The Illusion of Certainty

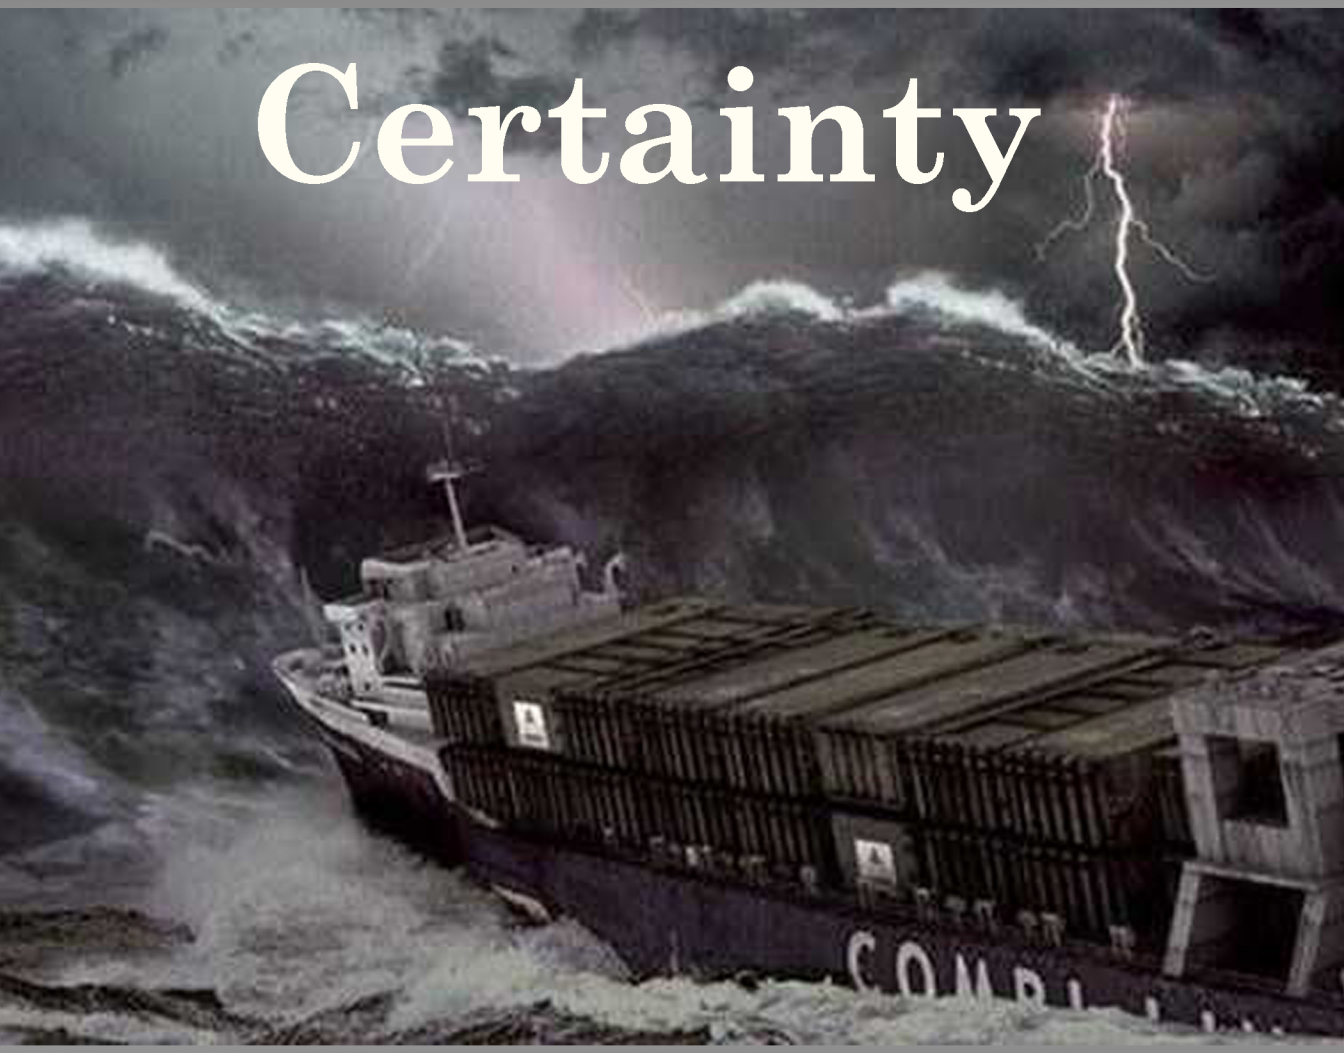

ELI SCHRAGENHEIM
and
H. WILLIAM DETTMER

THE ILLUSION OF CERTAINTY

Eli Schragenheim

and

H. William Dettmer

BOOKS BY ELI SCHRAGENHEIM AND BILL DETTMER

Schragenheim

- Management Dilemmas: The Theory of Constraints Approach to Problem Identification and Solutions (1999)

Schragenheim and Dettmer

- Manufacturing at Warp Speed: Optimizing Supply Management Financial Performance (2001)
- Supply Chain Management at Warp Speed, Integrating the System from End to End (2009)

Schragenheim, Camp, and Surace

- Throughput Economics, Making Good Management Decisions (2019)

Dettmer

- Goldratt's Theory of Constraints (1996)
- Breaking the Constraints to World-Class Performance (1998)
- Strategic Navigation (2002)
- The Logical Thinking Process (2006)
- He Said, She Said (2013)
- The Logical Thinking Process: Executive Summary (2016)
- Systems Thinking and Other Dangerous Habits (2022)
- Reminiscences (2023)
- Kingdom Come (2025)

The Illusion of Certainty

How to Deal Effectively with Uncertainty

> The combat leader almost never has the full picture or a clear and certain understanding of the enemy's actions or reactions, nor even the knowledge of the immediate consequences for momentary decisions . . .
>
> [L]eaders cannot be paralyzed by fear. That results in inaction. It is critical for leaders to **act decisively amid uncertainty**; to make the best decisions they can **based on only the immediate information available** . . . There is no 100 percent right solution. The picture is never complete. Leaders must be comfortable with this and be able to **make decisions promptly**, then be ready to **adjust those decisions quickly** based on evolving situations and new information. Intelligence gathering and research are important, but they must be employed with realistic expectations and must not impede **swift decision making** that is often the **difference between victory and defeat**. Waiting for the 100 percent right and certain solution leads to delay, indecision, and an inability to execute. Leaders must be prepared to make an educated guess based on previous experience, knowledge of how the [opponent] operates, likely outcomes, and whatever intelligence is available in the immediate moment.
>
> This "incomplete picture" principle . . . applies to virtually every aspect of our individual lives, such as personal health-care decisions or whether or not to evacuate from the predicted path of a major storm. It particularly applies to **leadership and decision making in business**. While business leaders may not generally face life or death situations, they are certainly **under intense pressure**. With capital at risk, markets in flux, and competitors actively working to outmaneuver opponents, professional careers and paychecks are at stake. Outcomes are never certain; success never guaranteed. Even so, **business leaders must be comfortable in the chaos and act decisively amid such uncertainty**.
>
> — Willink and Babin
> *Extreme Ownership [1]*

PART I. UNCERTAINTY: THE CHALLENGE IT PRESENTS

Introduction

As the quotation above so persuasively states, we all deal with uncertainty in our lives, both personal and professional. Most times we recognize this. Sometimes we don't. But dealing with uncertainty can be a tough job, and whether they recognize it or not, everyone faces uncertainty every day, all the day long.

No matter how clever we are, we could regret any move we might make, calling the decision a "mistake." Even trivial daily decisions such as choosing what to eat for lunch could lead to negative consequences. But most of us are not terribly inconvenienced by small deviations from what we expect in the short-term.

However, while we learn how to deal with daily uncertainty, getting over small disappointments, as humans we are often distressed by more threatening uncertainty in the future, especially the long-term outcomes. This is uncertainty that causes us to fear for our lives, and for our loved ones. Still, most of us learn to live even with this burden, continu-

ing to evaluate our current situations and deciding how to protect our lives from the bigger unknowns.

When it comes to managing organizations, however, we have a different attitude toward uncertainty, which is the key focus of this article. When managers and executives evaluate their decisions on behalf of their organizations, another factor enters the picture*: the fear of how others might judge the decision based on after-the-fact results*. This is a common fear. The organizational culture may lead people to ignore uncertainty. Moreover, leaders frequently tend to evade making problematic decisions.

When there is no way to delay a decision, the clear choice is to choose the course that seems safer, regardless of the potential gain that might have been achieved. In other words, when evaluating new initiatives and business opportunities, the personal fear of negatives results, including those with very limited real damage to the organization, often produces too conservative a strategy. Ironically, this might actually open the door to new threats to the organization.

Another difference between organizational and personal decisions is that most decisions within an organization affect other decisions, including ones not necessarily made by the same decision makers. For instance, when a client requests a product or service, a sequence of actions within the organization is created. Several different employees, perhaps in different departments, are likely to be involved in reaching an organizational objective, including decisions on how each part of the company's business process should be done.

A dependent sequence of tasks, carried out by multiple employees, requires considering *common and expected uncertainty*, in order to ensure smooth and stable fulfillment of the organization's objectives, including its commitments to its customers.

Organizations must plan for long-term as well as short-term objectives. However, uncertainty often permeates every detail in the plan, forcing the employees in charge of the execution to re-evaluate the situation and introduce changes. By confronting uncertainty, both during planning and execution, the odds of achieving all, or most, of the key objectives of the original plan increase substantially.

We all need to engage uncertainty. The challenge of doing so requires tools and heuristics that can make for better decisions. Evading the topic is not productive. Let's start with the effects of uncertainty on the individuals. Then we'll move on to the impact on managing organizations.

Recognizing Uncertainty and its Ramifications

Let's say you take a flight from New York to London. The scheduled landing time is 16:35. Would you be extremely surprised if the flight didn't actually arrive until 17:01? Would you be highly surprised if the flight, after several announced delays, was canceled and rescheduled for two days later? There could be significant consequences of such a delay to your trip. How many meetings or other actions did you synchronize with the originally planned arrival time? Would you prefer not to plan ahead and instead react to the delay in real time?

We all understand that reality comes with a considerable amount of uncertainty. Living in a state of uncertainty can be frightening. We can't always know what's going to happen, and most of us don't know how to deal with uncertainty, thus improving the quality and duration of our lives.

The real problem with uncertainty is that any decision might negatively impact us, sometimes very severely. Nevertheless, we *still* have to make daily decisions under conditions of uncertainty.

Living with uncertainty can create fear and tension. This can drive people to a couple of behaviors that can result in considerable "unpleasantness."

- **Relying on superstitious beliefs** that promise to influence, or even know *a priori*, what's going to happen. For instance, going to a fortune teller, believing in our sixth sense to see the future, or praying to God while rolling the dice.
- **Ignoring the uncertainty in order to reduce the fear.** When we ought to have a frightening medical test, we might "forget" to actually take the test. Politicians and managers typically state future predictions and concepts with perfect confidence that totally ignores the possibility for any deviation.

Mitigating Damage

A practical common tool for helping us to minimize potential loss is insurance. We often pay a considerable amount of money, hoping it will mitigate the financial effects of a possible bad occurrence. Insurance doesn't protect from the occurrence of a frightening event itself; rather, it compensates us for the financial loss that results.

Home appliance warranties are another example. Who hasn't used such insurance to protect against the uncertain costs of simpler and much less damaging incidents, such as paying a technician to repair a refrigerator? Is this wise? No matter how much such a technician charges, the damage is usually not so high as to justify paying so much more for insurance over the long term, just to be protected against minimal damage.

More on the Role of Fear in Decision Making

Fear is a strong negative emotion. We prefer to avoid it, if at all possible. But in itself, fear can add value by forcing us to exercise caution and take preventive steps to avoid loss or unpleasantness in the future.

For instance, many of us don't invest all of our money on just *one* promising opportunity. We do this to protect ourselves from a financial disaster. Understanding the uncertain risk prompts us to distribute our money among many smaller, but quite different, investments. Each might only deliver a moderate gain but with much less risk of a loss. Another possible reaction to the fear is the "ostrich approach"— simply ignoring it, pretending there is no risk. This is simply gambling. The refusal to face reality usually doesn't end well.

Understanding the role of fear raises another issue. Besides the direct good or bad consequences of a decision, there could be also ramifications because of the judgments of others after the fact, whether the decision was "right" or "wrong." How might this perception affect the image or reputation of a decision-maker?

Fear from criticism after the fact, whether justified or not, is typical in organizations, where the blame for an undesirable outcome of a decision can be critical to the decision maker. Just think of an initiative that seemed very promising when it was evaluated and the decision to go ahead was made. Let's assume the initiative eventually didn't deliver on its promise, and whoever suggested the initiative competes with others for promotion. The failure, maybe even communicated by the competitors as a "major failure," may prevent the initiator's promotion.

Was the decision to go ahead with the initiative a "major mistake?" Does the fact that the expected results didn't materialize make it clear that it was a mistake? What if, when objectively evaluated at the time, there was a good chance of a major success? (The three consecutive failures of Elon Musk's Falcon 1 rocket booster come to mind.) Such initiatives are subject to very high uncertainty. The failure might have been just bad luck, or the accumulation of uncertain factors.

Here's a related question: At the time of proceeding with the initiative, did anyone raise the possibility of failure and analyze it? We should expect such an analysis, but it seems not to be a very common practice. Thus, the fear of unjust after-the-fact judgment [2] can kill innovative ideas, because people fear negative, possibly unfair, judgment after the fact.

In mathematics a detailed *theory of probability* arose. This was an attempt to predict the likelihood would be for one particular outcome versus another. Still, probability theory can't determine what the outcomes of specific cases might be.

Statistical analysis, based on probability theory, provided a means to learn from past occurrences, and to be able to predict the global impact of events. One common example is forecasting the results of an election based on surveying a relatively small sample of prospective voters.

There is a pressing need to translate general mathematics into simple heuristics for decision makers in specific cases, where most of the time the probabilities are unknown.

Part of the overall problem of dealing with uncertainty is that we're all subject to many biases. These biases often push us to make decisions that result in clearly undesirable consequences. Here are a couple of examples.

People hate losing money more than they like the gains of winning. This implies that the actual financial loss or gain might not be the decisive factor for many people. Consider a lottery in which one invests $100 and either loses it or wins $200 (including the original $100, so there is a net gain of $100). If the odds were 50%-50%, as they would be in flipping a coin, most people would refuse to participate. Even if the lottery offers a payback of $210 against the $100 bet (with the same 50% probability of winning), the majority would still reject the lottery, even though, according to pure economic calculations, they should accept (and even repeat the lottery gamble many times).

Since uncertainty cannot, in most real-world cases, be expressed with precise probabilities, most people evaluate the uncertainty according to their emotions. Consequently, some people vastly exaggerate the optimistic view of their odds, while others exaggerate the level of risk.

Fear and biases impact all our decisions, but naturally decisions that clearly have a significant impact are much more difficult to make. Another factor that bears on decision making is the difference between decisions that are largely limited to our own lives and families, and decisions we must make as a part of working for an organization. There are two key differences between personal and organizational decisions.

Personal Versus Organizational Decisions

The first difference relates to the goal of the organization. In this situation, the decision is supposed to serve the goal. Most of the time there are performance measurements to ascertain whether the decision truly improved performance. The universal objective of all

managers in an organization is (or should be) to make decisions that support better performance relative to the goal.

This is not the case with the personal decisions made by individuals. Those decisions concern their lives: whom to marry, what to study, which employment offer to accept, which restaurant to choose for lunch, and what TV program to watch. Most individuals don't have *one* clearly stated goal. They may have many different things that they value. So most individuals don't know how to assess how good their decisions are. Most don't even think about it until after the fact, especially when their decision turns out to produce less than desirable outcomes.

Uncertainty in Organizational Decisions

Now, when it comes to managing organizations, we have a different attitude toward uncertainty, which is the key focus of this paper. Which brings us to the second key difference between individual and organizational decision making: organization leaders' attitude toward uncertainty. Mostly, they just ignore it.

When managers, executives, and even lower-level supervisors assess the organizational decisions they must make, they have two very different concerns. First, how will the decision affect the performance of the organization? And second, how will the decision be judged within the organization, based on subsequent results?

While most people recognize that the environment is uncertain, and perhaps even consider it in their personal decisions, when it comes to decisions within an organization, fear plays a much larger role. Every manager in an organization must make decisions, for which the criterion *should be* maximizing the achievement of the goal of the organization. Once the actual outcomes of decisions are known, it's fairly easy to determine whether a specific decision truly advanced the goal. Actually, in most real-world cases the net impact of a particular move on the bottom line is not straightforward. In fact, determining the net contribution of just one decision, when so many other factors influenced the outcome, is open to debate — and manipulation.

It's easy to see this kind of after-the-fact judgment as unfair criticism, especially when it ignores the uncertainty at the time the decision was made.

A Basketball Example

Consider a decision basketball coaches often make at critical moments during an important game: to substitute one player with another. This is a common example of a decision made under uncertainty, as the success or failure of the move depends on many real-time factors: How does the entering player feel? How will the opposing team react to the new player? For the coach this is a highly uncertain decision. Critics of the game would describe the move as "brilliant," or "foolish," based on the actual score ultimately achieved, without considering what the coach knew, or didn't know at the time of the decision.

Assessing Personnel Performance

In most organizations leaders evaluate the performance of individual employees, including managers and executives. This practice is deeply embedded within the underlying culture of most organizations. What motivates this need for personal assessment? It's that "the system" needs to identify those who don't perform acceptably, as well as those who excel. In order to assess personal performance, management typically defines specific

"targets" that employees are expected to achieve.

This use of such personal performance measurements motivates employees to try to set targets low enough so that, even in the face of situational variation, they'll be confident that they can meet these targets. In practicality, this means that while targets are met most of the time, only seldom they are outperformed, lest top management set higher targets. (Today's exceptional performance becomes tomorrow's standard.)

This culture actually creates distrust in employees. It also imposes extraordinary pressure when decisions that might produce a minimal loss (but also might result in a very high gain) are considered. Consequently, many individuals forego decision options they might otherwise choose, in fear of after-the-fact criticism that fails to consider the uncertainty involved.

In practice, this culture of distrust and judgment-after-the-fact produces an organizational tendency to ignore uncertainty. Why? Because it becomes difficult, if not impossible, to judge how good (or lackluster) an employee's true performance is.

What Happens When Uncertainty Is Ignored?

A typical example of ignoring uncertainty is widespread reliance on single-number discrete forecasts of future sales. Any rational forecast should include not just the quantitative average (a single number), but also a reasonable deviation from that number. The fact that most organizations use just single-number forecasts is evidence of the illusion of certainty.

There is a tendency to ignore uncertainty and evade making problematic decisions. When there is no way to delay a decision, the clear choice is to choose the course that seems safer, no matter what potential gain that might have been achieved. In other words, when evaluating new initiatives, business opportunities, or organizational changes, the personal fear of negative results, even those with very limited possible real damage to the organization, leads an overly conservative a strategy. Ironically, this opens the way for new threats to the organization.

Organizations typically plan for long-term objectives as well as for the short-term. A plan requires many individual decisions regarding different stages, inputs or ingredients. All such decisions together are expected to lead to the achievement of the objective. But uncertainty typically crops up in the execution of every detail in the plan. This forces the employees in charge of the execution to re-evaluate the situation and introduce changes, which may well impact the timely and quality of the desired objective.

Inability to predict reliably the timely and quality achievement of critical objectives can be vastly improved with a better understanding of uncertainty and how to deal with it. Quality issues have long been addressed by Total Quality Management (TQM). Specific quality deficiencies are resolved, significantly reducing quality deviations.

But more effectively accommodating all other types of uncertainty, both during planning and execution, would improve the odds of achieving all, or most, of the key objectives in the original plan.

Ultimately, we all need to address the challenge of accommodating uncertainty effectively. Doing so requires useful tools and heuristics that can lead to much better decisions, rather than trying to evade the topic. Let's look at a couple of examples of uncertainty in action, one personal (Case #1) and one organizational (Case #2).

The role of common and expected uncertainty is very big in organizations, because their processes are exposed to uncertainty at every step. Eventually the clients of the organizations are the ones that suffer from this uncertainty.

This situation described in the sidebar on car restoration (Case #1) is probably not rare. Delays such as this likely happen often on jobs as large as this one. To that extent, one might say that such delays might be *expected*. And the length of those delays can't be predicted with any certainty. Moreover, the incidents that caused the delays are not *uncommon*. People get sick and injured. Employees decide to leave for other employment. And skilled workers are often difficult to find and hire.

In summary, we can consider the project described in Case #1 to be afflicted by *common and expected uncertainty*, where the overall effect is the accumulation of all the smaller incidents. The accumulation effect is what makes common and expected uncertainty truly damaging. But most practices for managing uncertainty ignore that effect.

Independent uncertain incidents usually tend to adversely affect the overall outcome. *We typically respond to this kind of uncertainty by spreading out our investments to reduce the overall risk.* Still, the effects of the various incidents often accumulate because of interdependencies that might be difficult to observe.

The Decision Continuum

When someone undertakes any kind of decision, there are three circumstances under which they might be working: *certainty, risk,* and *uncertainty*. (**Figure 1**)

Naturally, anyone would prefer to be deciding under conditions of *certainty*. There would be little possibility of error. Unfortunately, this is rarely the case for most decisions.

Decision under *risk* implies that the outcome (or outcomes) of a particular decision can be estimated with some degree of confidence. In other words, probabilities can be assigned. These may be precise mathematical probabilities (as in gambling in a casino), but more likely they can only be estimated with varying degrees of confidence.

Decisions under conditions of *uncertainty* are the most difficult — and, unfortunately, the most common. These are the kinds of decisions that many people face daily.

The Three Categories of Decisions Under Uncertainty

There are three types of uncertainty: *common and expected, emerging opportunities and threats,* and potentially *catastrophic events*. (See **Figure 2**)

Dealing with uncertainty requires the use of specific tools to predict the possible consequences of a decision. Note, these are categories of *decisions* — not categories of the uncertainty *causes*, or the specific type of uncertainty such as variability, uncertainty, complexity, and ambiguity (VUCA). The objective is to make better decisions under conditions of uncertainty; hence, the categories below focus on decisions under uncertainty.

These categories and their tools are especially critical for managerial decisions. Most organizational cultures strive for optimization, which is practically impossible due to uncertainty. Note that the academic definition of uncertainty specifies that the probabilities (or the parameters of the stochastic equation) are unknown. This description fits the majority of practical cases in reality. The three possible categories below encompass almost all cases of decision uncertainty.

CASE #1: COMMON AND EXPECTED UNCERTAINTY

In 2023, one of the authors experienced first-hand the vagaries of uncertainty. Bill Dettmer engaged a medium-sized automotive repair and body shop to work on his car. Here is his account.

In 2020, I inherited a classic car: a 2001 Cadillac DeVille that had several after-market trim modifications and additions, including a canvas top that resembled a convertible top. From a superficial look at both the exterior and interior, the car appeared to be in superb condition, and it garnered positive comments from other drivers when I drove it around town.

But I later discovered the hidden problem: rust. A local collision repair shop manager pointed out that there was body rot (visible only from the inside or underneath) in the fenders and doors, and worst of all, the rocker panels (which are part of the frame). I had realized that this might be a problem, just not as bad a problem as it turned out to be.

After a concerted search, I located an automotive body shop several hundred miles away that was experienced in doing restorations on cars like mine. I had the car shipped to that shop for repair.

After conferring with the body shop owner, we agreed that the restoration and repainting should require about four months to complete. And I gave the shop manager approval to proceed.

Sounds simple enough. The body panels and doors get pulled off, all of the interior upholstery is removed, the interior rust gets sanded down to the bare metal, the pieces are painted, and the car is reassembled — job done!. "Not so fast . . ."

As the body technicians got the car disassembled, they discovered the rust was far more extensive than it had appeared. The car had "lived most of its life in a state with cold winters, where salt was liberally used on roads to melt ice and improve traction. And salt corrodes steel at an extraordinary rate.

That was uncertainty's first "hit." It really shouldn't have been a "surprise," given where the car had come from. However, there were more unanticipated obstacles.

The first thing the body shop manager advised me was that some of the pieces (two doors and two fenders) were too far gone to repair. Replacements would be required for those. Unfortunately, replacement body parts for a 22-year-old car are hard to find. To make matters worse, the rusted rocker panels were also too far gone to save. These are part of the car's frame.

Finding replacement parts, then how much time and cost it should take, added another element of common and expected uncertainty.

The body shop owner suggested locating a "donor car" at a wrecking yard. If one could be found in good condition, it would be disassembled and its "good" parts would be "transplanted" to my car to replace the rusted ones. As it turned out, one of those "transplanted" parts was the entire frame of the car!

A donor car was located and acquired within about two months. Work proceeded on sanding and priming the composite car, inside and out. There was no way we would make the original four-month delivery date, but six or seven months wasn't out of the question.

However, then two of the most experienced body shop technicians left for other employment opportunities. It required a couple of months to get suitable quality replacements on board. Another not-too-surprising blow. Then the lead sanding technician broke his foot. Six weeks to get him back on the job. By now we're up to nine months. Then the body shop manager had a nervous breakdown, and the owner went through two more hired managers before finding one that was both qualified and willing to stay.

By this time, it's been more than a year since they started the project. Note that none in this series of incidents is too catastrophic. But their individual impacts accumulate throughout the whole sequence.

More bodywork technicians left the company, and it required several more months to get new ones on board. Then extended illness struck. More time lost. We are now eighteen months since the project that was supposed to have taken four months began. And finally, it appears that, as of this writing, the painting is underway. (See images, below) By next month reassembly of the car will start. But who knows for sure . . . ?

Painting is almost complete … after 18 months!

CASE #2: UNCERTAINTY AND PARKINSON'S LAW

Eli Schragenheim was invited by a sophisticated engineering company to investigate why a project that had been planned to finish in *one year* actually required *five years* to complete! Here is Eli's story:

On my first visit at the organization headquarters, I interviewed five high-level engineers who participated in the Pegasus (not its real name) project. All five (!) were deeply offended when I mentioned the purpose of our engagement: determining what caused the completion of Pegasus to be delayed for so long? Each one of the five reacted in a similar way, though their verbalizations were different. Here's a typical answer:

> *"Eli, finishing such a breakthrough project in five years is a failure? Really??? You must be out of your mind. This project is by far the biggest achievement our organization has ever had! In the U.S., they've already been working on it for the past ten years, and they are not even close to what we have achieved! This is a state-of-the art new technology that we developed. I can't tolerate an inquiry into why it took us as long as five years. You should be questioning how we managed to complete it in only five years!"*

This response pointed to what happened and why a project was estimated to take only one year, when everyone involved knew it would take much longer. Suppose you're a senior engineer in a company that has conceived a truly innovative new product requiring breakthrough state-of-art technology. You're definitely eager to participate in projects like this!

Naturally, management insists on a firm projection of the time, cost, and resources needed. They also want to know when the company can commit to deliver the new product to clients. And they insist on a detailed project plan, specifying the required resources and time.

As an experienced professional responsible for a specific task in most projects, you might reasonably estimate between 5 and 10 days of "hands on" time for the task. But this project plan asks for a fixed number of days. You're inclined to quote 10 days for your part of the plan. You might have said 5 days, but if the work actually took 9 days, your superiors wouldn't be happy.

Specifying one number to perform a discrete task actually ignores uncertainty, but even a routine task typically experiences unexpected "incidents." If you specify a conservative number, such as 10 days to complete, it will almost never finish earlier. Professional engineers know they'll be held to a shorter time in the future — say 8 days — just because they finished early last time. This is an example of Parkinson Law: *Work expands so as to fill the time available for its completion.*

But what if the task is not routine, but something innovative, never attempted before? Every developer wants to be part of the "latest, greatest thing." Your intuition, from prior experience, tells you that *at best* the task will require 10 days to complete, and some measure of luck would be required to meet *that* time. But when unanticipated complications occur, actual task completion might take 30 days or much more. But if you say the task could take up to 30 days, the other tasks might also be estimated conservatively. There's a realistic chance that management might scrap the whole project, because of excessive cost. So for truly innovative projects, engineers' behavior is typically much more daring, likely quoting the most optimistic task times, so as not to scare management.

Here's an important observation: When the Pegasus project was optimistically planned to finish at one year, the engineers didn't consider that "cheating!" One question that was impossible to answer: "How long will it take to finish such a new and revolutionary project?" And management's expectations are *one number*, such as "It will take one year." At this point, it's up to the technical professionals, well aware of the wide range of possible answers, to decide whether to quote the shortest time, or the relatively long time. And this is what they do, depending of how strongly they might desire to work on such a project!

Management's failure to recognize uncertainty, insisting on a fixed, specific time-to-completion or cost, results in employee behavior that protects them from the inherent uncertainty in any technical project. Eliyahu M. Goldratt famously observed, "Tell me how you measure me, and I'll tell you how I'll behave." Though employees don't fully explain their behaviors, they're a direct consequence of managerial policies. Ignoring uncertainty can have very severe consequences on the behavior of employees.

Decisions for Common and Expected Uncertainty

This category covers myriad decisions that are made daily, but which are affected by daily small "happenings" that cause minor disruptions. For example, the lack of something trivial such as coffee, or a short, sudden disruption in an Internet connection.

The importance of this category is that it forces all of us, individuals and organizations alike, to waste money and efforts trying to prevent more substantial damage. In an organizational process with several steps, and when the steps are interdependent, relatively small disruptions accumulate and threaten the planned outcome of the process.

People learn from experience and common sense that having assured access to a resource might prevent significant regret later. For instance, at one time or another we all have bought more food than we might immediately need at our supermarket, eventually having to dump it if it goes bad.

Most organizations view any surplus of time, cost, and capacity as a "waste." The problem, as most people are aware, is that a certain part of the "waste" is absolutely required to prevent a bigger problem if the unforeseen occurs.

For example, a critical decision might be how much manpower should be employed. Maintaining the precise number that is absolutely necessary ignores the situations in which a surge in work occurs. The adverse consequences could be excessively late deliveries, causing dissatisfied clients to take their business to a competitor. Misunderstanding what excess (or surplus) is really needed can ultimately lead to failure to hold market share, and the financial damage that results from it.

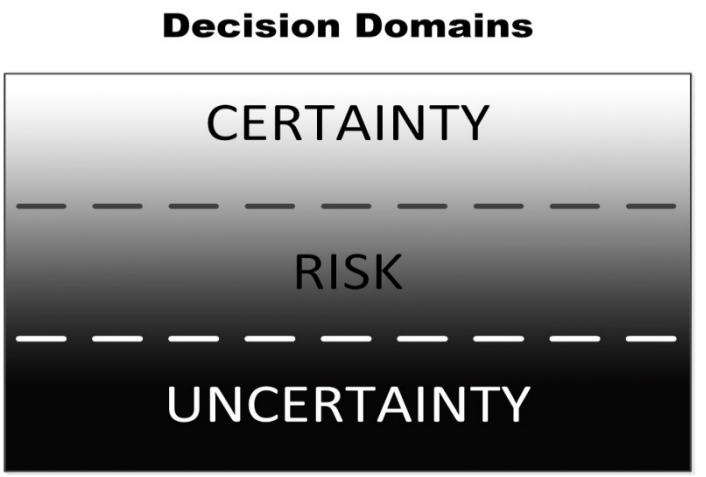

Figure 1. Decision Domains

Organizations, both retail and manufacturing, might carry too much stock, despite their desire for optimization. This is part of the very expensive price of ignorance in how to deal with common and expected uncertainty. Overall, not knowing how to manage common and expected uncertainty can cause a lot of damage — mostly financial — to most organizations. It can also impact the stability of process outcomes. Understanding this type of uncertainty is crucial to organizations maintaining firm control of their operations. Knowledge of the logic and tools for managing under uncertainty is necessary to maintaining stability and growth at the

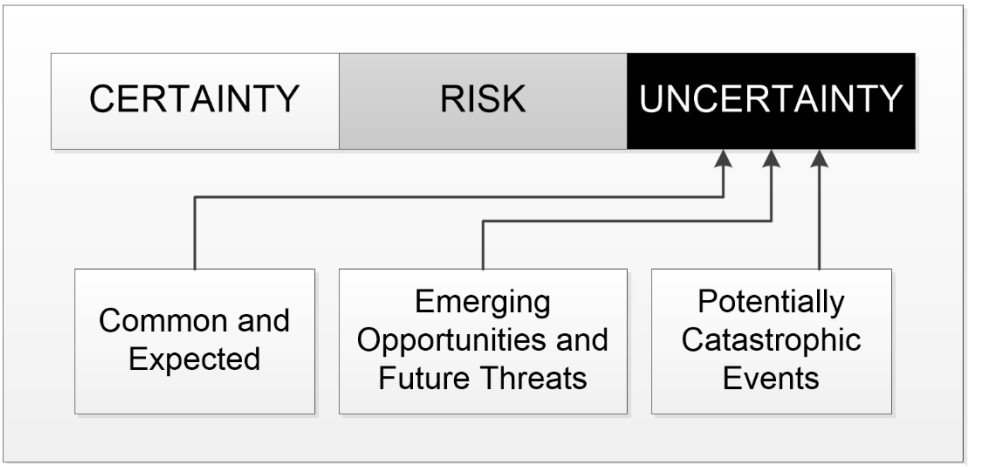

Figure 2. Types of Uncertainty

same time. Part 3 addresses tools and heuristics for successfully "beating" uncertainty in more detail.

Decisions Concerning Emerging Opportunities and Threats in the Future

By now, most professionals are familiar with the SWOT analysis model: strengths (S), weaknesses (W), opportunities (O), and threats (T). This category focuses primarily on decisions concerning the opportunities and threats that shape our medium- and long-term future. While common and expected uncertainty concerns routine decisions, this one considers decisions that are "beyond our comfort zone," where we normally have enough confidence in our intuition.

The point here is to be able to effectively express the *difference between the potential gain and the potential loss* from an opportunity that might be on the table, so as to make the best decision. Only when the decision maker is aware of the possible range of outcomes between the worst loss and the best gain can an effective decision be made.

This raises the issue of predicting/forecasting the future. For most managers it's common practice to arrive at *one* number or outcome of future reality. This is likely to be a distorted picture of the future! We know that some future outcomes might reasonably happen. Our intuition also might show us outcomes that theoretically could happen, but only in very rare cases, which we would likely not consider.

Decisions concerning potential events that could be catastrophic

This category is often characterized as "black swan" events. [3] The associated risks are so low that they seem negligible. These very low probability events are not intuitive, and assessing them statistically is extremely difficult. Just consider that a probability of one-in-a million is, mathematically, five times more probable than one-in-five million. In practicality, however, both are pretty much equally improbable.

Then there's the question of potential damage. The chance of a very low probability catastrophe happening may be impossible to predict. Its damage may also be difficult to quantify. The 2025 wildfires in southern California are a tragic example. The damage was beyond disastrous, but the probability of it happening, though low, could easily have been antici-

pated. Such fires happen every few years there, though none have been as devastating as these recent ones.

The first decision about potentially catastrophic events is: Do we consider these to be such rare cases, that we should simply ignore them? After all, most of us fly on airliners, consciously ignoring what might happen. For example, how often does a helicopter, flying at 300 feet altitude at night, collide with a commercial airliner, killing all aboard both aircraft?

Another consideration should be: what can we do to significantly prevent these rare, but immensely catastrophic events? Should we simply refrain from flying? Would the effect of doing so on our quality of life likely be too high?

When we consider wildfires, we, as citizens, expect the government to institute effective ways to prevent such devastating fires, even when the weather actually accelerates the spread of such fires. More, assuming that, in some cases the wildfire *will* spread, then there should be sufficient water in reservoirs to help the fire fighters contain the fire. (This did not happen in the 2025 Los Angeles wildfires.)

PART II. THE PSYCHOLOGY OF DECISION-MAKING

In Part I we examined the critical role of uncertainty in decision-making. Actually, "role" is probably not the best choice of words. When it comes to uncertainty, it's not a discrete thing with well-defined borders. It's a part of the environmental landscape.

Part I explained the nature of uncertainty and its characteristics. In it, we offered some possible tactics for coming to grips with uncertainty, and possibly reducing it in making decisions.

When all is said and done, however, making a decision is the sole purview of a human being. [4] And humans are subject to emotion. This is true of everyone, no matter how rational and dispassionate they believe themselves to be. The discipline that is most suited to understanding (and perhaps controlling) human emotion is psychology. With that in mind, let's take a "shallow dive" into the role of psychology in decision-making.

Motivation

What motivates people to make the decisions that they do? Many readers will be familiar with Abraham Maslow's hierarchy of needs. [5] (**Figure 3**) Maslow theorized that humans have needs that they strive to satisfy. Further, Maslow suggested that it's *unsatisfied* needs that motivate people to action.

Maslow also suggested that human needs are hierarchical. This means that satisfying needs lower in the hierarchy pyramid captures a person's attention until they are largely (though not necessarily completely) satisfied. At that point, the these lower level needs become less of a motivator than unsatisfied higher level needs. The person in question will then bend most of his or her efforts to fulfilling those needs.

Now, if at some point, the person's lower level needs fail to be satisfied to an acceptable level, both the person's attention and behavior will, temporarily, revert to that lower level need. Satisfying higher level needs takes second place until the lower level need satisfaction is secured again.

For example, consider typical white-collar office employees. Once employees are hired, their lower level needs for food, water, shelter, and for safety and security are met. At some point thereafter, the lower of the higher level needs — love and belonging — starts to become more important to them, and not coincidentally assumes greater influence in motivating their behavior.

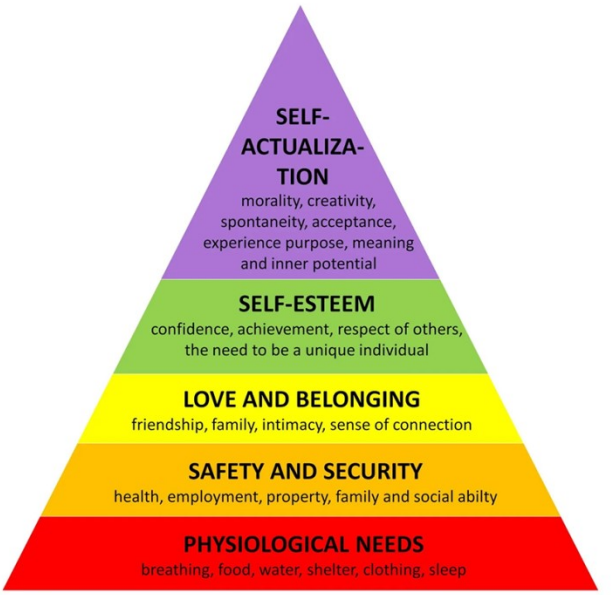

Figure 3. Maslow's Hierarchy of Needs

Now, assuming no reversion to the lower level needs, once the love-and-belonging needs are largely fulfilled, the employee will start thinking more about self-esteem and self-actualization. And these needs will begin to exert more motivational influence than the rest of the needs below them.

It's worth noting here that while everyone makes decisions in both their daily lives and their employment, the higher a person rises in an organization's hierarchy, the more these top two layers of Maslow's Hierarchy are likely to predominate in motivating a person's behavior.

Think about the executives you've known or observed. Clearly, their needs to achieve, to have the respect of their peers and subordinates, and their need to be unique (all self-esteem components) are important to them. At this point, their desire to be creative and spontaneous, their need to feel a purpose or have meaning to their lives, and to realize their inner potential (all self-actualization components) drive their motivation almost above everything else.

Near the end of his life, Maslow said that the order in the hierarchy is not nearly as rigid as he might have implied earlier in his career. Moreover, he also noted that in individual cases,

> …the order of needs might be flexible based on external circumstances or individual differences… for some individuals the need for self-esteem is more important than the need for love. For others, creative fulfillment may supersede even the most basic needs. [6]

Since Maslow's original conception of five-level hierarchy of needs, later pundits have suggested that there are three additional layers. However, discussing those is beyond the scope of this book.

Security Versus Satisfaction

For our purposes, let's simplify Maslow's hierarchy to two components: *Security and satisfaction*. We're setting aside the parts of Maslow's hierarchy that are negligible in consideration by executive decision makers. At the point that they've successfully climbed the corporate ladder and achieved a position of significant influence and authority, their behavior-

al motivators change. Professionally all they're really concerned about is a) not losing what they've achieved (security) and b) enhancing their record of achievement (and perhaps their image) by doing something new, different, and perhaps original (satisfaction).

Efrat's Cloud

In the 1990s, Efrat Goldratt-Ashlag [7] conceived a conflict diagram that explained the often agonizing choice people must at times make between their security and their satisfaction. This conflict diagram (or Evaporating Cloud) has proved to have such universal applicability that it's become known as "Efrat's Cloud." [8] (**Figure 4**)

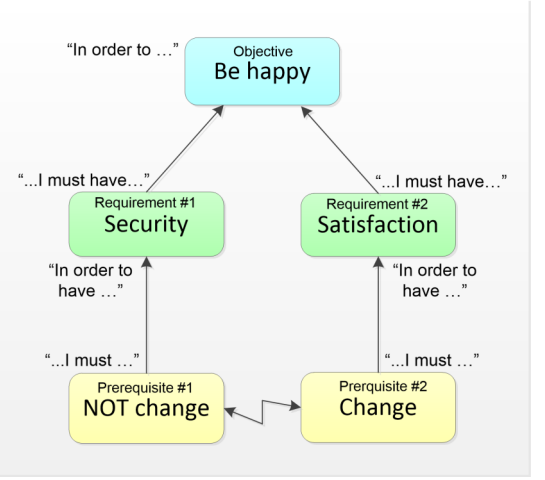

Figure 4. Efrat's Cloud

If one's objective is to be happy, then two requirements must be satisfied: security and satisfaction. To feel secure, one *must not* change. But to feel satisfied, one *must* change.

Some readers are probably thinking, "Not necessarily. I can change and still feel secure. Or I can <u>not</u> change and still be satisfied." On the face of it, this could be true. However, the distinction lies in diving a little deeper into how Efrat defined security and satisfaction. And it's a critical distinction.

Efrat defined security as the feeling of comfort and confidence one realizes from the consistency of doing the same things day after day: getting up at the same time, going to work and coming home at the same time, doing the same things in the same ways day after day. In other words, one knows *what to expect*, and these regularly met expectations instill a sense of comfort and confidence … and *security*.

For Efrat, satisfaction was the feeling of fulfillment and accomplishment one realized from achieving a difficult objective, when there was substantial doubt about the chances of success. For example, a project manager might bring a project to successful completion ahead of schedule and under budget. Or an athletic team might defeat a favored opponent when nobody thought that could happen.

The difference between security and satisfaction is that to remain secure, one steadfastly avoids doing anything different or out of the ordinary. In other words, don't take any risks. But the satisfaction of achieving a difficult objective typically isn't realized by doing the same old things in the same old way. It usually requires creative thinking or doing something most people wouldn't expect (or maybe even think of). In other words, taking a risk, when failure might be a distinct possibility. This choice presents a dilemma for a decision-maker.

Another way of expressing this dilemma is: Are you a risk-*taker*, or are you *risk-averse*? Risk-aversion leans toward the security side of the dilemma. Risk-taking favors the satisfaction side. (See **Figure 5**)

Think about hang gliders or wing-suit flyers. Does their need for security outweigh their need for satisfaction? No! Both of these activities are very risky. So risky, in fact, that if

you killed yourself in pursuing them, your life insurance company would likely not pay out to your survivors. Or think about an entrepreneur, such as Elon Musk. Is he risk-averse? Most assuredly not. He might not be as extreme a risk-taker as a wing-suit flyer. But he surely isn't going to hide in his shell like a tortoise, either.

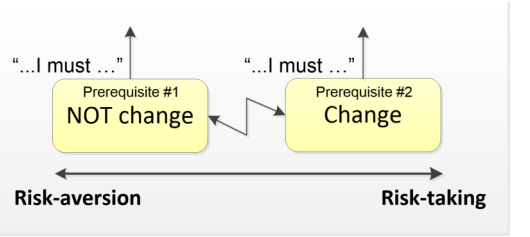

Figure 5. Risk-Aversion vs. Risk-Taking

The Technology Adoption Life Cycle

Risk-taking and risk aversion can be framed another way. In 1991, Geoffrey Moore proposed a technology adoption life cycle. As a consultant for Regis McKenna Associates in silicon valley, he observed different reactions in people's acceptance or resistance to the introduction of new technologies. Specifically, his audience was senior level managers and executives.

Moore conceived the concept of the technology life cycle. [9] He expressed it as a normal distribution curve, as shown in **Figure 6**.

At the left tail of the curve are the *innovators*, who are determined to try the latest, new-

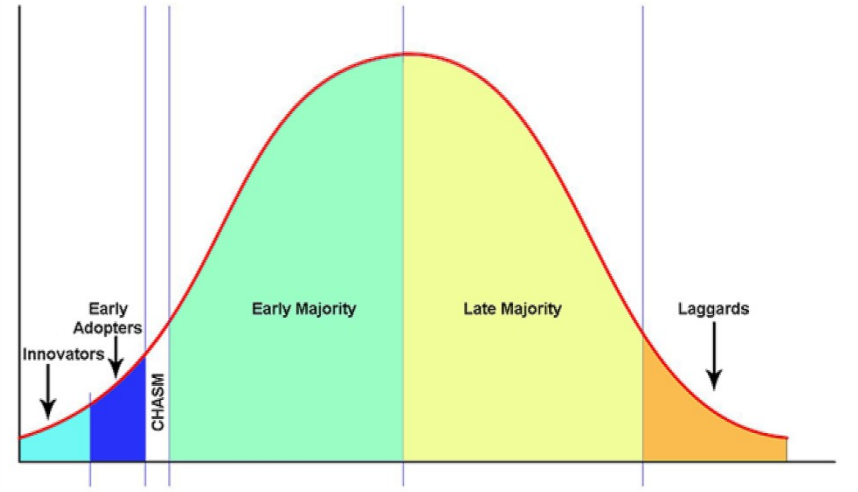

Figure 6. Technology Adoption Life Cycle (Moore, 1991)

est things to enter the market. They don't really care about practicality; they just want to be on the "cutting edge."

Moving to the right, we see the *early adopters*. These are forward-looking people whose underlying concern is that the new technology have some practical application to their domain. They're willing to take a risk on it, knowing that it might fail, but also that it might deliver disproportionate positive results, too.

Then there's the "chasm." Moore suggested that there was a gap over which new technology had a very difficult time leaping. On the other side of that gap was what he called the majority. This is part of the mainstream population where the largest percentage of prospective adopters of new technology lies. As you can see from **Figure 6**, the majority represents about one standard deviation left and right of the mean. Moore further divided the

majority into *early* and *late*.

The *early majority* is sensitive to the introduction of new technology, but it's much more conservative than the early adopters. Decision-makers in this group are cautiously optimistic about new technologies, but they want to be convinced. Convincing them requires showing them where someone (or some organization) else has already succeeded with the new technology. More than that, they want to see that the technology has succeeded in *their* type of industry or activity.

The *late majority* is even more conservative than the early majority. Left to themselves, they would likely not adopt any new technology. But they do so only because they see everybody else doing so, and they don't want to be left behind. Call them "reluctant risk-takers."

Finally, the last group, at the right tail of the curve, is the *laggards*. Like the innovators and early adopters, they're small in number, but they're well "dug in." They never make the change to a new technology unless they're on the verge of organizational/ business "death."

The Technology Life Cycle and Risk-Taking

While Moore's technology adoption life cycle was intended to address the question of marketing high technology to the business world, it can also be more broadly applied to individuals in their daily lives. **Figure 7** shows Moore's graph with the labels modified. Call it the technology adoption life cycle applied to risk-taking. [10]

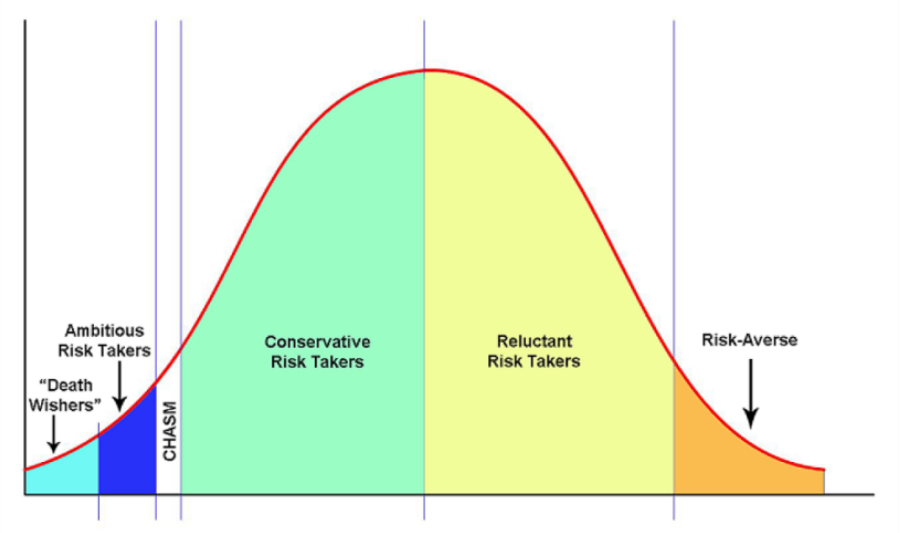

Figure 7. Technology Adoption Life Cycle Adapted for Risk-Taking

Instead of innovators, we have "death wishers," who eagerly engage in highly risky behaviors. Next are the ambitious risk-takers, who are willing to stumble a bit and might even accept odds of 50% for success. After the chasm are the conservative risk-takers and the reluctant risk-takers. Conservative risk-takers can be convinced to try something new, though left to their own devices, they'd just sit back and watch. Reluctant risk-takers don't voluntarily accept change, and they certainly don't seek it out. And the totally risk averse are typically "left behind."

Uncertainty and Risk-Taking

By now, it should be obvious: there's a direct relationship between uncertainty and the willingness of decision-makers to take risks. Now, why would this relationship exist in the first place? The answer is that it's emotional.

And the underlying emotion that connects the two is *fear*. This fear springs from two "roots." The first of these roots is endemic to the individual. The second is related to the larger system. Let's look at the causality behind an executive's reluctance to make a specific decision. (**Figures 8** through **10**) Read these diagrams from bottom to top.

Figure 8 explains the personal motivation influencing many executives' reluctance to make decisions that might have a significant impact on the organization. Now, recall our earlier discussion of Efrat's Cloud (**Figures 5** and **6**). The degree of risk aversion affects executives' predisposition to take chances. And Moore's Technology Adoption Life Cycle (**Figure 6**) suggests that the majority of decision-makers (approximately one standard deviation either side of the mean) are naturally risk-averse to some degree.

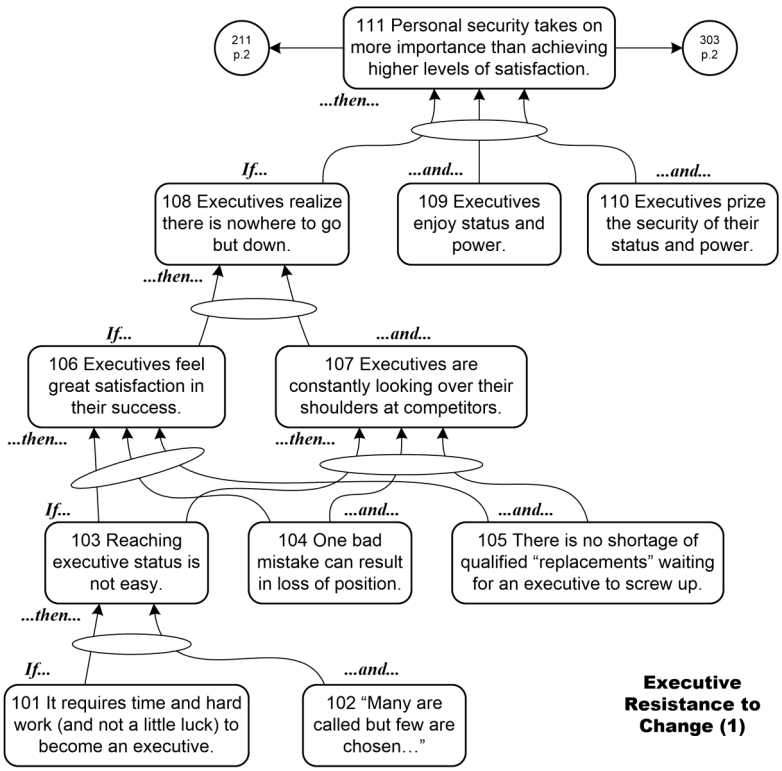

Figure 8. Executive Resistance to Change (1)

But **Figure 9** introduces another factor in an executive's risk aversion: the organization, or systemic, viewpoint. Look specifically at the blocks numbered 205, 209, 210 and 211. This part of the logic tree emphasizes the concern most executives will have for the well-being of the company's employees and its stockholders. This is a very real concern, and it tempers an executive's inclination to take risks.

So in reality, there are two independent causes of an executive's reluctance to make major decisions. One is fear of the personal consequences. And even if a particular executive

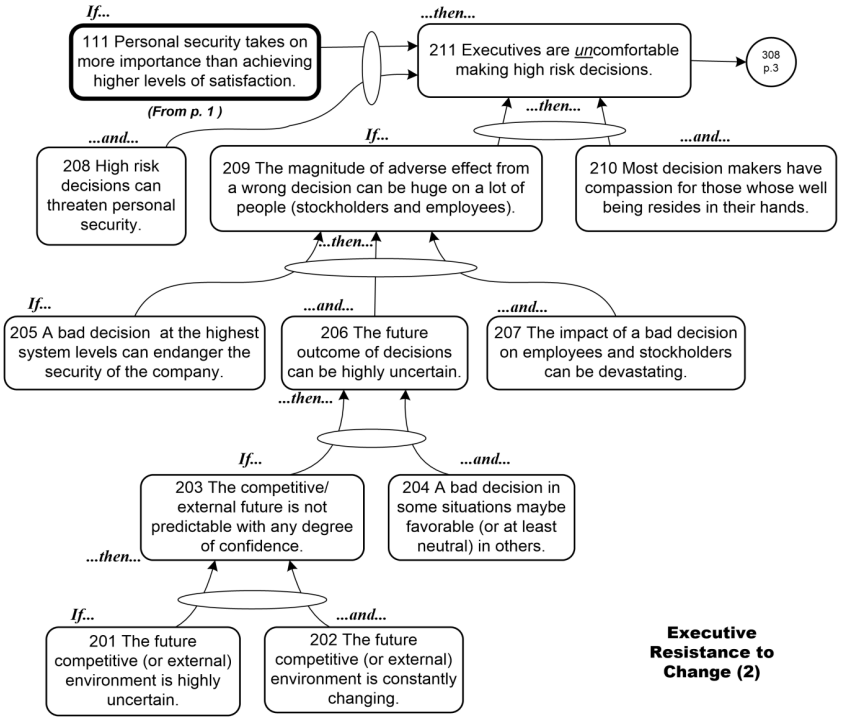

Figure 9. Executive Resistance to Change (2)

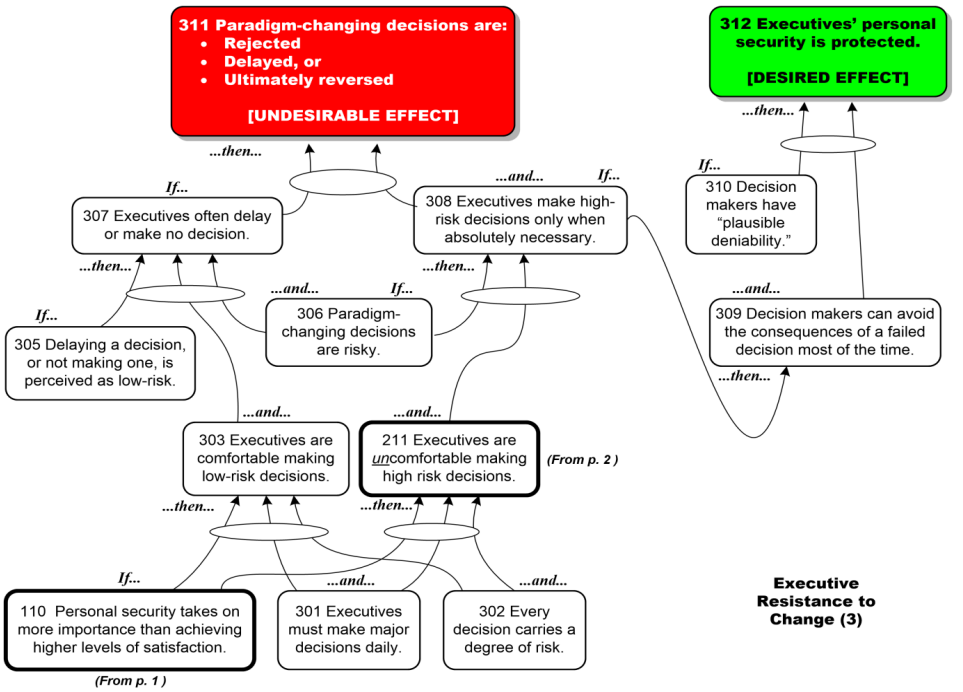

Figure 10. Executive Resistance to Change (3)

isn't terribly concerned about his own security, the responsibility — the obligation — for the well-being of employees and stockholders alone is enough to instill some level of risk aversion in a decision-maker. Now, take a look at **Figure 10**, below, to see the ultimate outcome of an executive's risk-aversion.

As an executive considers the gravity of the decision he or she is expected to make, low-risk decisions are eminently easier to make than high-risk decisions. That's why they're comfortable making low-risk decisions. Not so much with high-risk decisions. As a result, high-risk-decisions — perhaps paradigm-changing — are either rejected or delayed (possibly until the change in question dies a natural death). In the event that a high-risk decision *is* undertaken, it's at risk of being reversed if the executive "gets cold feet." The ultimate outcome is that the executive's personal security is preserved.

PART 3. TOOLS AND HEURISTICS

Practical Steps for Dealing with Uncertainty

What tools can we use to deal with uncertainty, especially in organizations, when the hard data necessary for managing it aren't available?

In mathematics and statistical analysis, when the basic data *are* available, the focus is on calculating, or at least estimating, the *expected value* and the *standard deviation*. In forecasting, the accepted expressions are the single-number forecast and the forecasting error. This opens the way to define *a reasonable range* around the expected value, usually called the *confidence interval*. The likelihood that the actual uncertain value falls somewhere within the range of the confidence interval can be fairly high, perhaps as high as 95%.

Many (most?) people don't fully understand the term "standard deviation." It's not intuitive. Moreover, in most practical common cases, there's no way to calculate or even estimate these values, because not enough reliable historical data are available that truly apply to the situation.

We suggest several steps to deal with uncertainty in a pragmatic, effective way, even without reliable past data:

1. Classify various types of decisions under uncertainty according to the potential damage that might occur, especially when a certain decision could prevent or reduce the damage. Certain decisions could result in a small loss, or a small delay in the plan. Other decisions might cause much more serious damage, especially when the decision-maker considers actions that are out of the decision-maker's comfort zone.

Another type of decision deals with facing a risk for a disastrous event, such as an earthquake, but usually with very low probability, so the dilemma is whether to invest in costly steps that would face the disaster and reduce its damage. Another way to categorize decisions under uncertainty is by inquiring whether certain information, which can be found, could assist in the estimation of the total uncertainty. For instance, when farmers need to decide what to plant in their fields, analyzing the weather forecast information might be helpful, even though there are additional uncertain impacts, such as the world prices for certain agriculture products when the crops will be ready. It is possible that in the near future artificial intelligence (AI) could be used to estimate the uncertainty that is relevant for certain aspects of the decision at hand.

2. Outline the reasonable boundaries of uncertainty that are useful for our decisions and actions. The idea is to estimate a fairly conservative potential outcome of the decision, and a more optimistic outcome. The difference between the two encompasses a wide range of what reasonably might happen. For instance, when we plan to go from New-York to London, our reasonably optimistic assessment might be an on-time arrival, or maybe landing a half hour early. Our reasonably conservative expectation might be landing three hours later than scheduled. In this example the possibility of landing two days late would not be considered "reasonable," because its likelihood is too rare to include it in our plans.

3. Estimate an acceptable picture of what we can reasonably expect the worst damage and the best gain to be for the decision under consideration**.** Use the boundaries of uncertainty, and the key alternatives for action (which would be based mainly on our own intuition). When we consider the New-York London flight, being half-an-hour early might not be all that useful, but the original plan can be executed as is. However, let's say our arrival is three hours late. If the initial departure decision included an important meeting taking place perhaps two hours after the planned landing, then the meeting would have to be postponed.

4. Make the decision based on any partial information available. In the above example, when we decide to book a limousine driver to meet us, we would start with the official airline schedule, assuming the limousine driver monitors arrival information at the airport. However, for arriving at an important meeting, we should be more conservative: We might allow for being late by three hours, plus the time in might take (again using a conservative assessment) to reach the meeting location.

Tools and Heuristics: How to Deal Effectively with Each Category

Decisions, including short-term planning and execution, are often impacted by common and expected uncertainty (discussed earlier).

Let's focus here on organizations, where the adverse consequences of ineffectively addressing common and expected uncertainty can often be much higher than most decision makers realize. In planning a series of actions required to achieve an important objective, an organization must anticipate many undesirable — but not necessarily surprising — situations that might eventually compromise the quality and timely achievement of the objective. In any sequence of time-sensitive actions, delays along the sequence accumulate.

The key tool for preventing the adverse consequences from accumulating fluctuations is to *protect the important outcome with a buffer — but NOT every single action*. When the objective is to deliver on time, then that buffer has to be time. In practice, this means that when a series of actions takes two weeks, *on average*, then a promised delivery in two weeks cannot be reliable. When reliability is important, the promise has to include a time buffer.

When a single buffer protects the outcome of a series of actions, there is no need for each action to have its own buffer. However, when management suffers from the illusion of certainty, the people responsible for action often feel that they have to apply a time buffer on their own. Typically, they do this by claiming they require more time than the average, just for the activity itself. (And often they consume that extra "buffer" time, even when there is no need to). These are hidden buffers. Their purpose is really *to protect the workers from management, not to protect an on-time delivery*. Or, more precisely, to protect workers from management's failure to recognize that the time required for a certain action can't reliably be a fixed number, due to natural uncertainty. An insidious characteristic of hidden buffers is that most of the time the buffer will be fully consumed, whether needed or not. As

a result, that lost time isn't available to protect the completion of another action, which may have had a real need for extra time.

Single time-buffers could be used to ensure the on-time delivery of many different individual missions, where a mission might contain different tasks performed by different human resources. This time control can be achieved by the using a priority system to identify the few missions that are "almost late." To do so, we define a "red-line" time. This comprises the last 20 to 33 percent of the total time given to the mission.

When the mission is still not completed, but the total time passes the red-line, a warning -flag is set, coloring the mission as 'red,' which means very high priority— immediate action is necessary to save completion from being late. Management should be aware of the missions in red and put pressure to push the mission to completion within the red-line time, thus keeping the delivery of the mission objective on-time. Having a list of "red missions," given proper response to the red-orders, guarantee on time performance, in spite of all the common and expected uncertainty "messing with" the flow of work.

Refer to case #2 for an example of the damage to management when buffer management is not used.

Tools for Evaluating New Opportunities

What we've discussed up to this point is part of the *planning*. In dealing with common-and-expected-uncertainty the required decisions and desired outcomes are clear enough. However, during *execution*, we need to protect ourselves from uncertain incidents that could seriously degrade our system's performance. At the very least, in the event of unexpected obstacles, we want to end up no worse off than when we started. Noting the "red missions" and giving them a high priority over other missions supports the right focus that allows to all the missions to complete at the planned times.

When it comes to evaluating a new opportunity, the decision itself is on the table: To go for it, or not? Such a decision might include considering several alternatives.

When evaluating a new opportunity, the key uncertain variable is how much value would be gained (or lost) as a result of going ahead with the opportunity? This is basically predicting/forecasting the total impact of the outcome on the decision maker and/or the organization. (Why is the decision maker part of the consideration? Part II answers that particular question.)

Instead of making a clear and *precise* prediction of the outcome — an impossibility — a much more effective approach is to estimate *the reasonable range* of the possible outcomes, based on the best intuition plus relevant information we can gather.

Let's say you have to make an airline flight at a specific time. You must decide when to leave your home or hotel. Your final decision considers the possibility of traffic jams, the time it might take to check in, going through immigration control, security checks, and only then proceeding to the gate.

Most of us allow more time than the absolutely minimum necessary time. The difference between the shortest time required to reach the gate and a time that seems "way more than enough" is the reasonable range. Assuming it's critical that you not miss the flight, you'd usually prefer to depart for the airport at a time that represents "way more than enough." Of course, if it happens to you that you missed a flight in spite of taking much more than enough time, then this feedback would lead you to reassess *how much* "way

more than enough" should be for similar situations in the future. In other words, *we need to update our estimation of what might happen frequently. By doing so, we improve our capability to protect ourselves in the future*.

The previous example demonstrates the common longest-yet-reasonable value of the reasonable range. Now here's a different decision that would take the more conservative approach.

Let's say that one of your subordinates in your organization asks for an urgent meeting on a day that you know is going to be very busy. You think this employee is a bit of a "snooze," and that the topic is trivial. You can't really refuse to meet this person, but you'd like to confine the meeting to only 15 minutes, which you consider the shortest time that still satisfies common courtesy. But certainly not 30 minutes, much less a full hour, which you might allocate for meetings of greater importance.

Assessing the reasonable range as a formal tool is critical for taking uncertainty into account. So, even if the not-too-reliable employee surprises you with a worthy idea, 15 minutes might still be enough to grasp the potential and pursue it in more detail later.

The reasonable range raises the question of "what is reasonable." One of the complications of uncertainty is that certain outcomes might have only a very low probability. And it may be impossible to precisely assess this risk.

What kind of uncertain event might we *not* consider? This is where the term "reasonably" gives us a clue, though not a precise one, about something we should consider.

The importance of clearly stating a reasonable range for any important decision is that it outlines for us what *we know we don't know*. In retrospect, any actual outcome that falls within the reasonable range we have already considered signifies that we predicted it correctly. We don't need to consider all cases that fall within the range. Most of the time it is virtually impossible anyway. It's more than adequate just to consider the two extremes, and pose the question: *If we prepare for the optimistic side of the range, but reality materializes closer to the maximum point, what would the impact be?*

On the other hand, if we favor the optimistic side, but reality unfolds closer to the conservative side, what might the adverse consequences be? The following example demonstrates the need to determine the range of consequences of the actions we might take, based on our estimation of the reasonable range. We hope to provide the decision maker the best information available to make the final decision.

An Investment Example

Let's assume you're given an opportunity to invest in a specific, potentially lucrative initiative, led by someone you truly respect. The investment can only be in multiples of $100,000. You will be able to recoup your relative share of the results no sooner than two years after you make the investment.

This is a risky investment. Your best estimation is that in the most reasonably optimistic case, the venture could return *ten times* your original investment in just two years. The reasonably pessimistic case is that the whole investment would be lost.

Mathematically, assuming the chance of the investment realizing a profit is much more than 50 percent, then this is a desirable bet. But if you're considering investing $500,000,

and eventually you lose it all, the real financial (and other) damage to you might be intolerable. This implies a basic difference between the absolute value of winning and the absolute value of losing. If you cannot tolerate losing $500,000, then don't invest $500,000, even with some probability of gaining $5,000,000!

But if losing just $100,000 is tolerable, then maybe you should invest, based on your assessment that the reasonable range is between losing it all and gaining 1,000 percent (ten times the investment). Your decision is based on both the reasonable range, and the limit you set for how much loss you can tolerate.

We have no special insight regarding preventing a very low-probability catastrophe. The question is more about how tolerable measures to prevent it might be. After all, we *do* need enough to be able to live life in a desirable way. So, most of us drive our cars, in spite of the chance for deadly accidents. Most of us take airline flights, in spite of the risk of crashes. But most of us also insure our homes against damage from earthquakes. The insurance rate is generally tolerable, and an earthquake would destroy our home. So, we're content to mitigate the risk of unacceptable loss this way.

Conclusions

We must be continually aware of *uncertainty!* Beware of the *illusion of certainty!* This is the key message, and it applies to both individuals and organizations.

The tendency to ignore uncertainty is more deeply embedded in organizations than in individuals. This has more to do with the seemingly unfair (and frightening) after-the-fact judgment of decisions made by managers (or executives), ostensibly for the good of the organization. This kind of judgment invariably ignores the state of uncertainty at the time of the decision.

The first thing to do is always try to evaluate the level of uncertainty, even when it is based just on intuition. Actually, facing uncertainty without the use of well-defined probabilities is by far the more common occurrence. Actual, precise probabilities are found mostly in gambling and theoretical models in mathematics. In the vast majority of actual real-world situations, one has to assume the range of *reasonable uncertainty*, partially based on cases from the not-too-long past, but mainly on intuition, which is usually also based on the experience of the decision maker.

Translating intuition into sensible reasonable range allows us to examine the outcomes of the extreme limits of the range. Eventually this produces the best input for making a decision knowing both the largest potential damage and the largest gain.

Overcoming the fear of unjust after-the-fact judgment of decisions (which were made under significant uncertainty) requires *documenting the predicted reasonable range, and explaining the key information that led to the decision*. When important organizational decisions are required, getting an agreement on the reasonable rage becomes crucial to good decisions. Moreover, the decision maker also has reasonable protection against after-the-fact judgement.

There is one more significant advantage to maintaining ongoing estimations of the reasonable range. It allows collecting usable feedback on your estimations. When many decisions have been taken and documented, it is possible to verify the "track record" of your predictions. Eventually reality tells us what actually happened and where the actual result falls: In what percentage of our estimations the actual result was within the reasonable

range, and how many were out of the range.

This feedback should be used to understand our own biases. Are we too optimistic, meaning that many actual results are below our "reasonable" pessimistic prediction? People who are too pessimistic would often see results that lie beyond their optimistic side of the range. That feedback should be used in the future to overcome personal bias and arrive at better intuitive assessments of the reasonable range .

We can predict that over time our understanding of the true impact of uncertainty will significantly improve. Imagine what this could do for the quality of *your* decisions. ♦

ENDNOTES

1. Willink, Jocko, and Leif Babin. *Extreme Ownership: How U.S. Navy SEALS Lead and Win*. NY: St. Martin's Press (2017)

2. This is sometimes referred to as "Monday morning quarterbacking," meaning criticism after-the-fact for what was not obvious ahead of time.

3. Taleb, Nassim Nicholas. *The Black Swan: The Impact of the Highly Improbable*. NY: Random House (2007)

4. Or at least it should be; these days many humans abdicate their decision authority to computers.

5. https://www.simplypsychology.org/maslow.html

6. Ibid.

7. Efrat is the daughter of Eliyahu M. Goldratt, creator of the Theory of Constraints.

8. Moore, Geoffrey L. *Crossing the Chasm*. NY: Harper Collins, 1991.

9. Dettmer, H.W. *Systems Thinking and Other Dangerous Habits*. College Station, TX: Virtual Bookworm Publishing, 2021, p. 339.

10. The full method of maintaining effective priorities between different orders and missions, is called Buffer Management. Many Theory of Constraints books describe it in detail, e.g, *Supply Chain Management at Warp Speed* (Schragenheim and Dettmer, 2007).

ELI SCHRAGENHEIM

Eli Schragenheim brings a background in mathematics to the field of decision making. He studied film-making and was active as a TV director for nine years before striking out into computers and software programming. Eli worked closely with the late Dr. Eliyahu Goldratt, creator of the Theory of Constraints (TOC). Eli developed computer simulators that are widely used as educational tools for managers. These simulators demonstrate key logical insights for managing semi-complex systems. His Management Interactive Case Study Simulator (MICSS) served as a laboratory for learning real cause-and-effect on manufacturing shop floors. Eli contends that ignoring uncertainty is, by far, the biggest core problem of every organization.

H. WILLIAM DETTMER

H. William "Bill" Dettmer is the author of seven books on constraint and complex system management, including *The Logical Thinking Process* and *Systems Thinking and Other Dangerous Habits*. Over his 25-year consulting career, his clients include United Health Group, Raytheon Missile Systems, the U.S. Defense Microelectronics Agency, Bayer GmbH (Germany), Bosch GmbH (Germany), Deloitte Touche Tomatsu and Aerosud SA (South Africa), and the Centre Recherche Henri Tudor (Luxembourg). A retired U.S. Air Force command pilot, Bill has 4,000 hours of multi-engine jet flying time and 75 combat missions to his credit. He holds a BA degree from Rutgers University And a Master of Science in Systems Management from the University of Southern California, where he also taught graduate courses in the same program for seven years.