of starting his own. When he initially made a pro-con list, this seemed like a great idea. There were many pros (the chance to work with founders changing the world, the potential for extremely high compensation, the opportunity to work on startups in a high-leverage way without the risk and stress of being the founder, etc.) and no obvious cons.

However, there were several cons that he just didn't fully appreciate or just didn't know about yet (the relentless socializing involved—not good for a major introvert—the burden of having to constantly say no to people, the difficulty of breaking into the field, the fact that much of your time is spent with struggling companies, etc.). While certainly a great career for some who get the opportunity, venture capital was not a good fit for Gabriel, even if he didn't realize it at first. With more time and experience, the full picture has become clear (the grass isn't greener, at least for him), and he has no plans to make that career change.

This anecdote is meant to illustrate that it is inherently difficult to create a complete pro-con list when your experience is limited. Other mental models in this chapter will help you approach situations like these with more objectivity and skepticism, so you can uncover the complete picture faster and make sense of what to do about it.

You've probably heard the phrase *If all you have is a hammer, everything looks like a nail*. This phrase is called **Maslow's hammer** and is derived from this longer passage by psychologist Abraham Maslow in his 1966 book *The Psychology of Science*:

> I remember seeing an elaborate and complicated automatic washing machine for automobiles that did a beautiful job of washing them. But it could do only that, and everything else that got into its clutches was treated as if it were an automobile to be washed. I suppose it is tempting, if the only tool you have is a hammer, to treat everything as if it were a nail.

The hammer of decision-making models is the pro-con list; useful in some instances, but not the optimal tool for every decision. Luckily, there are other decision-making models to help you efficiently discover and evaluate your options and their consequences across a variety of situations. As some decisions are complex and consequential, they demand more complicated mental models. In simpler cases, applying these sophisticated models would be overkill.

It is best, however, to be aware of the range of mental models available so that you can pick the right tool for any situation.

WEIGHING THE COSTS AND BENEFITS

One simple approach to improving the pro-con list is to add some numbers to it. Go through each of your pros and cons and put a score of −10 to 10 next to it, indicating how much that item is worth to you relative to the others (negatives for cons and positives for pros). When considering a new job, perhaps location is much more important to you than a salary adjustment? If so, location would get a higher score.

Scoring in this way helps you overcome some of the pro-con list deficiencies. Now each item isn't treated equally anymore. You can also group multiple items together into one score if they are interrelated. And you can now more easily compare multiple options: simply add up all the pros and cons for each option (e.g., job offers) and see which one comes out on top.

This method is a simple type of **cost-benefit analysis**, a natural extension of the pro-con list that works well as a drop-in replacement in many situations. This powerful mental model helps you more systematically and quantitatively *analyze* the *benefits* (pros) and *costs* (cons) across an array of options.

For simple situations, the scoring approach just outlined works well. In the rest of this section, we explain how to think about cost-benefit analysis in more complicated situations, introducing a few other mental models you will need to do so. Even if you don't use sophisticated cost-benefit analysis yourself, you will want to understand how it works because this method is often used by governments and organizations to make critical decisions. (Math warning: because numbers are involved, there is a bit of arithmetic needed.)

The first change when you get more sophisticated is that instead of putting relative scores next to each item (e.g., −10 to 10), you start by putting explicit dollar values next to them (e.g., −$100, +$5,000, etc.). Now when you add up the costs and benefits, you will end up with an estimate of that option's worth to you in dollars.

For example, when considering the option of buying a house, you would start by writing down what you would need to pay out now (your down payment, inspection, closing costs), what you would expect to pay over time (your mortgage payments, home improvements, taxes . . . the list goes on), and what you expect to get back when you sell the house. When you add those together, you can estimate how much you stand to gain (or lose) in the long term.

As with pro-con lists, it is still hard to account for every cost and benefit in a cost-benefit analysis. However, it is important to note that this model works well only if you are thorough, because you will use that final number to make decisions. One useful tactic is to talk to people who have made similar decisions and ask them to point out costs or benefits that you may have missed. For instance, by talking to other homeowners, you might learn about maintenance costs you didn't fully consider (like how often things break, removing dead trees, etc.). Longtime homeowners can easily rattle off this hidden litany of costs (said with experience!).

When writing down costs and benefits, you will find that some are intangible. Continuing the house example, when you buy a house, you might have some anxiety around keeping it up to date, and that anxiety can be an additional "cost." Conversely, there may be intangible benefits to owning a home, such as not having to deal with a landlord. In a cost-benefit analysis, when faced with intangibles like these, you still want to assign dollar values to them, even if they are just rough estimates of how much they are worth to you. Doing so will help you create a fair quantitative comparison between the courses of action you are considering.

Writing down dollar values for intangible costs and benefits can seem strange—how do you know what it's worth to you to not have to deal with a landlord? But if you think about it, this is no different than scoring a pro-con list. In the scoring method, if the extra amount you'd have to pay monthly rated a −10 (out of 10) and landlord avoidance rated a +1 (out of 10), then you have a quick way to start an estimate: just take the extra payment amount and divide it by 10. Say the excess monthly payments are expected to be $1,000 per month; then you could estimate it is worth $100 per month to avoid a landlord. Of course, you can pick any numbers that make sense to you.

You can get hung up here because it can feel arbitrary to write down specific values for things that you don't know exactly. However, you should know that doing so truly helps your analysis. The reason is that you really do have some sense for how valuable things are and putting that (even inexact) sense into your analysis will improve your results. And, as we will see in a moment, there is a method for testing how much these values are influencing your results.

So far, you've moved from scoring to dollar values. Next, you graduate to a spreadsheet! Instead of a column of costs and a column of benefits, now you want to arrange the costs and benefits on a timeline. Give each item its own row, and each column in the timeline will now list the cost or benefit created by that item in a given year. So, the first column holds all the costs and benefits you expect this year (in year 0), the next column in year 1, then year 2, and so on. The row for a $2,000-per-month mortgage payment would look like −$24,000, −$24,000, −$24,000, for as many years as the life of the mortgage.

The reason it is important to lay out the costs and benefits over time in this manner (in addition to increased clarity) is that benefits you get today are worth more than those same benefits later. There are three reasons for this that are important to appreciate, so please excuse the tangent; back to the cost-benefit analysis in a minute.

First, if you receive money (or another benefit) today, you can use it immediately. This opens up opportunities for you that you wouldn't otherwise have. For instance, you could invest those funds right now and be receiving a return on that money via a different investment, or you could use the funds for additional education, investing in yourself. (See *opportunity cost of capital* in Chapter 3.)

Second, most economies have some level of **inflation**, which describes how, over time, prices tend to increase, or *inflate*. As a result, your money will have less purchasing power in the future than it does today. When we were younger, the standard price for a slice of pizza was one dollar; now a slice will run you upward of three dollars! That's inflation.

Because of inflation, if you get one hundred dollars ten years from now, you won't be able to buy as much as you could if you had the same amount of money today. Consequently, you don't want to regard an amount of money in ten years as the equivalent amount of money available today.

Third, the future is uncertain, and so there is risk that your predicted benefits and costs will change. For instance, benefits that depend on currencies, stock markets, and interest rates will fluctuate in value, and the further you go into the future, the harder they are to predict.

Now back to cost-benefit analysis. As you recall, you have a spreadsheet that lays out current and future costs and benefits across time. To account for the differences in value between current and future benefits, you use a mental model we introduced back in Chapter 3: the *discount rate*. You simply *discount* future benefits (and costs) when comparing them to today. Let's walk through an example to show you how it works.

Cost-benefit analysis is arguably most straightforward with simple investments, so let's use one. Bonds are a common investment option, which operate like a loan: you invest (loan) money today and expect to get back more money in the future when the bond matures (is due). Suppose you invest $50,000 in a bond, which you expect to return $100,000 in ten years. Feel free to make a spreadsheet and follow along.

Cost-Benefit Analysis

	Time line						
	Year 0	Year 1	Year 2	Year 3	Year 4	...	Year 10
Costs	$(50,000)	—	—	—	—	...	—
Benefits	—	—	—	—	—	...	$100,000
Discounted (6%)	$55,859					...	

Cost-Benefit Analysis

Net benefit	\$5,839					...	

The only cost today (year 0) is \$50,000, to purchase the bond. The only benefit in the future (year 10) is \$100,000, what you get back when the bond matures. However, as noted, that benefit is not actually worth \$100,000 in today's dollars. You need to discount this future benefit back to what it is worth today.

Using a discount rate of 6 percent (relatively appropriate for this situation—more on that in a bit), you can use a *net present value* calculation (again see Chapter 3 if you need a refresher) to translate the benefit of \$100,000 in ten years into today's dollars given the 6 percent discount rate. The formula is $\$100{,}000/1.06^{10}$ and you get the result of \$55,839.

That's all you need for a relatively sophisticated cost-benefit analysis right there! To finish the analysis, just add up all the discounted costs and benefits in today's dollars. You have the discounted benefit of \$55,839 minus the initial cost of \$50,000, netting you \$5,839.

You want the net benefit to be positive or else the deal isn't worth doing, since you'd end up worse off (in today's dollars). In this case, the net benefit is positive, so the investment is worth considering among your other options.

A central challenge with cost-benefit analysis is that this end result is sensitive to the chosen discount rate. One way to show this sensitivity is through a **sensitivity analysis**, which is a useful method to *analyze* how *sensitive* a model is to its input parameters. Using the \$50,000 bond example, let's run a sensitivity analysis on the discount rate. To do so, you just vary the discount rate and calculate the net benefit for each variation.

Sensitivity Analysis

Discount rate	Net benefit

0%	$50,000
2%	$32,033
4%	$17,556
6%	**$5,839**
8%	-$3,680
10%	-$11,446
12%	-$17,803
14%	-$23,026
16%	-$27,332

Notice how a seemingly small difference in the discount rate can represent a huge difference in the net benefit. That is, the net benefit is very sensitive to the discount rate. While the net benefit is positive at a 6 percent discount rate, it is three times more positive at 4 percent, and it is negative at 8 percent. That's because at higher discount rates, the future benefit is discounted more. Eventually, it is discounted so much that the net benefit drops into negative territory.

Running a sensitivity analysis like this can give you an idea of a range of net benefits you can expect under reasonable discount rates. You should similarly run a sensitivity analysis on any input parameter about which you are uncertain so that you can tell how much it influences the outcome.

Recall how earlier we discussed the difficulties around putting dollar values to intangible costs and benefits, such as how much not having a landlord is worth. You could use sensitivity analysis to test how much that input parameter matters to the outcome, and how a range of reasonable values would directly influence the outcome.

In general, sensitivity analysis can help you quickly uncover the key drivers in your spreadsheet inputs and show you where you may need to spend more time to develop higher accuracy in your assumptions. Sensitivity analysis is also common in statistics, and we actually already presented another one in Chapter 5 when we showcased how sample size is sensitive to *alpha* and *beta* when designing experiments.

Given that the discount rate is always a key driver in cost-benefit analyses, figuring out a reasonable range for the discount rate is paramount. To do so, consider again the factors that underlie the discount rate: inflation (that the purchasing power of money can change over time), uncertainty (that benefits may or may not actually occur), and opportunity cost of capital (that you could do other things with your money). Since these factors are situationally dependent, there is unfortunately no standard answer for what discount rate to use for any given situation.

Governments typically use rates close to their interest rates, which normally move with inflation rates. Large corporations use sophisticated methods that account for their rates of borrowing money and the return on investment seen from previous projects, together resulting in a rate that is usually significantly higher than government interest rates. New businesses, which are highly speculative, should be using much higher discount rates still, since it costs them a lot to borrow money and they are often in a race against time before they run out of money or get eaten by competitors. Thus, the range of acceptable rates can vary widely, from close to the inflation rate all the way up to 50 percent or higher in an extremely high-risk/high-reward situation.

One decent approach is to use the rate at which you can borrow money. You would want your investment returns to be higher than this rate or else you shouldn't be borrowing money to invest. Note that this rate would typically have the inflation rate already built into it, since credit rates move with interest rates, which typically move with inflation. That is, people loaning you money also want to be protected from inflation, and so they usually build an expected inflation rate into their lending rates.

"I don't know, Jack. This 'magic beans' startup
sounds kind of risky."

As investments can look very different based on different discount rates, there are many open debates about which discount rates are most appropriate to use in differing situations, especially when it comes to government programs. Different discount rates can favor one program over another, and so there can be a lot of pressure from different lobbying groups to choose a particular rate.

Another problem occurs in situations where the costs or benefits are expected to persist far into the future, such as with climate change mitigation. Because the effects of the discount rate compound over time, even rather small rates discount far-future effects close to zero. This has the effect of not valuing the consequences to future generations, and some economists think that is unfair and potentially immoral.

Even with this central issue around discount rate, cost-benefit analysis is an incredibly valuable model to frame a more quantitative discussion around how to proceed with a decision. As such, many governments mandate its use when evaluating policy options. In 1981, U.S. President Ronald Reagan signed Executive Order 12291, which mandated that "regulatory action shall not be undertaken unless the potential benefits to society from the regulation outweigh the potential costs to society." This language has been tweaked by subsequent U.S. presidents, though the central idea of it continues to drive policy, with the U.S. federal government conducting cost-benefit analyses for most significant proposed regulatory actions.

One final issue with cost-benefit analysis to keep in mind is the trickiness of comparing two options that have different time horizons. To illustrate this trap, let's compare our theoretical bond investment from earlier to another bond investment. Our bond investment from before cost $50,000 and returned $100,000 in ten years, which at a 6 percent discount rate resulted in a net benefit in today's dollars of $5,839.

Our new investment will also be a $50,000 bond investment, though instead of returning $100,000 in ten years, it pays back $75,000 in just six years. The cost today (year 0) for this second bond is again –$50,000. Using the same 6 percent discount rate, the $75,000 benefit six years from now discounted back to today's dollars would be worth $52,872, for a net benefit of $2,872 ($52,872 – $50,000). This net benefit is less than the net benefit of the first bond investment opportunity of $5,839 and so it seems the first bond is a better investment.

However, if you purchased the second bond, your $75,000 would be freed up after six years, leaving you four more years to invest that money in another way. If you were able to invest that money in a new investment at a high enough rate, this second bond is potentially more attractive in the end. When making a comparison, you therefore must consider what could happen over the same time frame.

In other words, cost-benefit analysis is only as good as the numbers you put into it. In computer science, there is a model describing this phenomenon: **garbage in**, **garbage out**. If your estimates of costs and benefits are highly inaccurate, your timelines don't line up, or your discount rate is poorly reasoned (*garbage in*), then your net result will be similarly flawed (*garbage out*).

On the other hand, if you take great care to make accurate estimates and perform relevant sensitivity analyses, then cost-benefit analysis can be a first-rate model for framing a decision-making process, and is in most cases a desirable replacement for a pro-con list. Next time you make a pro-con list, at least consider the scoring method to turn it into a simple cost-benefit analysis.

TAMING COMPLEXITY

When you can list out your options for a decision, and their costs and benefits are relatively clear, then cost-benefit analysis is a good starting point for approaching the decision. However, in many cases, your options and their associated costs and benefits are not very clear. Sometimes there is too much uncertainty in potential outcomes; other times, the situation can be so complex that it becomes difficult even to understand your options in the first place. In either case, you'll need to use some other mental models to navigate such complexity.

Consider a relatively common situation that homeowners face: the expensive repair. Suppose you want to repair your pool equipment before the summer swimming season. You get bids from two contractors. One bid is from your usual dependable pool service, but it seems high at $2,500. The second bid comes in at a lower cost of $2,000, though this contractor is a team of one, you don't have a history with them, and they also seem like they might be a little out of their depth.

As such, you get the impression that there is only a 50 percent chance that this contractor will finish at the quoted cost in a timely manner (in one week). If not, you estimate the following scenarios:

- A 25 percent chance that they will be one week late at an extra cost of $250 for the extra labor
- A 20 percent chance that they will be two weeks late at an extra cost of $500
- A 5 percent chance that they will not only take longer than three weeks to complete the job, but also that some of their work will need to be redone, totaling extra costs of $1,000

This situation (multiple bids with timing/quality concerns) is very common, but because of the uncertainty introduced in the outcome, it's a bit too complex to analyze easily with just cost-benefit analysis. Luckily, there is another straightforward mental model you can use to make sense of all these potential outcomes: the **decision tree**. It's a diagram that looks like a *tree* (drawn on its side), and helps you analyze *decisions* with uncertain outcomes. The branches (often denoted by squares) are decision points and the leaves represent different possible outcomes (often using open circles to denote chance points). A decision tree that represents this pool situation could look like the figure below.

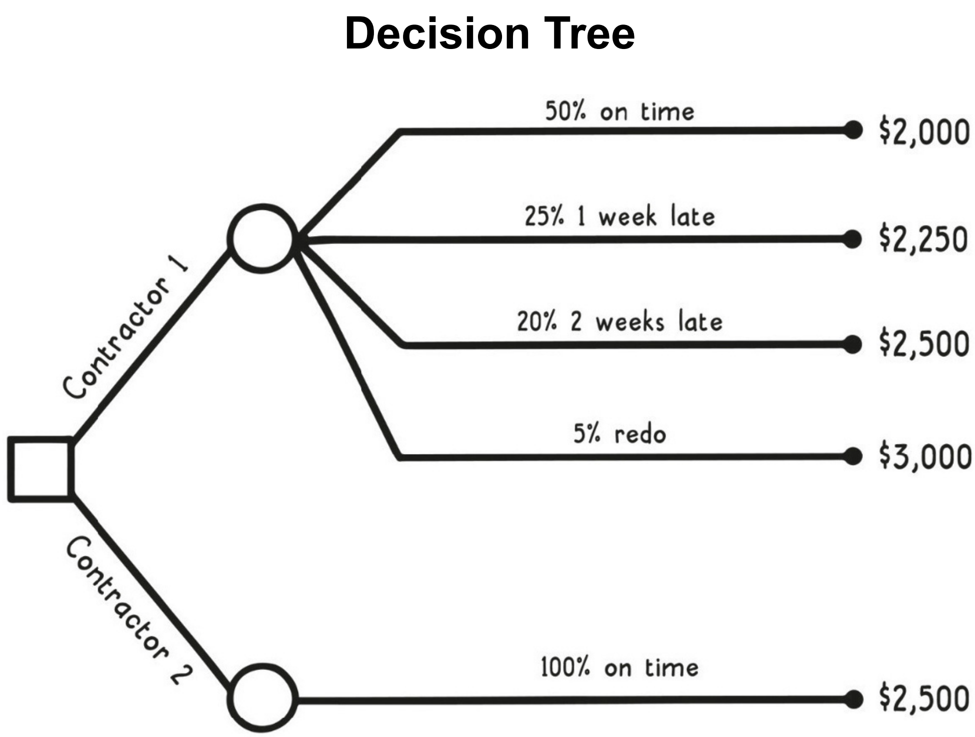

The first square represents your choice between the two contractors, and then the open circles further branch out to the different possible outcomes for each of those choices. The leaves with the closed circles list the resulting costs for each outcome, and their probabilities are listed on each line. (This is a simple *probability distribution* [see Chapter 5], which describes how all the probabilities are distributed across the possible outcomes. Each group of probabilities sums to 100 percent, representing all the possible outcomes for that choice.)

You can now use your probability estimates to get an **expected value** for each contractor, by multiplying through each potential outcome's probability with its cost, and then summing them all up. This resulting summed *value* is what you would *expect* to pay on average for each contractor, given all the potential outcomes.

The expected value for your usual contractor (Contractor 2 in the decision tree) is just $2,500, since there is only one possible outcome. The expected value for the new contractor (Contractor 1 in the decision tree) is the sum of the multiplications across their four possible outcomes: $1,000 + $562.50 + $500 + $150 = $2,212.50. Even though the new contractor has an outcome that might cost you $3,000, the expected value you'd pay is still less than you'd pay your usual contractor.

Expected Value

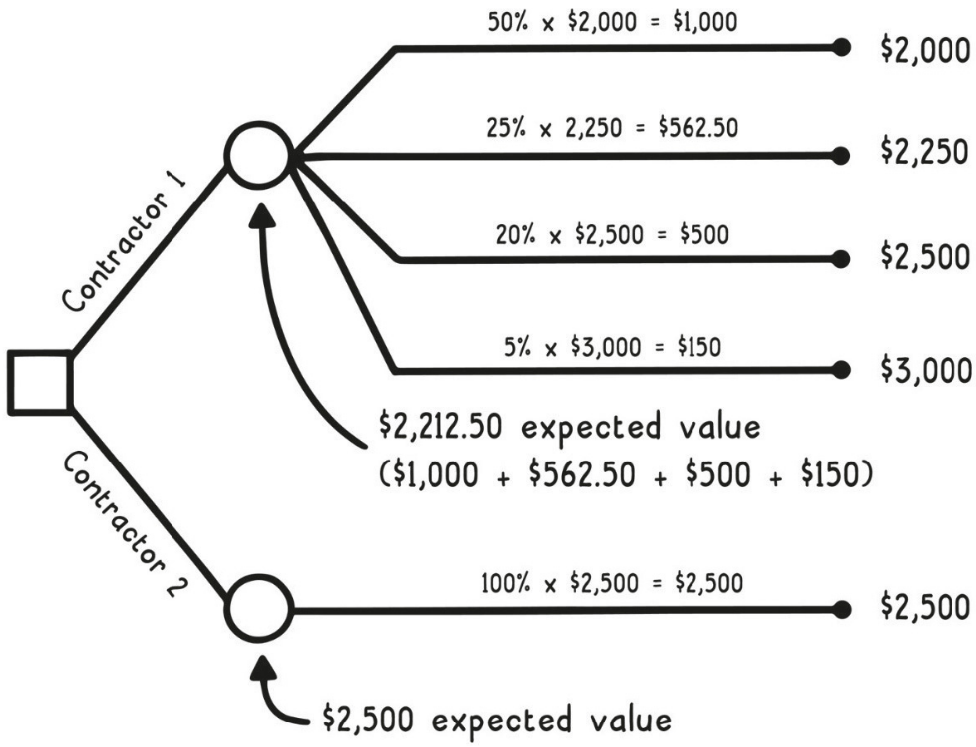

What this means is that if these probabilities are accurate, and you could run the scenario one hundred times in the real world where you pick the new contractor each time, your average payment to them would be expected to be $2,212.50. That's because half the time you'd pay only $2,000, and the other half, more. You'd never pay exactly $2,212.50, since that isn't a possible outcome, but overall your payments would average out to that expected value over many iterations.

If you find this confusing, the following example might be helpful. In 2015, U.S. mothers had 2.4 kids on average. Does any particular mother have exactly 2.4 kids? We hope not. Some have one child, some two, some three, and so on, and it all averages out to 2.4. Likewise, the various contractor payment outcomes and their probabilities add up to the expected value payment amount, even though you never pay that exact amount.

In any case, from this lens of the decision tree and the resulting expected values, you might rationally choose the new contractor, even with all their potential issues. That's because your expected outlay is lower with that contractor.

Of course, this result could change with different probabilities and/or potential outcome payments. For example, if you thought that, instead of a 5 percent chance for a $3,000 bill, there was a 50 percent chance you could end in this highest outcome, then the expected value for the new contractor would become higher than your usual contractor's bid. Remember that you can always run a sensitivity analysis on any inputs that you think might significantly influence the decision, as we discussed in the last section. Here you would vary the probabilities and/or potential outcome payments and see how the expected values change accordingly.

Additionally, consider another way the decision could change. Suppose you've already scheduled a pool party a few weeks out. Now, if the lower-bid contractor pushes into that second week, you're going to be faced with a lot of anxiety about your party. You will have to put pressure on the contractor to get the job done, and you might even have to bring in reinforcements to help finish the job at a much higher cost. That's a lot of extra hassle.

To a wealthier person who associates a high opportunity cost with their time, all this extra anxiety and hassle may be valued at an extra $1,000 worth of cost, even if you aren't paying that $1,000 directly to the contractor. Accounting for this possible extra burden would move up the two-week-late outcome from $2,500 (previously a $500 overrun) to $3,500 (now a $1,500 overrun).

Similarly, if this new contractor really messes up the job and you do have to bring in your regular contractor to do most everything over again on short notice, it will cost you the extra $1,000 in anxiety and hassle, as well as literally more payment to the other contractor.

So, that small 5 percent chance of a $3,000 outcome might end up costing the equivalent of an extra $2,000, moving it to $5,000 in total.

By using these increased values in your decision tree, you can effectively “price in” the extra costs. Because these new values include more than the exact cost you’d have to pay out, they are called **utility values**, which reflect your total relative preferences across the various scenarios. We already saw this idea in the last section when we discussed putting a price to the preference of not having a landlord. This is the mental model that encapsulates the concept.

Utility values can be disconnected from actual prices in that you can value something more than something else, even though it costs the same on the open market. Think about your favorite band—it’s worth more to you to see them in concert than another band that offers their concerts at the same price, simply because you like them more. You would get more utility out of that concert because of your preference. In the pool case, the stress involved with scrambling to fix the pool before your party is an extra cost of lost utility in addition to the actual cost you would have to pay out to the contractors.

In terms of the decision tree, the outcome values for the leaves can become the utility values, incorporating all the costs and benefits (tangible and intangible) into one number for each possible outcome. If you do that, then the conclusion now results in a flipped decision to use your usual contractor (Contractor 2 in the decision tree below).

Utility Values

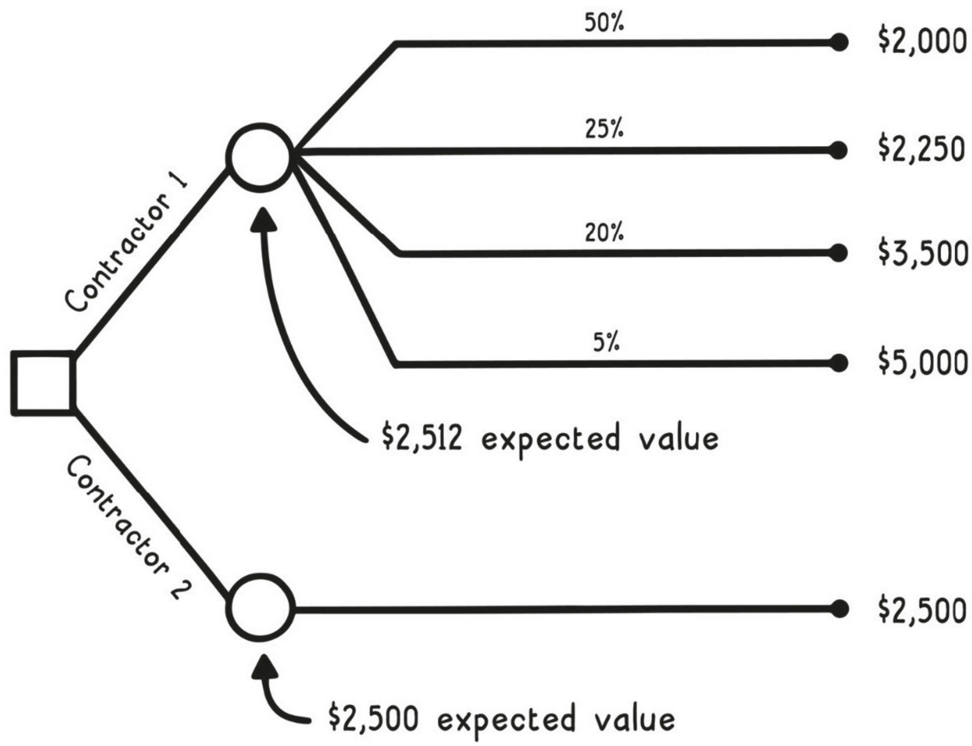

However, note that it is still a really close decision, as both contractors now have almost the same expected value! This closeness illustrates the power of probabilistic outcomes. Even though the new contractor is now associated with much higher potential “costs,” 50 percent of the time you’d still expect to pay them a much smaller amount. This lower cost drives the expected value down a lot because it happens so frequently.

Just as in cost-benefit analysis and scoring pro-con lists, we recommend using utility values whenever possible because they paint a fuller picture of your underlying preferences, and therefore should result in more satisfactory decisions. In fact, more broadly, there is a philosophy called **utilitarianism** that expresses the view that the most ethical decision is the one that creates the most *utility* for all involved.

Utilitarianism as a philosophy has various drawbacks, though. Primarily, decisions involving multiple people that increase overall utility can seem quite unfair when that utility is not equally distributed among the people involved (e.g., income inequality despite rising standards of living). Also, utility values can be hard to estimate. Nevertheless, utilitarianism is a useful philosophical model to be aware of, if only to consider what decision would increase overall utility the most.

In any case, decision trees will help you start to make sense of what to do in situations with an array of diverse, probabilistic outcomes. Think about health insurance—should you go for a higher-deductible plan with lower payments or a lower-deductible plan with higher payments? It depends on your expected level of care, and whether you can afford the lower-probability scenario where you will need to pay out a high deductible. (Note that the answer isn't obvious, because with the lower-deductible plan you are making higher monthly premium payments. This increase in premiums could be viewed as paying out a portion of your deductible each month.) You can examine this scenario and others like it via a decision tree, accounting for your preferences along with the actual costs.

Decision trees are especially useful to help you think about unlikely but severely impactful events. Consider more closely the scenario where you have a medical incident that requires you to pay out your full deductible. For some people, that amount of outlay could equate to bankruptcy, and so the true cost of this event occurring to them is much, much higher than the actual cost of the deductible.

As a result, if you were in this situation, you would want to make the loss in utility value for this scenario extremely high to reflect your desire to avoid bankruptcy. Doing so would likely push you into a higher-premium plan with a lower deductible (that you can still afford), and more assurance that you would avoid bankruptcy. In other words, if there is a chance of financial ruin, you might want to avoid that plan even though on average it would lead to a better financial outcome.

One thing to watch out for in this type of analysis is the possibility of **black swan events**, which are extreme, consequential events (that end in things like financial ruin), but which have significantly higher probabilities than you might initially expect. The name is derived from the false belief, held for many centuries in Europe and other places, that black swans did not exist, when in fact they were (and still are) common birds in Australia.

As applied to decision tree analysis, a conservative approach would be to increase your probability estimates of low-probability but highly impactful scenarios like the bankruptcy one. This revision would account for the fact that the scenario might represent a black swan event, and that you might therefore be wrong about its probability.

One reason that the probability of black swan events may be miscalculated relates to the *normal distribution* (see Chapter 5), which is the bell-curve-shaped probability distribution that explains the frequency of many natural phenomena (e.g., people's heights). In a normal distribution, rare events occur on the tails of the distribution (e.g., really tall or short people), far from the middle of the bell curve. Black swan events, though, often come from **fat-tailed distributions**, which literally have *fatter tails*, meaning that events way out from the middle have a much higher probability when compared with a normal distribution.

Fat-Tailed Distribution

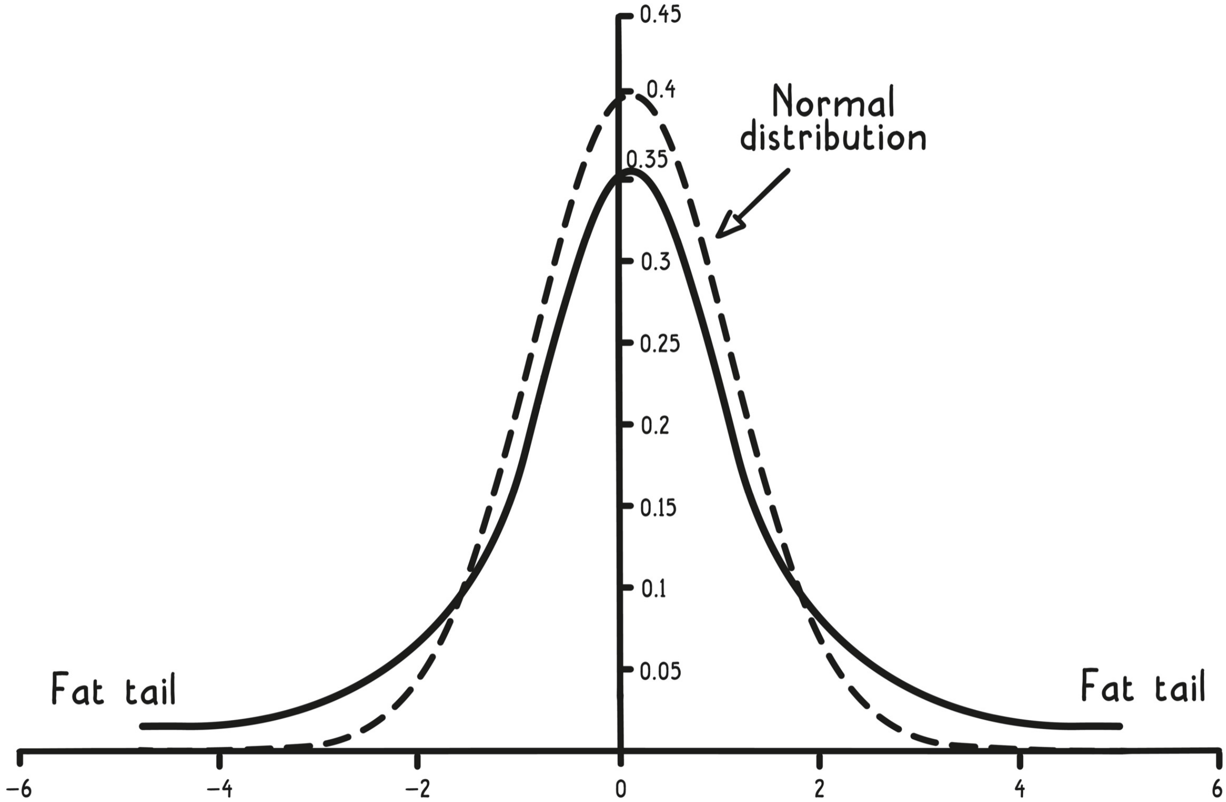

There are many naturally occurring fat-tailed distributions as well, and sometimes people just incorrectly assume they are dealing with a normal distribution when in fact they are dealing with a distribution with a fatter tail, and that means that events in the tail occur with higher probability. In practice, these are distributions where some of the biggest outliers happen more often than you would expect from a normal distribution, such as occurs with insurance payouts, or in the U.S. income distribution (see the *histogram* in Chapter 5).

Another reason why you might miscalculate the probability of a black swan event is that you misunderstand the reasons for its occurrence. This can happen when you think a situation should come from one distribution, but multiple are really involved. For example, there are genetic reasons (e.g., dwarfism and Marfan syndrome) why there might be many more shorter or taller people than you would expect from just a regular normal distribution, which doesn't account for these rarer genetic variations.

A third reason is that you may underestimate the possibility and impact of *cascading failures* (see Chapter 4). As you recall, in a cascading-failure scenario, parts of the system are correlated: if one part falters, the next part falters, and so on. The 2007/2008 financial crisis is an example, where the failure of mortgage-backed securities cascaded all the way to the banks and associated insurance companies.

Our climate presents another example. The term *one-hundred-year flood*, denotes a flood that has a 1 percent chance of occurring in any given year. Unfortunately, climate change is raising the probability of the occurrence of what was once considered a one-hundred-year flood, and it no longer has a 1 percent chance in many areas. The dice are loaded. Houston, Texas, for example, has had three so-called five-hundred-year floods in the last three years! The probabilities of these events clearly need to be adjusted as the cascading effects of climate change continue to unfold.

To better determine the outcome probabilities in highly complex systems like banking or climate, you may first have to take a step back and try to make sense of the whole system before you can even try to create a decision tree or cost-benefit analysis for a particular subset or situation. **Systems thinking** describes this act, when you attempt to *think* about the entire *system* at once. By thinking about the overall system, you are more likely to understand and account for subtle interactions between components that could otherwise lead to unintended consequences from your decisions. For example, when thinking about making an investment, you might start to appreciate how seemingly unrelated parts of the economy might affect its outcome.

Some systems are fairly simple and you can picture the whole system in your head. Others are so complex that it is too challenging simultaneously to hold all the interlocking pieces in your head. One solution is literally to diagram the system visually. Drawing diagrams can help you get a better sense of complex systems and how the parts of the system interact with one another.

Techniques for how to effectively diagram complex systems are beyond the scope of this book, but know that there are many techniques that you can learn, including *causal loop diagrams* (which showcase feedback *loops* in a system) and *stock and flow diagrams* (which showcase how things accumulate and *flow* in a system). Gabriel's master's thesis involved diagraming the email spam system. The picture on the next page is one of his causal loop diagrams—you aren't meant to understand this diagram; it's just an example of what these things can end up looking like. Just know now that it was really helpful in gaining a much better understanding of this complex system.

Email Spam Causal Loop Diagram

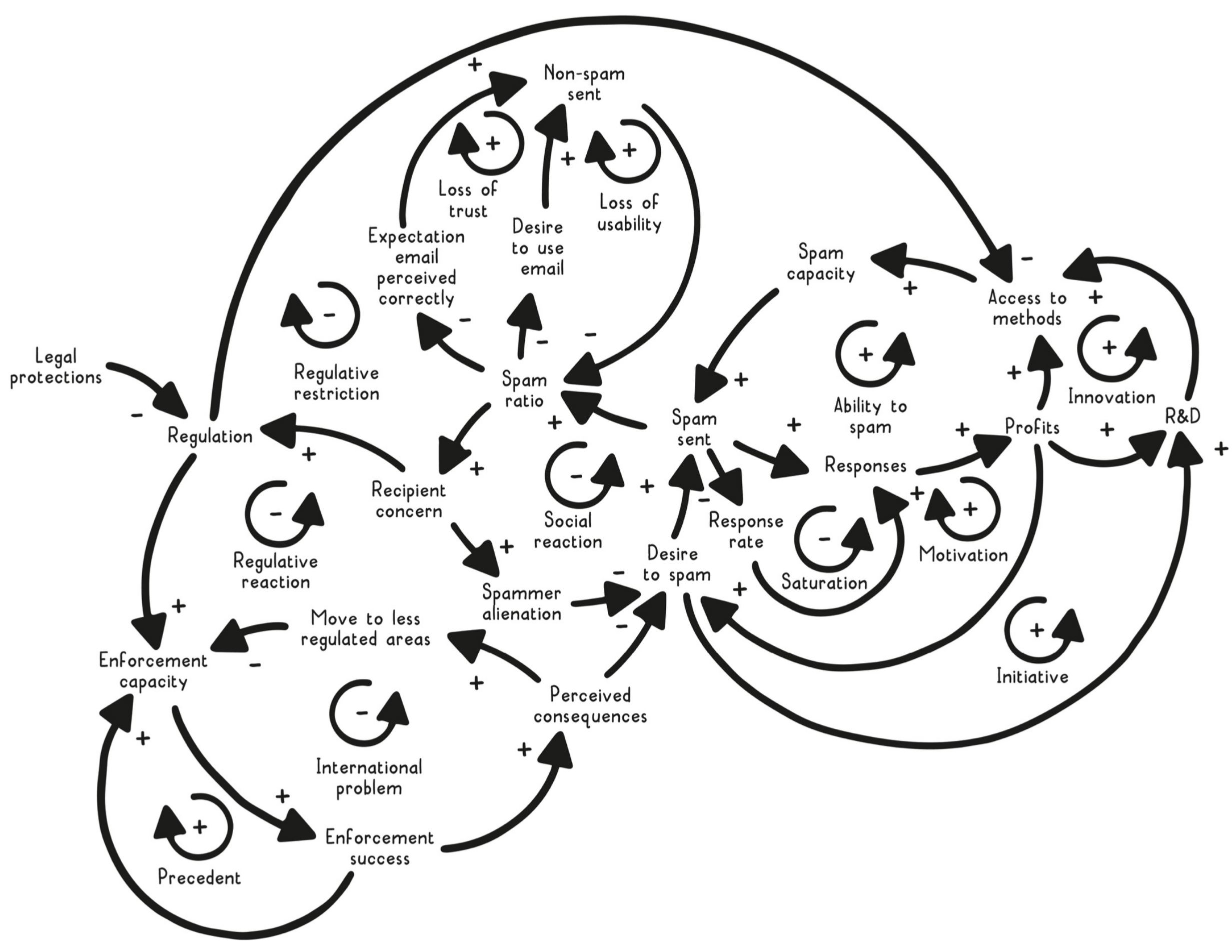

As a further step, you can use software to imitate the system, called a *simulation*. In fact, software exists that allows you to compose a diagram of a system on your screen and then immediately turn it into a working simulation. (Two such programs that do this online are Insight Maker and True-World.) In the process, you can set initial conditions, and then see how the system unfolds over time.

Simulations help you more deeply understand a complex system and lead to better predictions of black swans and other events. Simulations can also help you identify how a system will adjust when faced with changing conditions. **Chatelier's principle**, named after French chemist Henri-Louis Le Chatelier, states that when any chemical system at equilibrium is subject to a change in conditions, such as a shift in temperature, volume, or pressure, it readjusts itself into a new equilibrium state and usually partially counteracts the change.

For example, if someone hands you a box to carry, you don't immediately topple over; you instead shift your weight distribution to account for the new weight. Or in economics, if a new tax is introduced, tax revenues from that tax end up being lower in the long run than one would expect under current conditions because people adjust their behavior to avoid the tax.

If this sounds like a familiar concept, it's because Chatelier's principle is similar to the mental model *homeostasis* (see Chapter 4), which comes from biology: recall how your body automatically shivers and sweats in response to external conditions in order to regulate its internal temperature. Chatelier's principle doesn't necessarily mean the system will regulate around a predetermined value, but that it will react to externally imposed conditions, and usually in a way that partially counteracts the external stimulus. You can see the principle in action in real time with simulations because they allow you to calculate how your simulated system will adjust to various changes.

A related mental model that also arises in dynamic systems and simulations is **hysteresis**, which describes how a system's current state can be dependent on its history. Hysteresis is also a naturally occurring phenomenon, with examples across most scientific disciplines. In physics, when you magnetize a material in one direction, such as by holding a magnet to another piece of metal, the metal does not fully demagnetize after you remove the magnet. In biology, the T cells that help power your immune system, once activated, thereafter require a lower threshold to reactivate. Hysteresis describes how both the metal and the T cells partially remember their states, such that what happened previously can impact what will happen next.

Again, this may already seem like a familiar concept, because it is similar to the mental model of *path dependence* (see Chapter 2), which more generally describes how choices have consequences in terms of limiting what you can do in the future. Hysteresis is one type of path dependence, as applied to systems.

In engineering systems, for example, it is useful to build some hysteresis into the system to avoid rapid changes. Modern thermostats do this by allowing for a range of temperatures around the set point: if you want to maintain 70 degrees Fahrenheit, a thermostat might be set to turn the heater on when the temperature drops to 68 degrees and back off when it hits 72 degrees. In this way, it isn't kicking on and off constantly. Similarly, on websites, designers and developers often build in a lag for when you move your mouse off page elements like menus. They build their programs to remember that you were on the menu so that when you move off, it doesn't abruptly go away, which can appear jarring to the eye.

You can use all these mental models around visualizing complex systems and simulating them to help you better assess potential outcomes and their associated probabilities. Then you can feed these results into a more straightforward decision model like a decision tree or cost-benefit analysis.

A particular type of simulation that can be especially useful in this way is a **Monte Carlo simulation**. Like *critical mass* (see Chapter 4), this is a model that emerged during the Manhattan Project in Los Alamos in the run-up to the discovery of the atomic bomb. Physicist Stanislaw Ulam was struggling with using traditional mathematics to determine how far neutrons would travel through various materials and came up with this new method after playing solitaire (yes, the card game). In his words, quoted in *Los Alamos Science*:

> The first thoughts and attempts I made to practice [the Monte Carlo method] were suggested by a question which occurred to me in 1946 as I was convalescing from an illness and playing solitaires. The question was what are the chances that a Canfield solitaire laid out with 52 cards will come out successfully? After spending a lot of time trying to estimate them by pure combinatorial calculations, I wondered whether a more practical method than "abstract thinking" might not be to lay it out say one hundred times and simply observe and count the number of successful plays.

A Monte Carlo simulation is actually many simulations run independently, with random initial conditions or other uses of random numbers within the simulation itself. By running a simulation of a system many times, you can begin to understand how probable different outcomes really are. Think of it as a dynamic sensitivity analysis.

Monte Carlo simulations are used in nearly every branch of science. But they are useful outside science as well. For example, venture capitalists often use Monte Carlo simulations to determine how much capital to reserve for future financings. When a venture fund invests in a company, that company, if successful, will probably raise more money in the future, and the fund will often want to participate in some of those future financings to maintain its ownership percentage. How much money should it reserve for a company? Not all companies are successful, and different companies raise different amounts, so the answer is not straightforward at the time of the initial investment. Many funds use Monte Carlo simulations to understand how much they ought to reserve, given their current fund history and the estimates of company success and size of potential financings.

More generally, making the effort to understand complex systems better through systems thinking—whether it be by using diagrams, running simulations, or employing other mental models—not only helps you get a broad picture of the system and its range of outcomes, but also can help you become aware of the best possible outcomes. Without such knowledge, you can get stuck chasing a

local optimum solution, which is an admittedly good solution, but not the best one.

If you can, you want to work toward that best solution, which would be the **global optimum**. Think of rolling hills: the top of a nice nearby hill would be a good success (local optimum), though in the distance there is a much bigger hill that would be a much better success (global optimum). You want to be on that bigger hill. But first you have to have a full view of the system to know the bigger hill exists.

Local vs. Global Optimum

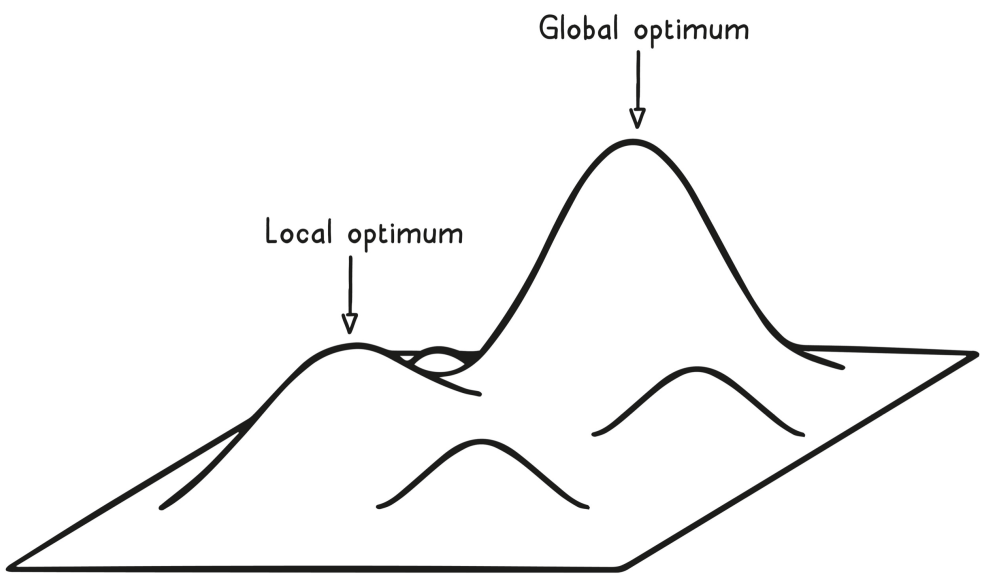

BEWARE OF UNKNOWN UNKNOWNS

In 1955, psychologists Joseph Luft and Harrington Ingham originated the concept of **unknown unknowns**, which was made popular by former U.S. Secretary of Defense Donald Rumsfeld at a news briefing on February 12, 2002, with this exchange:

> Jim Miklaszewski: In regard to Iraq, weapons of mass destruction, and terrorists, is there any evidence to indicate that Iraq has attempted to or is willing to supply terrorists with weapons of mass destruction? Because there are reports that there is

no evidence of a direct link between Baghdad and some of these terrorist organizations.

Rumsfeld: Reports that say that something hasn't happened are always interesting to me, because as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns—the ones we don't know we don't know. And if one looks throughout the history of our country and other free countries, it is the latter category that tend to be the difficult ones.

The context and evasiveness of the exchange aside, the underlying model is useful in decision making. When faced with a decision, you can use a handy *2 × 2 matrix* (see Chapter 4) as a starting point to envision these four categories of things you know and don't know.

Knowns & Unknowns

	Known	Unknown
Known	What you know you know	What you know you don't know
Unkown	What you don't know you know	What you don't know you don't know

This model is particularly effective when thinking more systematically about risks, such as risks to a project's success. Each category deserves its own attention and process:

- *Known knowns:* These might be risks to someone else, but not to you since you already know how to deal with them based on your previous experience. For example, a project might require a technological solution, but you already know what that solution is and how to implement it; you just need to execute that known plan.
- *Known unknowns:* These are also known risks to the project, but because of some uncertainty, it isn't exactly clear how they will be resolved. An example is the risk of relying on a third party: until you engage with them directly, it is unknown how they will react. You can turn some of these into known knowns

by doing *de-risking* exercises (see Chapter 1), getting rid of the uncertainty.

- *Unknown knowns:* These are the risks you're not thinking about, but for which there exist clear mitigation plans. For example, your project might involve starting to do business in Europe over the summer, but you don't yet know they don't do much business in August. An adviser with more experience can help identify these risks from the start and turn these into known knowns. That way they will not take you by surprise later on and potentially throw off your project.
- *Unknown unknowns:* These are the least obvious risks, which require a concerted effort to uncover. For example, maybe something elsewhere in the organization or in the industry could dramatically change this project (like budget cuts or an acquisition or new product announcement). Even if you identify an unknown unknown (turning it into a known unknown), you still remain unsure of its likelihood or consequences. You must then still do de-risking exercises to finally turn it into a known known.

As you can see, you enumerate items in each of the four categories, and then work to make them all known knowns. This model is about giving yourself more complete knowledge of a situation. It's similar to systems thinking, from the last section, in that you are attempting to get a full picture of the system so you can make better decisions.

As a personal example, consider having a new baby. From reading all the books, you know the first few weeks will be harrowing, you'll want to take some time off work, you'll need to buy a car seat, crib, diapers, etc.—these are the known knowns. You also know that how your baby might sleep and eat (or not) can be an issue, but until the baby is born, their proclivities remain uncertain—they are known unknowns. You might not yet know that swaddling a baby is a thing, but you'll be shown how soon enough by a nurse or family member, turning this unknown known into a known known. And then there are things that no one knows yet or is even thinking about, such as whether your child could have a learning disability.

A related model that can help you uncover unknown unknowns is **scenario analysis** (also known as *scenario planning*), which is a method for thinking about possible futures more deeply. It gets its name because it involves *analyzing* different *scenarios* that might unfold. That sounds simple enough, but it is deceptively complicated in practice. That's because thinking up possible future scenarios is a really challenging exercise, and thinking through their likelihoods and consequences is even more so.

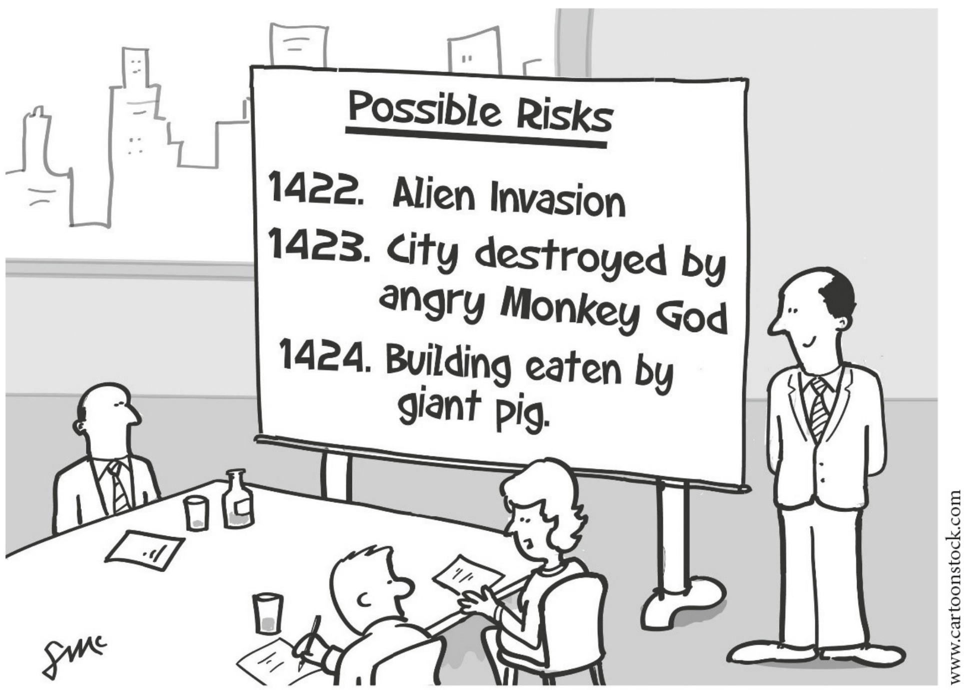

"Well he certainly does a very thorough risk analysis."

Governments and large corporations have dedicated staff for scenario analysis. They are continually thinking up and writing reports about what the world could look like in the future and how their citizenry or shareholders might fare under those scenarios. Many academics, especially in political science, urban planning, economics, and related fields, similarly engage in prognosticating about the future. And of course, science fiction is essentially an entire literary genre dedicated to scenario analysis.

To do scenario analysis well, you must conjure plausible yet distinct futures, ultimately considering several possible scenarios. This process is difficult because you tend to latch onto your first thoughts (see *anchoring* in Chapter 1), which usually depict a direct extrapolation of your current trajectory (the present), without challenging your own assumptions.

One technique to ensure that you do challenge your assumptions is to list major events that could transpire (e.g., stock market crash, government regulation, major industry merger, etc.) and then trace their possible effects back to your situation. Some may have little to no effect, whereas others might form the basis for a scenario you should consider deeply.

Another technique for thinking more broadly about possible future scenarios is the **thought experiment**, literally an *experiment* that occurs just in your *thoughts*, i.e., not in the physical world. The most famous thought experiment is probably "Schrödinger's cat," named after Austrian physicist Erwin Schrödinger, who thought it up in 1935 to explore the implications of different interpretations of the physics of quantum mechanics. From his 1935 paper "The Present Situation in Quantum Mechanics":

> A cat is penned up in a steel chamber, along with the following device (which must be secured against direct interference by the cat): in a Geiger counter, there is a tiny bit of radioactive substance, so small, that perhaps in the course of the hour one of the atoms decays, but also, with equal probability, perhaps none; if it happens, the counter tube discharges and through a relay releases a hammer that shatters a small flask of hydrocyanic acid. If one has left this entire system to itself for an hour, one would say that the cat still lives if meanwhile no atom has decayed. The first atomic decay would have poisoned it.

So, you have a cat in a box, and if a radioactive atom decayed in the last hour, it would have killed the cat. This thought experiment poses some seemingly unanswerable questions: Until you observe the cat by opening the box, is it alive or dead, or in an in-between state, as certain interpretations of quantum mechanics would suggest? And what exactly happens when you open the box?

Schrödinger's Cat Thought Experiment

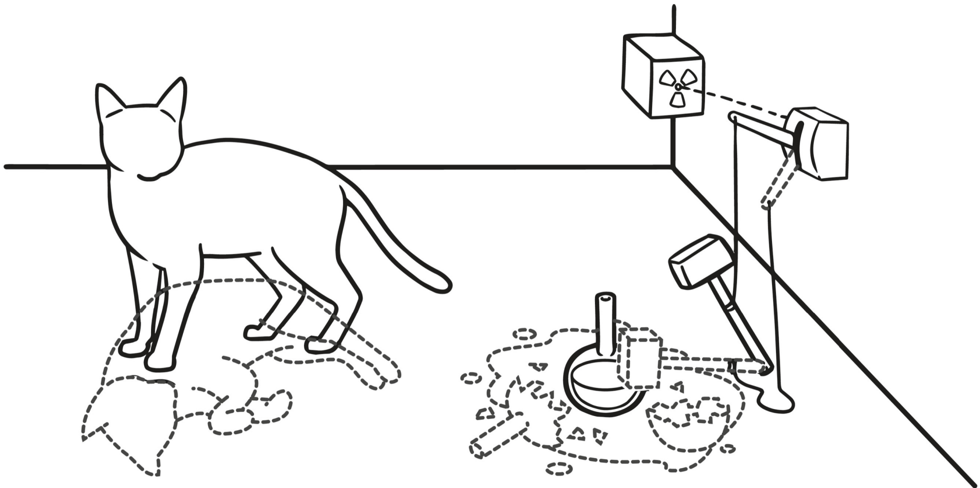

Answers to this thought experiment are beyond the scope of this book and were argued over for decades after it was posed. Therein lies the power of the thought experiment.

Thought experiments are particularly useful in scenario analysis. Posing questions that start with "What would happen if . . ." is a good practice in this way: What would happen if life expectancy jumped forty years? What would happen if a well-funded competitor copied our product? What would happen if I switched careers?

These types of what-if questions can also be applied to the past, in what is called **counterfactual thinking**, which means *thinking* about the past by imagining that the past was different, *counter* to the *facts* of what actually occurred. You've probably seen this model in books and movies about scenarios such as what would have happened if Germany had won World War II (e.g., Philip K. Dick's *The Man in the High Castle*). Examples from your own life can help you improve your decision making when you think through the possible consequences of your past decisions. What if I had taken that job? What if I had gone to that other school? What if I hadn't done that side project?

When reconsidering your past decisions, though, it is important not only to think of the positive consequences that might have occurred if you had made a different life choice. The *butterfly effect* (see Chapter 4) reminds us that one small change can have ripple effects, so when considering a counterfactual scenario, it is important to remember that if you change one thing, it is unlikely that everything else would stay the same.

Posing what-if questions can nevertheless help you think more creatively, coming up with scenarios that diverge from your intuition. More generally, this technique is one of many associated with **lateral thinking**, a type of *thinking* that helps you move *laterally* from one idea to another, as opposed to critical thinking, which is more about judging an idea in front of you. Lateral thinking is *thinking outside the box*.

Another helpful lateral-thinking technique involves adding some randomness when you are generating ideas. For example, you can choose an object at random from your surroundings or a noun from the dictionary and try to associate it in some way with your current idea list, laterally forming new offshoot ideas in the process.

No matter what techniques you use, however, it is extremely difficult to perform scenario analysis alone. Seeking outside input produces better results, as different people with different perspectives bring new ideas to the table.

It is therefore tempting to involve multiple people in brainstorming sessions from the get-go. However, studies show this is not the right approach because of **groupthink**, a bias that emerges because *groups* tend to *think* in harmony. Within group settings, members often strive for consensus, avoiding conflict, controversial issues, or even alternative solutions once it seems a solution is already favored by the group.

The **bandwagon effect** describes the phenomenon whereby consensus can take hold quickly, as other group members "hop on the *bandwagon*" as an idea gains popularity. More generally, it describes people's tendency to take social cues and follow the decisions of others. In this way, the probability of a person adopting an idea increases the more other people have already done so.

In some cases, this is rational behavior, as when you follow the bandwagon and adopt a product based on well-researched reviews from owners of the product. In other cases, though, fads and trends can be based on little substance.

"Put me down for whoever comes out ahead in your poll."

Groupthink is terrible for scenario analysis and can have much wider implications, leading to bad group decision making in general if not actively managed. There are many ways to manage groupthink, though, including setting a culture of questioning assumptions, making sure to evaluate all ideas critically, establishing *a Devil's advocate position* (see Chapter 1), actively recruiting people with differing opinions, reducing leadership's influence on group recommendations, and splitting the group into independent subgroups.

It is this last recommendation that is particularly relevant for scenario analysis, as it forms the basis for **divergent thinking**, where you actively try to get *thinking* to *diverge* in order to discover multiple possible solutions, as opposed to **convergent thinking**, where you actively try to get *thinking* to *converge* on one solution.

One tactic is to meet once without brainstorming at all, just to go over the goal of the scenario analysis. Then send everyone off individually or in small groups. You could give them a prompt to react to, such as survey data, or have them come up with their own thought experiments and scenario ideas from scratch (divergent thinking). Finally, you bring everyone back together to go over all the proposed scenarios in order to narrow them down to just a few scenarios to explore further (convergent thinking).

It is additionally likely that people close to you, such as those within your organization, share similar cultural traits, and therefore you should look beyond your normal contacts and venture outside your organization to get as much lateral and divergent thinking as you can. One way to do so is actively to seek out people from different backgrounds to participate. Another way, easily enabled by the internet, is to **crowdsource** ideas, where you seek (*source*) ideas quite literally from anyone who would like to participate (the *crowd*).

Crowdsourcing has been effective across a wide array of situations, from soliciting tips in journalism, to garnering contributions to Wikipedia, to solving the real-world problems of companies and governments. For example, Netflix held a contest in 2009 in which crowdsourced researchers beat Netflix's own recommendation algorithms.

Crowdsourcing can help you get a sense of what a wide array of people think about a topic, which can inform your future decision making, updating your prior beliefs (see *Bayesian statistics* in Chapter 5). It can also help you uncover unknown unknowns and unknown knowns as you get feedback from people with previous experiences you might not have had.

In James Surowiecki's book *The Wisdom of Crowds*, he examines situations where input from crowds can be particularly effective. It opens with a story about how the crowd at a county fair in 1906, attended by statistician Francis Galton, correctly guessed the weight of an ox. Almost eight hundred people participated, each individually guessing, and the average weight guessed was 1,197 pounds—exactly the weight of the ox, to the pound! While you cannot expect similar results in all situations, Surowiecki explains the key conditions in which you can expect good results from crowdsourcing:

- *Diversity of opinion:* Crowdsourcing works well when it draws on different people's private information based on their individual knowledge and experiences.
- *Independence:* People need to be able to express their opinions without influence from others, avoiding groupthink.
- *Aggregation:* The entity doing the crowdsourcing needs to be able to combine the diverse opinions in such a way as to arrive at a collective decision.

If you can design a system with these properties, then you can draw on the *collective intelligence* of the crowd. This allows you to glean the useful bits of information that might be hidden among a group of diverse participants. In the ox example, a butcher may notice something different than a farmer would and different yet than a vet would. All this knowledge was captured in the collective weight guessed. A more modern example of making use of collective intelligence would be an audience poll as done on the television show *Who Wants to Be a Millionaire?*

In general, drawing on collective intelligence makes sense when the group's collective pool of knowledge is greater than what you could otherwise get access to; this helps you arrive at a more intelligent decision than you would arrive at on your own. "The crowd" can help systematically think through various scenarios, get new data and ideas, or simply help improve existing ideas.

One direct application of crowdsourcing to scenario analysis is the use of a **prediction market**, which is like a stock market for predictions. In a simple formulation of this concept, the price of each stock can range between $0 and $1 and represents the market's current probability of an event taking place, such as whether a certain candidate will be elected. For example, a price of $0.59 would represent a 59 percent probability that the candidate would be elected.

If you think the probability is significantly higher than 59 percent, then you could buy a yes share at that price. Alternatively, if you think the probability is significantly lower than 59 percent, then you could buy a no share at that price. If the candidate actually gets elected, then the market pays out holders of yes predictions at $1 per share, and if they are not elected, then those yes shares become worthless. Conversely, if the candidate doesn't get elected, then the market pays out holders of the no predictions at $1 per share and the yes shares become worthless.

If more people are making yes predictions than no predications, then the price of the stock rises, and vice versa. By looking at the current prices in the prediction market, you can get a sense of what the market thinks will happen, based on how people are betting (buying shares). Many big companies operate similar prediction markets internally, where employees can predict the outcome of things like sales forecasts and marketing campaigns.

Several larger public prediction markets also exist, such as PredictIt, which focuses on political predictions in the manner described above. While this market has successfully predicted many election outcomes across the world, in 2016 it failed to correctly predict both the election of Donald Trump and the UK's Brexit vote. Retrospective analysis showed that diversity of opinion seemed lacking and that participants in the prediction market likely didn't have enough direct contact with Trump or Brexit supporters. In addition, predictors were not operating fully independently, instead being influenced by the initial outsized odds against Trump and Brexit.

Another project, called the Good Judgment Project, crowdsources predictions for world events. Its co-creator, Philip E. Tetlock, studied thousands of participants and discovered **superforecasters**, people who make excellent *forecasts*, repeatedly**.** He found that these superforecasters consistently beat the world's leading intelligence services in their predictions of world events, even though they lack classified intelligence that these services have access to!

In a book entitled *Superforecasting*, Tetlock examines characteristics that lead superforecasters to make such accurate predictions. As it happens, these are good characteristics to cultivate in general:

- *Intelligence:* Brainpower is crucial, especially the ability to enter a new domain and get up to speed quickly.
- *Domain expertise:* While you can learn about a particular domain on the fly, the more you learn about it, the more it helps.
- *Practice:* Good forecasting is apparently a skill you can hone and get better at over time.
- *Working in teams:* Groups of people can outperform individuals as long as they avoid groupthink.

- *Open-mindedness:* People who are willing to challenge their beliefs tend to make better predictions.
- *Training in past probabilities:* People who looked at probabilities of similar situations in the past were better able to assess the current probability, avoiding the *base rate fallacy* (see Chapter 5).
- *Taking time:* The more time people took to make the prediction, the better they did.
- *Revising predictions:* Forecasters who continually revised their predictions based on new information successfully avoided *confirmation bias* (see Chapter 1).

Using prediction markets and the techniques of superforecasters can help you improve your scenario analysis by making it more accurate and focusing it on the events that are actually more likely to occur. As we've seen in Chapters 2 and 4, many unpredictable changes will inevitably occur; however, by spending time with these mental models, you can be better prepared for these changes. Even if you cannot predict exactly what will happen, you may envision similar scenarios and your preparation for those scenarios will help you.

In this chapter as a whole, we've seen an array of decision models that surpass the simple pro-con list that we started with. When you've arrived at a decision using one or more of these mental models, a good final step is to produce a **business case**, a document that outlines the reasoning behind your decision.

This process is a form of *arguing from first principles* (see Chapter 1). You are laying out your premises (principles) and explaining how they add up to your conclusion (decision). You are making your *case*. Taking this explicit step will help you identify holes in your decision-making process. In addition, a business case provides a jumping-off point to discuss the decision with your colleagues.

A business case can range from very short and informal (a few paragraphs) to extremely detailed and formal (a massive report) and is often accompanied by a presentation. In its final form it is used to convince others (or yourself!) that the decision is the right one. By using the mental models from this chapter, you can put together compelling business cases to help you and your organization make excellent decisions.

And it's not just for business. We started this chapter by discussing a potential career change. Knowing what you know now, you can approach that same problem in a much better way. For example, you could do scenario analysis to better uncover and imagine how different possible career futures could unfold. You could then more systematically analyze the seemingly best possible career paths more numerically through cost-benefit analysis, or using a decision tree if some of the choices are more probabilistic in nature. Then, in the end, you can put all of it together into a succinct business case to lay out the argument for your next career move.

KEY TAKEAWAYS

- When tempted to use a **pro-con list**, consider upgrading to a **cost-benefit analysis** or **decision tree** as appropriate.
- When making any quantitative assessment, run a **sensitivity analysis** across inputs to uncover key drivers and appreciate where you may need to seek greater accuracy in your assumptions. Pay close attention to any **discount rate** used.
- Beware of **black swan events** and **unknown unknowns**. Use **systems thinking** and **scenario analysis** to more systematically uncover them and assess their impact.
- For really complex systems or decision spaces, consider **simulations** to help you better assess what may happen under different scenarios.
- Watch out for blind spots that arise from **groupthink**. Consider **divergent** and **lateral thinking** techniques when working with groups, including seeking more diverse points of view.
- Strive to understand the **global optimum** in any system and look for decisions that move you closer to it.