Chapter One

Learn the Bulletproof Problem Solving Approach

In the 1980s, when Charles was at business school, he wanted to understand the then-ascendant Japanese business practices better. He wrote to dozens of Japanese companies to see if they would host him for a summer internship. Most never replied, but just as he was thinking he might be unemployed for the summer, Charles received a letter from a Dr. Utsumi at Canon, the camera and printer company. Canon was prepared to hire Charles as its first western intern, and soon he was winging his way to Japan.

It sounds like a fun adventure, and it was, but it was also a huge shock. Charles was seconded to the production planning division in a distant Tokyo suburb, and assigned to a Canon men's dormitory, three train lines and 90 minutes away. He couldn't speak or read Japanese. He was assigned what seemed at first an impossible task: develop a model for how to site factories. He despaired—what did he know about where to put factories? It seemed like a specialist problem.

But, with the help of a translating colleague, he began to interview the team about their experiences in different factory location decisions around the world. Patterns began to emerge in his findings. He learned which variables were involved, from local authorities' incentives, to local taxation rates, wage levels, raw materials transportation cost, and so on, and eventually he figured out which were more or less important. Finally he built a logic tree that captured the variables, the direction or sign of impact, and the weight of the factors. He tested the model with data from past factory decisions and honed its accuracy with the senior team. In the end, this little model became the core tool used by the department to make complex factory siting decisions! The secret was that it was a single-page way of seeing complicated trade-offs that had previously been buried in dense reports. It made the logic of the criteria clear, and opened weighting of variables up to discussion.

It saved what might have been a disastrous internship, but more importantly, it convinced Charles of the decision-making power of relatively simple logical structures and processes in problem solving. That is the core focus of this book.

Problem solving means different things to different people. When Rob asked his seven-year-old granddaughter how school was going, she said to him, "Papa, I'm very good at problem solving." This of course was music to Rob's ears! Of course, she was really talking about doing math and logic problems in a school setting. Unfortunately, these essential problem solving building blocks are seldom taught as a systematic process and rarely in a way that addresses problems of everyday relevance and consequence. For us, problem solving means the process of making better decisions on the complicated challenges of personal life, our workplaces, and the policy sphere.

The magic of the *Bulletproof Problem Solving* approach we introduce here is in following the same systematic process to solve nearly every type of problem, from linear ones to problems with complex interdependencies. It sets out a simple but rigorous approach to defining problems, disaggregating them into manageable pieces, focusing good analytic tools on the most important parts, and then synthesizing findings to tell a powerful story. While the process has a beginning and end, we encourage you to think of problem solving as an iterative process rather than a linear one. At each stage we improve our understanding of the problem and use those greater insights to refine our early answers.

In this chapter we outline the overall *Bulletproof Problem Solving Process,* introducing you to the seven steps that later chapters will address in more detail. We demonstrate the use of logic trees to uncover the structure of problems and focus on solution paths. We provide several straightforward cases to get readers started. Later chapters will introduce advanced techniques for more complicated and uncertain problems.

The Bulletproof Problem Solving Cycle

The bulletproof problem solving process is both a complete process and an iterative cycle. This cycle can be completed over any timeframe with the information at hand. Once you reach a preliminary end point, you can repeat the process to draw out more insight for deeper understanding.

We often use the expression, "What's the one-day answer?" This means we ask our team to have a coherent summary of our best understanding of the problem and a solution path at any point in the project, not just at the end. This process of creating active hypotheses is at the heart of *Bulletproof Problem Solving*. It can even help you face the dreaded "elevator test." The elevator test is when you, as a junior team member, find yourself in an elevator with the most senior person in your organization and they ask, "How is your project going?" We have all had this happen. You panic, your mind goes blank, and you stammer out a nonsensical dog's breakfast of an answer. The bulletproof problem solving process in the following pages can help you beat this situation and turn the elevator test into an opportunity for promotion.

The kind of problem solving we describe can be done alone or in teams. If you're tackling a problem by yourself, we suggest building in review processes that you can use with family and colleagues to get the higher objectivity and other bias-fighting benefits of a team.

The seven steps are introduced in Exhibit 1.1.

EXHIBIT 1.1 The bulletproof problem solving cycle

STEP 1: DEFINE THE PROBLEM

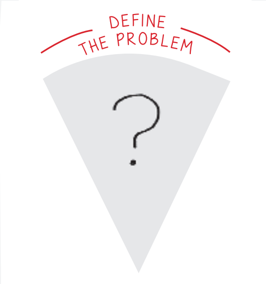

When a problem's context and boundaries aren't fully described, there is a lot of room for error. The first step in our process is to arrive at a problem definition that is agreed upon by those involved in making a decision. We test the problem definition against several criteria: that it is specific, not general, that we can clearly measure success, that the definition is bounded both in time frame and by the values of the decision maker, and that it involves definitive action being taken. This step may appear constraining, but it leads to the clarity of purpose essential for good problem solving.

STEP 2: DISAGGREGATE THE ISSUES

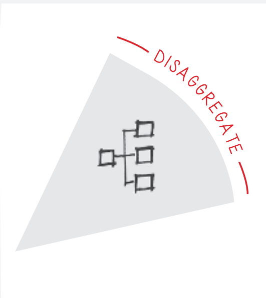

Once the problem is defined, it must be disaggregated (or broken down) into component parts or issues. We employ logic trees of various types to elegantly disassemble problems into parts for analysis, driving from alternative hypotheses of the answer. There is both an art and science to "cleaving" problems—revealing their fault lines—that drives better solutions. This is the stage at which theoretical frameworks from economics and science provide useful guides to better understanding the drivers of your problem solution. We usually try several different cuts at disaggregation to see which yields the most insight.

STEP 3: PRIORITIZE THE ISSUES, PRUNE THE TREE

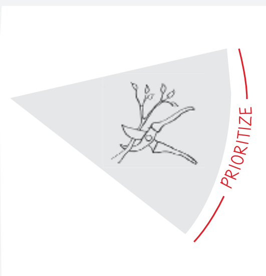

The next step is to identify which branches of the logic tree have the biggest impact on the problem, including which you can most affect, and focus your initial attention on these. We employ a simple matrix of size of impact of each lever and ability to move the lever as a way to prune our logic trees. Prioritizing analyses helps us find the critical path to the answer efficiently, making the best use of team time and resources.

STEP 4: BUILD A WORKPLAN AND TIMETABLE

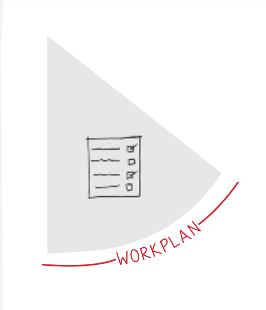

Once the component parts are defined and prioritized, you then have to link each part to a plan for fact gathering and analysis. This workplan and timetable assigns team members to analytic tasks with specific outputs and completion dates. We'll show you best practices in workplanning to move quickly and accurately to solutions. A good workplanning process also includes team norms around generating a diversity of views, use of experts, role-playing, and flattening team hierarchies to achieve better answers. Good team norms and process help us avoid common pitfalls and biases in decision making, including confirmation bias, sunk costs fallacies, and anchoring bias.

STEP 5: CONDUCT CRITICAL ANALYSES

Data gathering and analysis is often the longest step in the process. For speed and simplicity we start with simple heuristics—short cuts or rules of thumb—to get an order of magnitude understanding of each problem component, and to assess priorities quickly. This helps us understand where we need to do more work, and especially when and where to use more complex analytic techniques, including game theory, regression, Monte Carlo simulation, and machine learning. Don't worry! Complex techniques are rarely needed, and when they are, new online analytical tools make them much more accessible than you think. To keep the team on the critical path we make frequent use of one-day answers that summarize our best understanding in the form of situation, observations, and initial conclusions—and team review sessions to pressure-test these hypotheses.

STEP 6: SYNTHESIZE FINDINGS FROM THE ANALYSIS

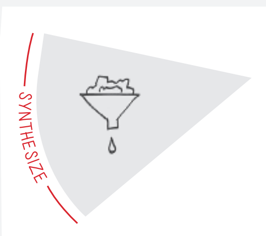

Problem solving doesn't stop at the point of reaching conclusions from individual analyses. Findings have to be assembled into a logical structure to test validity and then synthesized in a way that convinces others that you have a good solution. Great team processes are also important at this stage.

STEP 7: PREPARE A POWERFUL COMMUNICATION

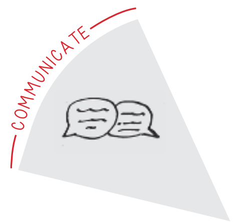

The final step is to develop a storyline from the conclusions that links back to the problem statement and the issues that were defined. A powerful communication will use a governing thought or argument that derives from your refined situation–observation–conclusion logic from earlier stages. This will be supported with your synthesized findings and assembled into component arguments that may follow inductive or deductive logic. It will either lead with action steps, or pose a series of questions that motivate action, depending on audience receptivity.

Prepare for an Avalanche of Trees!

We use logic or issue trees to visualize and disaggregate problems. We employ several types, including hypothesis trees and decision trees, as you will see in the cases we present throughout this book. We learned the power of logic trees at McKinsey and continue to find them essential to good problem solving. Why? Because they do the following:

- Provide a clear visual representation of the problem so that everyone can understand the component parts.
- Done correctly, they are holistic in the sense that everything relevant is captured in the tree.
- Lead to clear hypotheses that can be tested with data and analysis.

Our logic trees are sometimes simple and sometimes highly complex. But they all started on a sketchpad or a whiteboard.

Let's Start with Some Case Studies

To illustrate the bulletproof problem solving process, we chose some case studies that represent classes of problems that many of our readers will face, and that exhibit the power and utility of the process described in detail over the next several chapters.

1. Is Sydney airport capacity adequate for the future?
2. Should I install solar panels on my roof now?
3. Where should I move?
4. Should a start-up raise its prices?
5. Should I support a K–12 school education levy in my town?

These relatively simple cases will outline each of the seven problem solving steps, but with a focus on the use of logic trees to help you represent the problem and break it into manageable parts. Later chapters will go into the fine points of the other steps in more detail and for more complicated problems.

Case 1: Does Sydney Airport Have Adequate Capacity?

When Rob was the lead partner in recruiting for the Australian and New Zealand practice of McKinsey, the consulting firm made the decision to look beyond traditional hires with MBAs to try to attract clever physicists, scientists, lawyers, engineers, and liberal arts graduates. Discussing business cases in interviews put many of the potential hires at a disadvantage. So his recruiting team came up with a non-business case that they called the Sydney Airport case. It is pretty simple, but it is a good way to show the seven-steps method.

All of the candidates had flown into Sydney Airport and were aware of discussions in the newspapers about whether another airport was needed at that time. Sydney Airport has two of the 10 busiest air routes in the world, so this is a real-world example. At the interviews the candidates were given a simple problem definition (step 1 problem definition): "Will Sydney Airport capacity be adequate in the future?" and asked how they would think about that question. The problem statement was bounded around passenger airport capacity, so the candidates didn't have to spend a lot of time on policy factors that might warrant a second airport, such as greater accessibility, safety, or environmental factors like noise, or even alternatives like a very fast train link between major cities. As we'll see later, the boundaries on problem definition are really important to agree on up front.

Candidates would often ask a clarifying question or two and then outline their approach to addressing the issue. So what was Rob's team looking for? They wanted to see if the candidates used a logical structure to help them solve the problem. It's much easier to show the parts of the problem in written form, so we encouraged candidates to use a whiteboard or pad of paper. It is usually a trial and error process to get the breakdown right to solve the problem. This is step 2, problem disaggregation, and Exhibit 1.2 shows a simple first cut.

EXHIBIT 1.2

In this case, the simplest possible way to cleave the problem is to define airport capacity as supply (of landing slots) less demand. You could have a more complicated tree with competition from other ways to get to Sydney (and you might get extra credit for showing how those affect demand), but it probably isn't necessary in this relatively simple case.

A good candidate would dig a little deeper of course. Exhibit 1.3 shows one way of defining airport supply capacity (number of runways, capacity of each runway, and utilization) and demand (Sydney's share of regional demand). In the short term, the number of runways is fixed, and so is runway capacity (defined mostly by aircraft type).

EXHIBIT 1.3

Candidates would typically explain their approach to modeling demand growth by making different assumptions about gross domestic product (GDP) growth, fuel costs, and relative location attractiveness of Sydney relative to other destinations (see Exhibit 1.4).

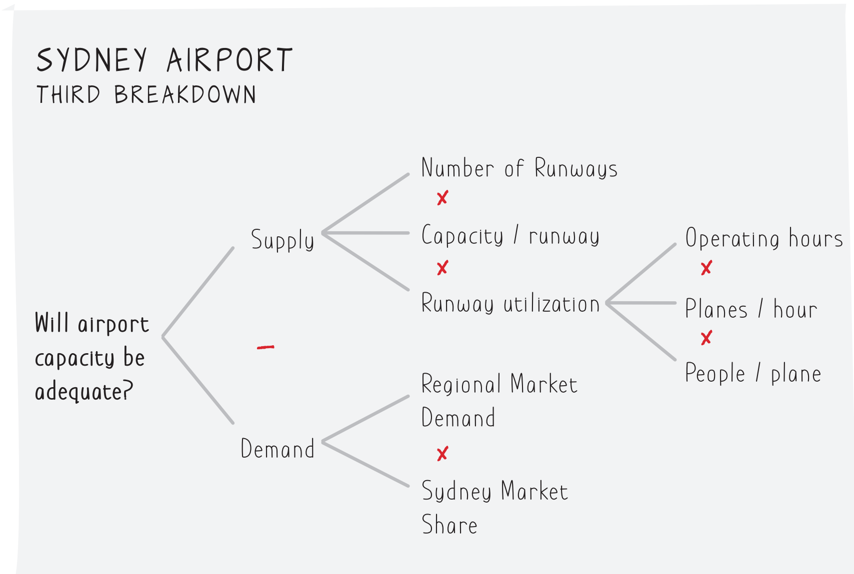

EXHIBIT 1.4

But the most productive approach to this problem is to go deeper into runway utilization, as it is one of the few variables that can be actively managed by transportation planners. Runway utilization is determined by hours of operation, spacing between aircraft movements, and the number of people per airplane. Hours of operation are limited by curfew periods, weather, and maintenance. Thinking about how you could vary these is the heart of steps 4 (workplanning) and 5 (analysis).

The answers Rob most liked were the ones where candidates would say something along these lines:

> Runway utilization is the key so I'd be looking at operating hours, planes per hour, and the people per plane. You probably can't do much with operating hours because there are curfew restrictions

between midnight and 6 a.m. because of the residents nearby. With planes per hour—a core variable for utilization—I'd want to see if they could safely reduce further the time between take-offs and landings. The third factor is the people per plane and that comes down to slot pricing favoring larger planes and policy about light aircraft use at peak hours (steps 6 and 7 synthesis and storytelling)

Good candidates might also propose to raise prices to curtail demand, a tool for airport capacity management, though one that could result in Sydney market share loss, which city economic planners might not embrace.

The branches on this kind of simple logic tree are joined together mathematically, so it is possible to model simple scenarios and show different alternatives by modifying the variables that planners could affect. A really outstanding candidate might show the impact of increasing utilization by 20% on passenger numbers, or employing larger planes.

What actually happened at Sydney Airport? Sydney got a third runway some years later and has managed the impact of significant traffic growth by working on the key variables identified in the case. Despite the current airport authority's opposition, Sydney is to get a second airport in the next decade.

Case 2: Should Rob Install Solar Panels on His Roof Now?

A few years ago Rob thought it might be time to install solar panels at their house in the Australian countryside. Rob and his wife Paula wanted to do something to offset their carbon footprint for some time, but were struggling to make a decision with reducing (and now eliminated) subsidies available from the power company, declining costs of installing solar PV, and questions over the future level of feed-in tariffs (the price at which the electricity company buys *from* you when you generate excess power at home). Was now the right time? He decided to approach it in the way he had learned at McKinsey and started with the hypothesis, "We should install solar PV now." He hadn't reached a conclusion by framing it this way, nor was he setting out to confirm it without regard to the facts. He was using the hypothesis to bring forth the arguments to either disprove it or support it.

Rob felt that the hypothesis would be supported if the following criteria could all be sustained:

- If the payback on the investment was attractive, something less than 10 years.
- If the decline in the cost of panels was slowing down such that he should not wait and make the investment later at substantially lower cost. Rob felt that if solar panel costs were going to continue to decline and be significantly cheaper in three years, he'd consider waiting.
- If the reduction in his CO_2 footprint was material, by which he meant 10% or more (other than air travel he is required to do and can offset independently).

Rob knows that constraining the scope of the problem with clear boundaries makes problem solving more accurate and speedy (step 1).

This kind of problem sounds quite complex at first, a jumble of unfamiliar terms like feed-in tariffs and avoided carbon. A logic tree helped Rob see the structure of his problem in one picture, and helped him break up the analyses into manageable chunks. He started by laying out the reasons and supporting facts that he would need to resolve the issue. You can also think of it this way—for Rob to answer the question affirmatively what would he have to be convinced of? What are the major reasons for going ahead and installing solar panels? Exhibit 1.5 is a first cut of Rob's logic tree (steps 2 and 4).

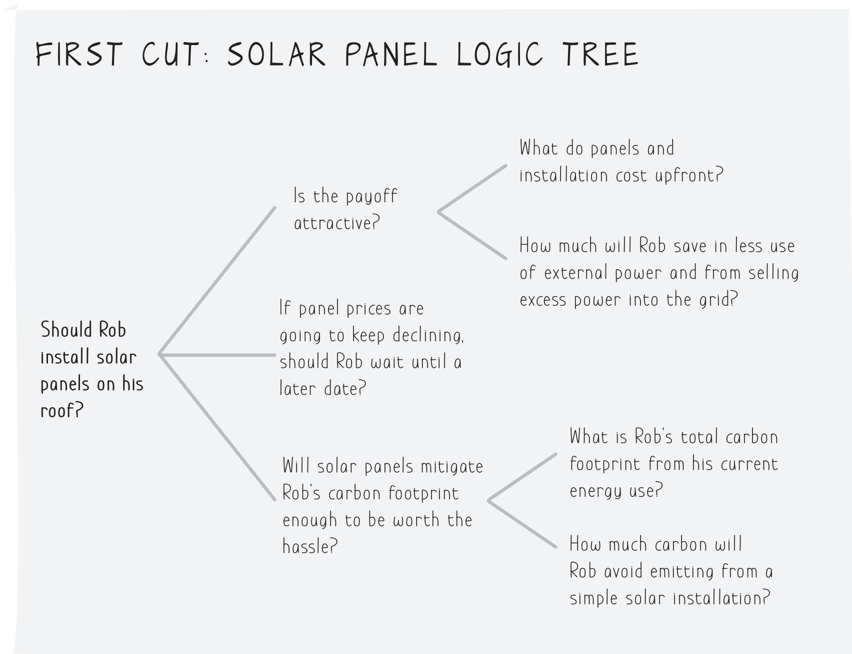

EXHIBIT 1.5

The first part he tackled was payback, because if the economics didn't work, the two other questions didn't need answering. Payback is pretty straightforward: the cost of the installed solar panels and inverter, divided by the annual electrical cost savings. The denominator in this analysis includes both estimating net savings from the installation from avoided electricity charges because he was using his own power, plus income from supplying electricity to the grid via feed-in tariffs. Most of this analysis can be done by online calculators that solar installers offer, once you know the size of the system, roof orientation, solar electric potential, and the efficiency in power generation. Rob simplified the analysis by leaving out battery storage options that add to cost but provide the opportunity to replace peak power charges. With an annual cost savings of around \$1,500 and investment costs of just over \$6,000, payback was attractive at about four years (step 5).

The next question was whether he should make the investment now, or wait, hoping for lower solar panel costs later. Rob was aware that the cost of a watt of PV had fallen almost 30% from 2012 to 2016, and almost 90% from the early days of solar PV. He wasn't sure whether this would continue in the future. With some simple Internet research, Rob learned that declining costs of equipment was still uncertain, but the cost per watt was unlikely to fall by more than 30% for at least the next three years. There is also uncertainty about future feed-in tariffs that have been set to encourage sales of solar PV. This has to be considered against rising retail prices for electricity customers.

At \$1,500 per year, the cost savings lost by waiting would be \$4,500 over three years, so the up-front cost of the solar PV installation would have to fall by 75% to make waiting worthwhile. Rob could have used a net present value analysis where the time value of money is considered rather than a simple payback. But in this case the simple method is fine: He felt comfortable with the four-year payback providing an implied rate of return of 25%. It was worth doing now.

Finally, he wanted to estimate how much of his CO_2 footprint he would reduce by going ahead. This depends on two things—one is what fuel source he is displacing (coal or gas in this case), and the second is the kilowatt hours (kWh) he is generating compared to his electricity use, which he knew from the first step. Rob simplified the analysis by looking at the carbon footprint of the average Australian citizen, and found that the avoided carbon from his little solar project could reduce his footprint by more than 20%. Since the payback as an investment is very solid in this case, Rob really could have pruned off this branch of the tree (step 3) and saved some time—but he and Paula had multiple objectives with this investment.

Whenever you do this kind of analysis, it is worth asking what could go wrong, what are the risks around each part of the thinking?

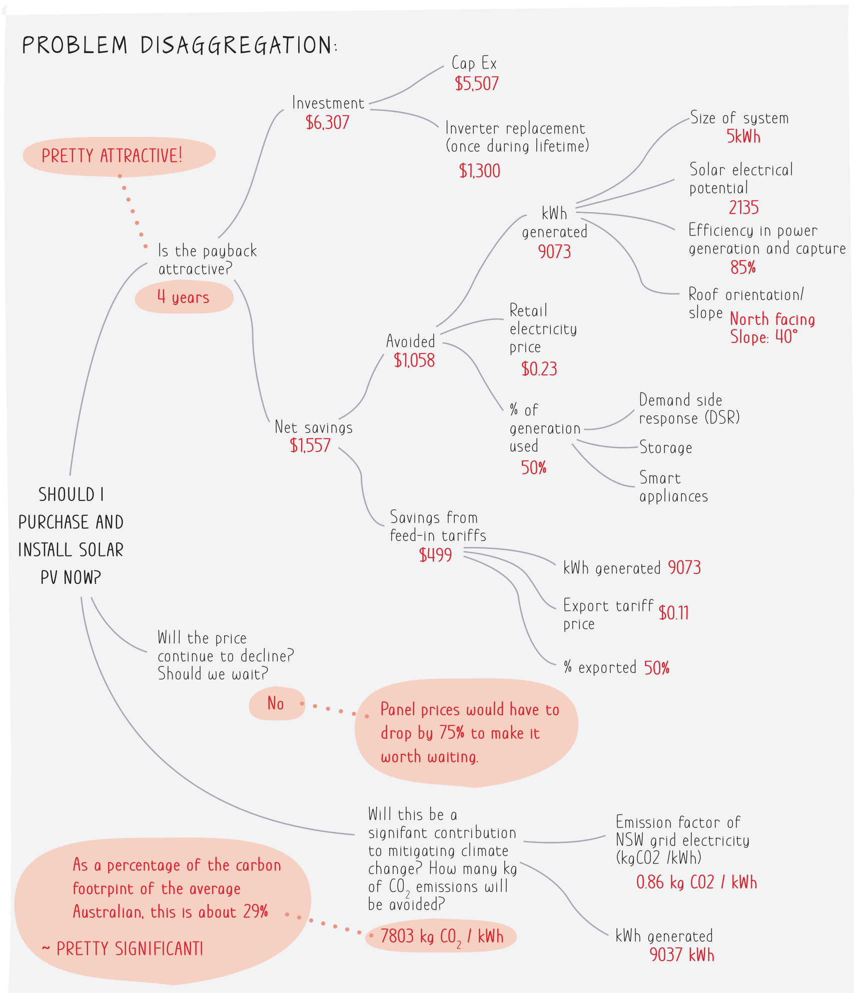

Exhibit 1.6
Assumptions: House type: detached bungalow, Roof orientation: north, Slope: 40°, Suitable roof area: 40m2, Installation size: 5kWh, Shading: none, Numbers for calculations are from 2017, Australia.

In this case there is a chance that the power company would reduce the subsidies for installing solar PV. This can be mitigated by acting quickly. The power company could also reduce the feed-in tariff rate at which it purchased any excess power produced by Rob—and in fact they did that later. But with a four-year payback the exposure is reasonably limited.

The result of Rob's analyses is shown in the more complicated tree shown in Exhibit 1.6.

With only a bit of online research, Rob was able to crack a relatively complicated problem. Rob should install solar panels now. The payback is attractive, and likely cost declines to install later are not enough to offset the savings he could earn now. As a bonus, Rob and Paula were able to reduce their carbon footprint by nearly 30% (steps 6 and 7).

The core of this good result was asking the right questions and disaggregating the problem into straightforward chunks.

Case 3: Where Should I Move?

In the early 2000s Charles was living in Los Angeles. Having recently sold the company he cofounded, his family wanted to move to a small-town environment where there would be more opportunities for recreation and really good schools. They liked the ski towns they had visited, and they had always enjoyed college towns. But how to choose? There are so many variables involved, and it is easy to get it wrong with only impressions from brief visits. Then Charles remembered the factory siting problem he worked on back at Canon in Japan and set up the decision-making effort in a similar way.

The whole family got involved in the problem solving brainstorming, kids included. They started by listing out what mattered to each of them, so their personal factors defined what it meant to be a good place to live. The family agreed on a weighting that favored the school system first, then the natural environment and recreation, and finally characteristics that made for

FAMILY BRAINSTORMING SESSION: WHERE TO LIVE?

ELEMENTS OF A GOOD LIFE	WHAT DOES THAT LOOK LIKE?
Really good schools for the kids	Great teachers Small class sizes Good taxpayer support for education School choices: public, charter, private Graduates get into good colleges
A clean environment and lots to do outside	Water and air quality high A four season climate Lots of sunny days, but sufficient rain Rivers to fish in Great hikes nearby Skiing and mountain biking
A cool, friendly town	A walkable town center Arts, theatre, libraries Not too much traffic Fun coffee shops and good restaurants Do any friends live there? Is it a university town? Is crime a problem?
Can you earn a living?	Cool, small companies Diverse local economy Not too far to West Coast work for Charles

EXHIBIT 1.7

a cool, in-town experience. Charles then added the elements of the ability to earn an income! These were agreed after lively debate with everyone involved (step 1). They planned to use the list to develop a set of towns to visit during family vacations (see Exhibit 1.7).

Charles began the analysis by breaking down the problem into the major elements the family said they valued, then identifying subelements, and finally measurable indicators or variables that captured each subfeature, such as sunny days or a comfort index (defined by temperature and humidity) for climate variables (step 4). It was a little work, but he discovered that most of the data was available online. With the family's input, he put a relative weight next to each variable, to reflect the importance of each element to their final decision.

He developed a tree of around 20 variables and gathered data for about a dozen towns (step 2). The tree he developed is shown in Exhibit 1.8, with the weightings shown in red.

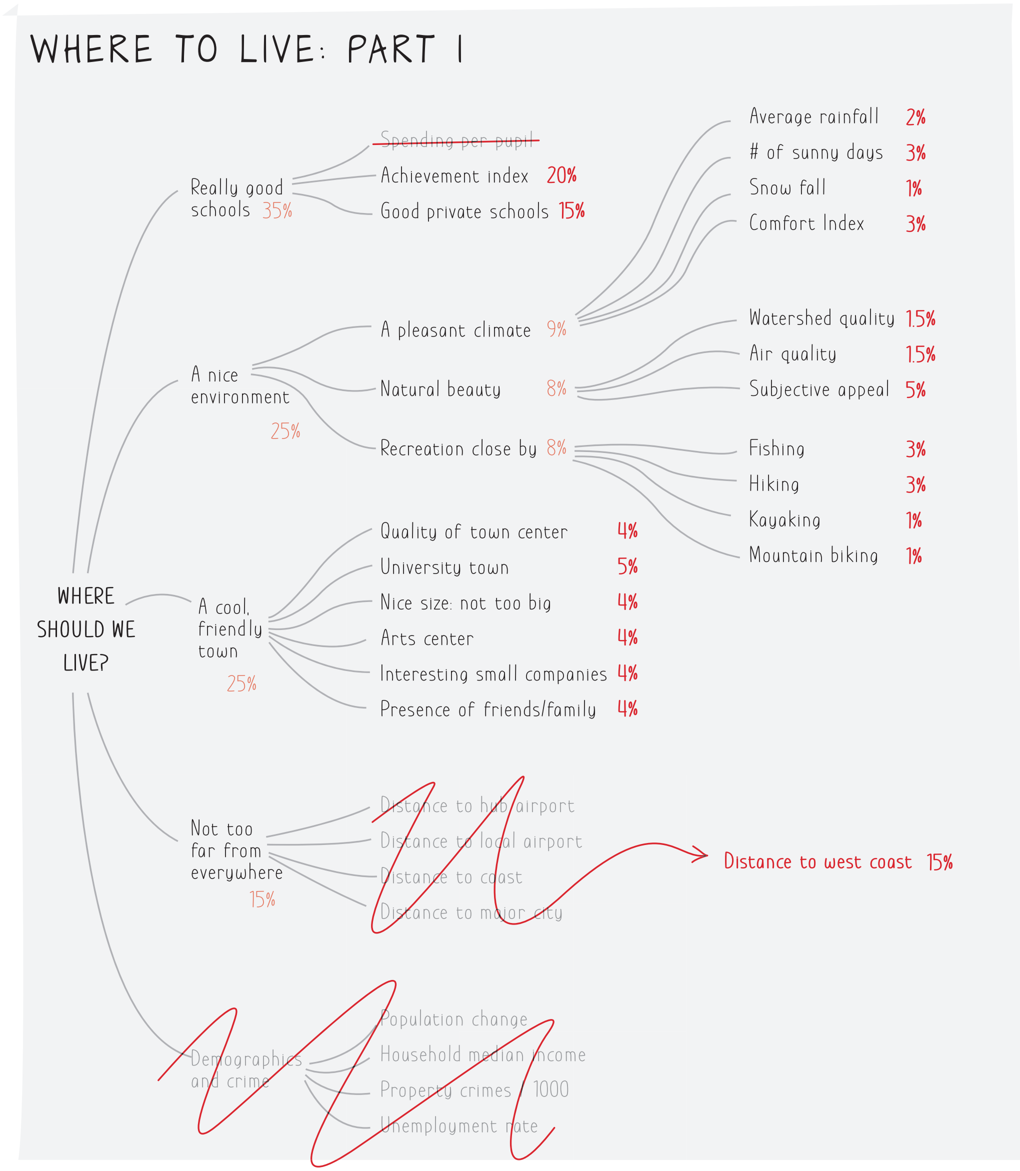

EXHIBIT 1.8

Once Charles gathered the data on a set of small college cities and mountain towns, it became clear that some of the variables were repetitive and others didn't really help distinguish between locations. He pruned his tree to make the analysis simpler and faster. It meant some locations didn't need to be visited. It also revealed that some factors he felt were important around airports and hubs could be encompassed in a single measure of commute time from each town to the West Coast, where most of Charles's work with young companies was located. A variable included early in the analysis was community safety or crime, but that turned out not to be a differentiating factor among the preferred communities, so it was also pruned (step 3).

He converted all the data for each factor to a common scale from 1 to 100, and then applied the weightings. There are various approaches to what is called normalizing data, but they are straightforward and can be found online. As you can see, some of the variables have a positive slope (for example, more days of sunshine is good), and some are negative (for example, longer commute times). So a 100 for Amherst, Massachusetts, on travel time is a negative weight in this case. If we were to get fancy, you could have more complex line shapes to the weighting for variables like rainfall, where you want some rain, but not too much! Exhibit 1.9 shows Charles's analysis (step 5).

In this case, the family was able to reach a conclusion to choose Ketchum, Idaho (steps 6 and 7). They all agreed on the choice because they had agreed on the factors that made for a good location and how to weigh or trade them off. There was a big trade-off in moving to Ketchum, Idaho that Charles was prepared to accept—the commute time to the West Coast for business was longer than other locations.

We sometimes hear a criticism with systematic problem solving like this example that you went through a phony process to prove what you had in mind from the beginning. In this case it isn't phony at all: Ketchum wasn't even on the list of the initial towns under consideration and was only added after a trip

AND THE BEST TOWN IS....

NORMALIZED TABLE

	Weight	Healdsburg, CA	Fort Bragg, CA	Bend, OR	Victoria, BC	Boulder, CO	Amherst, MA	Steamboat, CO	Ketchum, ID
Avg rainfall per annum	2%	71	24	0	58	82	29	100	92
# of sunny days	3%	100	91	39	0	74	35	79	52
Snowfall (inches)	1%	0	0	4	4	36	36	100	62
Watershed quality (best=100)	2%	0	7	45	41	100	35	94	57
Air quality (best=100)	2%	0	33	36	33	38	73	82	100
Comfort index (best=100)	3%	30	46	40	46	100	0	86	80
Air time to West Coast (worst=100)	15%	0	17	33	17	33	100	50	58
Availability of private schools (index)	15%	100	0	17	67	83	83	17	83
Achievement index	20%	46	0	38	65	97	100	78	76
Presence of friends/family	4%	100	0	0	0	0	0	0	100
University/college town	5%	0	0	100	100	100	100	0	50
Arts center	4%	50	25	0	100	100	75	0	75
Quality of town center	4%	67	33	0	67	67	50	50	100
Physical beauty of area	5%	25	50	0	50	50	0	100	100
Interesting small companies	4%	60	0	40	60	100	80	0	20
Size	4%	51	100	90	21	0	94	98	100
Recreation	8%	50	40	60	50	70	20	60	90
	100%								
TOTAL WEIGHTED INDEXED SCORE		48	20	34	51	70	70	52	76

EXHIBIT 1.9

to visit friends. But it is worth pointing out some risks in this analysis too. Boulder and Amherst are really close in total score to Ketchum, and quite small changes in how Charles ranked factors such as recreation or quality of town center—which are subjective—could have a big impact on the numerical conclusion. In this case the family could mitigate the risk by visiting each town and testing their feelings about the quantitative variables firsthand.

The where-to-live case illustrates how you can start with a simple list of issues or elements that are related to your problem statement, disaggregate the elements further into indicator variables, then finally add concrete measures and weights. The rest is straightforward arithmetic based on a considered ranking of features. This type of tree and analysis approach has applicability to many choice problems. Charles and Rob have used it to assess what apartment to buy, what employer to join—and, of course, where to put your factory.

Case 4: Making Pricing Decisions in a Start-up Company

In the past few years, one of Charles's friends started a company that makes an accessory for pick-up trucks that has a unique and clever design. The company, which we'll call Truckgear, sells around 10,000 units a year, a number that is growing quickly. It is at break-even on a cash basis (cash basis means not taking into account the accounting charge for depreciating assets). Charles invested in the company and helps devise its strategy.

Start-up companies face big and complex problems early on in the process and, compared to larger companies, they have limited cash resources and team members to address them. Truckgear had to make decisions on whether it should own its own manufacturing plant, which market segments to compete in (there are new and used truck segments and several sales channels to each), whether it should have its own sales force, how much to spend on marketing, and most fundamentally, how fast to grow given limited cash? No wonder start-up teams hardly sleep!

Recently the company had a big decision to make: Should it raise its prices (step 1)? It had held its initial pricing of around $550 for three years. Materials and manufacturing costs had increased as the product features were improved, crimping its margins and lowering the cash generated per unit. Obviously in young companies cash is even more critical than in established ones, as the sources of external financing are significantly fewer. The dilemma Truckgear faced was this—if the marketplace reacts negatively to the price increase, Truckgear growth would slow and perhaps even drop in unit sales.

There is no perfect answer to this kind of question, but we employed a particular kind of logic structure to assess it, a profit lever tree (step 2). We wanted to hone in on the key factors around the decision, and this kind of tree is mathematically complete, so we could use it to model different assumptions.

Exhibit 1.10 is a simple version of this kind of tree.

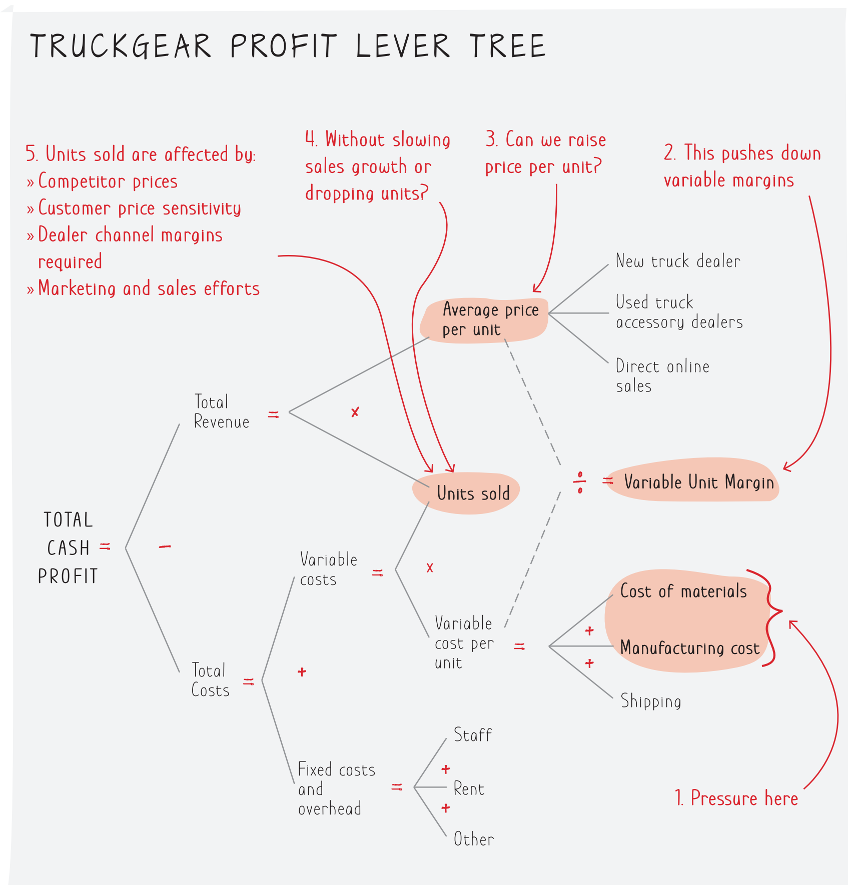

EXHIBIT 1.10

You can see how the tree makes Truckgear's problem visual: Pressure on costs pushes down variable margin per unit—can the company increase unit prices without slowing sales growth or even dropping volume?

Exhibit 1.11 displays the numbers for Truckgear.

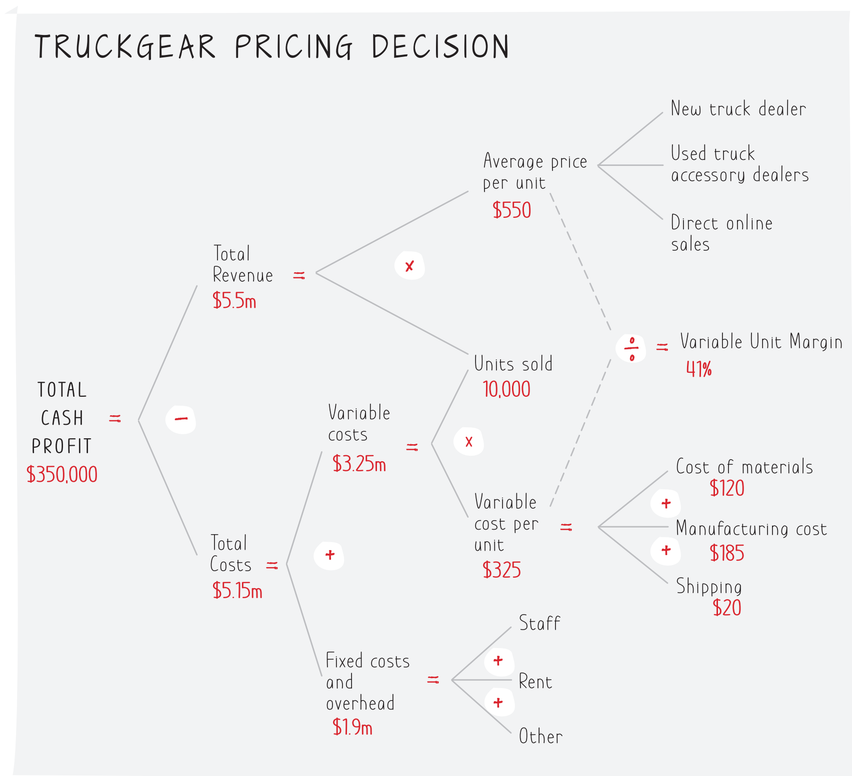

EXHIBIT 1.11

If the company could hold its current unit sales, a price increase of 7% would yield a $385,000 improvement in cash profitability, a substantial increase that could help fund additional marketing and sales programs. But you can also see that it would only take a drop in unit sales of 650 to neutralize the benefit of the price increase (step 5). What to do?

Whether a price increase leads to a loss in total cash profit (or, less seriously, slowing of growth) depends on competitor pricing, customer price sensitivity (which economists call price elasticity), whether the third-party dealer sales channels will absorb some of the price increase by accepting lower margins, as well as marketing and sales efforts. The company initiated a large phone survey of its recent customers and determined that:

- The largest volume customer segments were not sensitive to a modest cost-based increase.
- The competitor products were roughly comparable in price and quite different in functional characteristics.
- Dealers were not willing to reduce their margins to accommodate the price increase.

The company also evaluated whether it could achieve the same result by reducing its fixed overheads or taking manufacturing in-house. With few costs other than a lean staff and rent, the first was not an option. With limited current cash resources, investing in its own extremely expensive manufacturing presses and assembly also didn't make sense (step 3). On balance a small price increase to restore unit margins was worth the risk (step 6 and 7).

This kind of financial tree is particularly useful for solving problems that involve monetary trade-offs of alternative strategies. You can use it to track almost any kind of business problem. We'll show a number of more sophisticated versions in later chapters.

Case 5: Should Charles Support the Local School Levy?

In Charles's former hometown in Idaho, public education funding is supported principally by real estate taxes, which are levied as a percentage of property value each year, and by state sales taxes. When a local school board has bigger strategic investments to fund, it seeks approval through a vote by taxpayers for an additional levy to pay down a bond secured for that purpose. In the late 2000s the Blaine County School Board proposed a bond of more than $50 million to support a large set of investments across the county. With a population of only 20,000 people, depending on the size of your property, the levy could cost a homeowner thousands of additional dollars per year.

As taxpayers and citizens, people face this kind of decision all the time—a spending levy, a one-time referendum, or whether to support a candidate pitching a new state or national policy. The problems seem complex and are often swayed by partisan debate rather than informed by straightforward problem solving that any of us can do.

Charles was aware of press stories that K–12 education in the United States is lagging behind its peers globally, and so was in principle supportive of additional taxes on his house if it really improved local education. But he wanted to know whether voting for the school bond would really have an impact on closing the gap locally (Idaho ranked in the bottom half of US states in school-testing results).

Charles is no expert on education policy. He started by asking a simple question: What is the key problem with K–12 education in the US versus its peers, and does the proposed school levy address these issues (step one)? He knew the United States spends a lot per student in most jurisdictions, among the highest in the world, and he knew overall results were middle of the pack compared to peer countries. So he asked, is the problem:

- Per-student funding?
- IQ or demographics?
- Teachers and schools?

His research showed that the answer is mostly with *teachers and schools*. There was nothing in the data to suggest either that funding levels per student (higher in the United States than most countries) or student IQ levels (comparable to other countries) were the reason for poor student outcomes in the United States (step 3). Study after study shows that student outcomes on international test scores vary widely and are best explained by teacher characteristics and the school environment. Charles then proceeded to research which factors about teachers and schools had most impact on student outcomes and found four, ranked in approximate importance (step 4):

- Teacher numbers and classroom size
- Teacher quality (education, experience, training) and compensation
- School environment and facilities
- Technology

Next, he overlaid the extent to which the proposed bond would provide funding against these factors. It turned out that the funding was mainly directed at clean energy and school facilities, low-impact factors. There was little allocation of funding to the higher ranked factors, especially related to teacher recruitment, pay, and training (step 5). Charles decided not to support the bond on these grounds. Exhibit 1.12 shows what the analysis looks like in a decision tree format.

This analysis illustrates bulletproof problem solving for a societal-level policy problem: It takes a real issue, frames a question followed by a series of sub-questions, and guides research and analysis. Charles collected facts about education performance and the details of the school bond planned spending allocation. It took him only a few hours of framing and online research and helped move him from an emotional decision ("I support education and am concerned about local school performance,") to a reasoned one ("I can't support this levy, given how it is allocated.") (steps 6 and 7).

On the third attempted school levy a few years later, the measure contained funding for early childhood education and more teacher training—and the vote passed.

The next chapters will examine each of the seven steps to bulletproof problem solving in more detail, introducing more complex problems, and more sophisticated approaches to solving them.

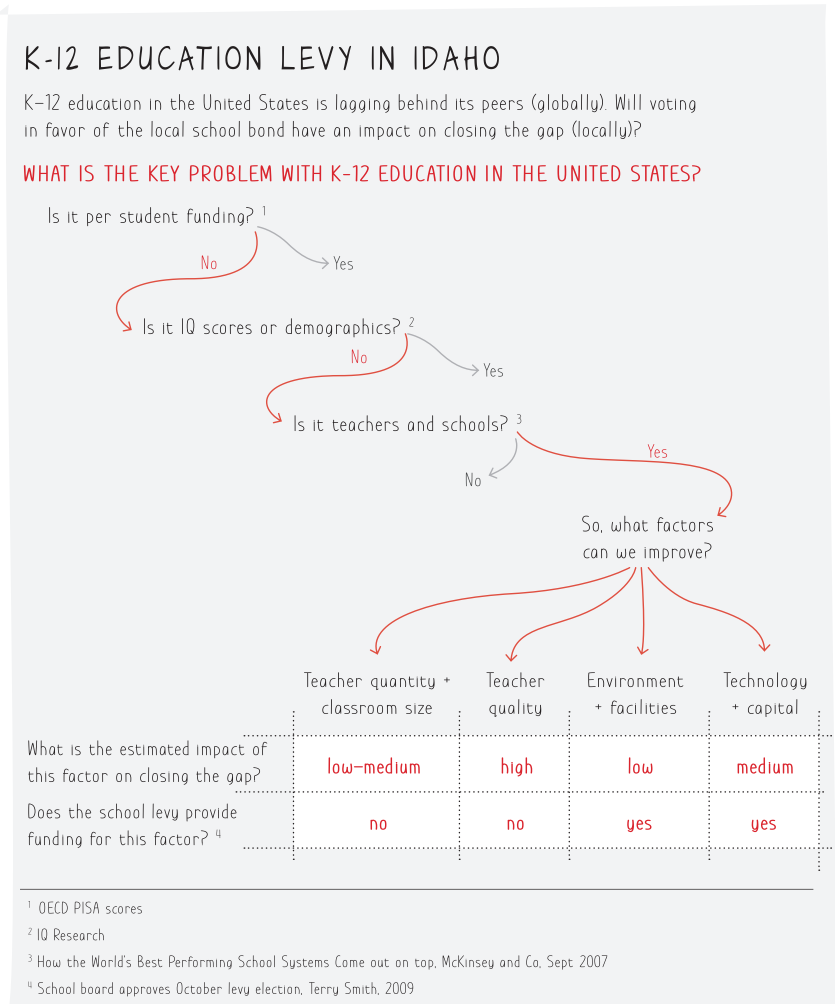

EXHIBIT 1.12

Chapter 1 Takeaways

- Good problem solving is a process, not a quick mental calculation or a logical deduction. It applies to even highly complex problems, like where to put factories around the world, as Charles experienced as an intern at Canon.
- The problem solving process we outline in the book is a cycle of seven steps that you work through; each step is important and lots of mistakes result from skipping steps.

- The most important step is to disaggregate the problem in a logical way into component parts, so you can isolate the most important analyses; logic trees are our prime work tool—they make it easy to see the structure of the problem.
- Prioritizing analyses is essential so that you avoid working on parts of the problem that don't contribute much to your answer—we call this being on the critical path.
- A well-defined work plan is needed to allocate analysis to team members and time frames for completion (in this chapter not much work planning was required because the problems were simple).
- How you go about analysis, using simple tools or sophisticated ones, is important to problem solving success; we always start with simple estimates and heuristics or rules of thumb.
- Problem solving isn't over until you can synthesize the results of your analysis and tell a story that convinces someone to act.
- We use different types of logic trees for different problems; in this chapter we have shown mathematically complete deductive logic trees for business problems, weighted factor analysis for making decisions, and decision trees for walking through complex choices.

Problems to Try on Your Own

1. Individual: Show the logic tree for whether you should make a job change—try a decision tree or a factor-weight tree (like where to live, but with job characteristics that are important to you); or try it for a new house or apartment, with weighted characteristics you would look for, and compare them against listings near you.
2. Business: Lay out the profit tree for your business or social enterprise to the third breakdown.
3. Societal: Draw a logic tree for Britain's decision whether or not to leave the European Union (the Brexit decision).