CHAPTER 5

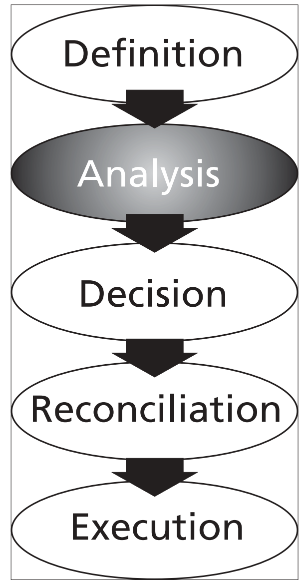

ANALYSIS CONCEPTS: UNCERTAINTY AND RISK

The habit of gambling contrary to reasonable calculations is a military vice which, as the pages of history reveal, has ruined more armies than any other cause.
—B.H. Liddell Hart: *Thoughts on War,* 1944

IDEALLY, WE WOULD LIKE to approach any decision fully able to predict its outcome with certainty. As we have said, during analysis we identify alternative courses of action to achieve our decision objective and compare them based on our criteria. Put another way, we ask ourselves what are the consequences of each alternative in terms of cost and effectiveness? And, what happens when we cannot predict these consequences with certainty? Depending upon the magnitude of our uncertainty and the importance of the uncertain issues, our ability to make a sound decision may be reduced substantially or eliminated entirely.

Taken to the extreme, when we have no sure knowledge, i.e., a situation of complete uncertainty, we may as well decide by chance. In most cases, this is undesirable—although there are certain decisions for which a coin flip is as good a method as any other. Uncertainty, in our framework, is the amount of doubt that we have about our projections of cost and effectiveness.

Risk rises from our uncertainty. Risk is the possibility of failure and of suffering loss or harm because of our unsure knowledge. The loss or harm is tangible and we can predict its effect although we may be less confident whether or not the effect itself will occur. In this chapter, we address the sources of uncertainty and risk and we present some methods for coping with them.

Objective Probabilities

Uncertainty and risk are very much about what is knowable and what is not. When we can identify a set of outcomes, when we are confident we know their behaviors in terms of cost and effectiveness, and we can predict the likelihood of any particular outcome statistically, we have objective probabilities. We can build objective probabilities whenever we can tabulate data from what has happened in a large number of prior cases that are similar to our present decision. Commonplace examples include weather forecasts and baseball batting averages. Each is calculated based on probabilistic information about what has happened over time with a set of initial conditions. Therefore, we can be confident of the likelihood we assign to different outcomes when the next occasion arises, so long as we believe that the future and the past will behave in a similar manner.

Weapons performance characteristics are a good military example of objective probabilities. When weaponeers calculate a missile's circular error probability vis-à-vis the aim point, they are establishing a typical objective probability-based measure. We cannot predict whether any particular missile will hit its target, but we can speak knowledgeably about the likelihood of a hit, as well as misses of various distances. In logistics, we measure the performance of an inventory system – and its possibility of failure—with objective probabilities. No inventory system is good enough to provide completely accurate information about every item in its database. But we do know with great confidence the likelihood of an inventory error. Precisely the same is true for quality control measures. While we cannot know whether an individual munition will explode, based on tests and experiments we can know how likely it is to be a dud and adjust our targeting plans and inventories accordingly.

Subjective Probabilities

In conditions of true uncertainty, one or both of the following conditions exist: we cannot predict the consequences of our alternatives with confidence or we cannot know their probabilities of occurrence. This happens when we are confronted with decisions that involve unique elements or at least elements too dissimilar to support a statistical probability based on the past. For example, suppose we must estimate the likelihood of a baseball batter getting a hit who has never batted before. We have no basis for assigning an objective probability to that estimate.

Our interactions with other people frequently fall into this category. As we grow to know someone, we can usually begin to see broad patterns in his or her behavior. However, people, situations, outside events, etc., change so continuously that we are seldom in a position to say that a certain individual has a 30 percent probability of doing any particular thing and mean it literally. At least, that expression of probability would have a very different level of meaning than a 30 percent estimate attached to a weather forecast or a gambling bet. Because these predictions are largely intuitive, they are subjective probabilities.

Consider the range of defense problems that fall into this second category in which uncertainty and risk are important elements. All our estimates of international behavior are based on our limited knowledge of the plans and intentions of others. We cannot say that Iraq has a .5 likelihood of launching an attack in the same sense that we can say that a typical 155mm artillery shell has a .9 probability of exploding. The U.S. plans military forces based on the two overlapping major theater war requirement. It would be enormously helpful if we could estimate the chance that a second major theater war might actually occur as we fight the first. We know that a second war is possible, but beyond that we cannot say. If we could say more, we could know much better whether or not planning for two overlapping major theater wars is a good use of our defense resources.

In the same sense, we cannot predict the likelihood that a Kosovo-type conflict or another crisis will occur in the upcoming year. If we could, we would have a much better understanding of what level of preparedness we require for such conflicts and force planning would be a science instead of an art. Because we do not know the probability of another or several similar events that will require deployment of U.S. forces, we must be ready for these scenarios whether their chances of happening are high or low.

Strategic planning and decision making are rife with uncertainty and the consequences of dealing with uncertainty are significant. Even at the operational level, can we know by objective probability how an adversary will respond to one battle plan versus another? For example, which course of action is more likely to produce an adversary's surrender: air attack or ground attack? Does the surrender of one adversary to one type of attack represent a universal truth or an exception? Does a new solution to an old problem represent new truth about the possibilities of the future or a fleeting aberration?

SILENT AIR ASSAULTS: GLIDERS, THE "OTHER" AIRBORNE

In 1940, during World War II, the Germans were the first ever to employ gliders in airborne assaults, allowing troops to land ready to fight and with unit integrity, at least at the squad level. The glider-delivered troops did not require parachute jump training and gliders could carry heavier equipment, like jeeps and anti-tank guns, than aircrews could shove out of aircraft in flight or land with a parachute. Also, because gliders could be released far from their destinations, the tow planes' motors did not alert defenders, whereas the parachutists' air transports were audible to them. Both parachutists and gliders were very vulnerable to ground fire and required open fields for landings. Because of their tactical advantages, many force planners thought glider troops would supplant paratroopers in airborne assaults. In their initial use in combat, German gliders landed engineers atop Fort Eban-Emael in Belgium and achieved complete tactical surprise; the strongest fort along Germany's western border fell quickly, unlocking supply lines into the Low Countries and France for the Panzers that had advanced through the Ardennes Forest.

After this success, how would you have reacted as a force planner in 1940, forced to choose between forming paratrooper or glider regiments? Which unit is more cost-effective? Allied and Axis planners hedged; their armies built both airborne and glider forces. It turned out the glider had a short life in combat, from 1940 to 1945. Casualties among gliders, glider pilots, troops, and cargos were high and for the most part they consumed material and manpower resources that could have been better allocated elsewhere, e.g., a pilot with the skill to land an unpowered aircraft on an unfamiliar, unimproved field at night might be better employed flying as his or her primary duty. A glider pilot without that skill was a hazard. Almost all major contemporary militaries still include paratroopers.

Should this result have been foreseeable? Could the force planners of 1940 have removed the uncertainty in their decision by doing more research or experimentation? Or was the glider an appropriate weapon for its albeit brief combat life? These same issues confront force planners today as the services strive toward Joint Vision 2020: should we press forward with new organizations and structures or use this "strategic pause" to do more experimentation and reduce uncertainty before committing to new, expensive paths?

If we could resolve issues about the efficacy of weapons or the effectiveness of air power with certainty, we could easily decide the current debate about the relative division of labor—and therefore resources—between air and ground forces. Logically, we suspect that striking certain kinds of targets with air power will incline an adversary to become conciliatory. But in any particular case, we cannot tell a Commander-In-Chief how likely it is that an adversary will indeed react this way.

Yet there is a helpful way we can characterize uncertainty. For instance, we may say about the baseball batter for whom we have no data, "He will not get a hit." By that, we mean that we rate his chances of hitting at less than 50 percent. This estimate of .5 is not an objective probability because no information exists about his previous batting performance. But there may be observations we can use to build a more refined estimate of his chances of hitting during his first at bat. We may evaluate the way he swings the bat or stands at the plate. We may see whether he appears confident or hesitant. While these clues are not the basis for an objective probability, they may support an expression of how likely *we believe* he is to get a hit. We generate a probability without data, without knowing the past (in a scientific sense), instead we evaluate our state of mind and determine how confident we are that the batter will or will not get a hit. We create a subjective probability.

We use subjective probabilities all the time in defense decision making to express our evaluations of uncertainty. For example, early in the Cuban Missile Crisis, President Kennedy estimated the chances of war between the U.S. and Soviet Russia as one in three. Subjective probabilities are less likely to be expressed numerically than objective probabilities; in fact, expressing them in numeric parlance can lead to misinterpretation. We believe the chances of two major theater wars overlapping one another are more than trivial; but when we say the chances of war are one in three, we do so knowing the chances of war are not measurable statistically like a batting average. Rather, this probability-phrased expression is shorthand for articulating a subjective estimate of likelihood based on experience.

In other words, we may not be able to calculate the statistical or objective probability of an event, but we still may have an idea for some reason of its likelihood. That idea may come from experiences which, although not identical, we believe are relevant to the probability we are trying to estimate. The estimate may also come from expert knowledge or intelligence we have about the specific circumstances of the event. Plainly put, the subjective probability does not measure the frequency with which something occurs, it captures how confident we are that something will or will not happen. This is important because as we analyze options and sort information provided by others, we need to know which data is subjective and which is objectively derived; the former is far more open to interpretation and dispute. In some cases, we may justify our assumptions based on subjective probabilities. When we do, we must be sure to inform our decision maker and be ready to be challenged by other stakeholders during reconciliation.

UNCERTAINTY AND FORCE PLANNING IN EUROPE

In 1994, the U.S. European Command knew it would continue to draw down from its 350,000 European Troop Strength (ETS). It also knew that the bulk of these troops would be Army and Air Force units stationed in Western Europe. ETS does not include rotational naval forces in the Sixth Fleet, typically comprising an aircraft carrier task group, submarines, patrol planes, and the Mediterranean Amphibious Readiness Group with its embarked Marine Expeditionary Unit-Special Operations Capable.

The force reduction process began opportunistically when the U.S. Army's VII Corps deployed for the Gulf War then re-deployed to the United States instead of returning to Europe. But, after their departure, how much deeper, if at all, should ETS have been cut? A General Accounting Office study looked at alternative force structures from 150,000 to 25,000 European Troop Strength in 25,000-person increments with some estimates on the influence and capability of each size force.

In 1994, the Commander-In-Chief of the U.S. European Command confronted a changing, and in many ways uncertain, mission. Nonetheless, he had to submit a force structure architecture for his command to the Secretary of Defense, the President, and Congress soon after the Gulf War, cognizant that many political leaders expected a peace dividend in part demonstrable by lowering ETS. The CinC defined the problem for his staff by using the National Military Strategy and focusing on the largest variable, the Army component, and by identifying several boundaries:

- Overall ETS would be more or less 100,000.
- The U.S. would retain leadership of NATO and that implied an assemblage of ground forces that equated to an army corps.
- U.S. forces would need to respond to several different kinds of crises, probably nearly simultaneously.
- Interest in peacetime engagement and the exercise tempo would increase, with outreach programs to Central Europe and more activity in Africa and the Mediterranean Sea than before.

By definition, an army corps has at least two divisions and a set of supporting forces. A standard army corps, however, would consume most of the 100,000 European Troop Strength goal by itself. The force planners therefore needed to see how much of the corps actually had to be stationed in Europe. Reviewing the regional situation and the Defense Planning Guidance, they identified four likely near-simultaneous scenarios requiring U.S. ground forces: (1) a peace operation in Former Yugoslavia (one division); (2) a peace operation in the Middle East (one brigade); (3) a humanitarian disaster (one battalion plus many support units), and; (4) a non-combatant evacuation operation in Africa (Southern European Task Force plus aviation elements). The forces not committed to crisis response would meet the exercise and engagement commitments.

The result was that the two divisions in V Corps in Germany have only two of their three brigades in Germany. Also, only about two-thirds of the corps-level units, those chosen for their utility in crisis response, are stationed in Europe. The resulting force, then known as the "Credible Corps," was sufficient to sustain U.S. leadership of NATO's European command. V Corps troops in Europe, plus independent and Echelon-Above-Corps Army support units, total 65,000. The Air Force sized their tactical force to support the Credible Corps and to support what is now Operation Northern Watch, which has been flown from Turkey into Iraq since 1992 (and flown primarily by rotational units from outside the theater). They packaged their logistic force for immediate reaction to the crisis scenarios, arriving at 35,000 Air Force personnel in theater. The Navy, already essentially headquarters and logistic organizations, maintained 14,000 people in Europe for a total ETS of 114,000 that was within the CinC's goal.

How well did the CINC and force planners assess uncertainty in 1994? They did pretty well. In 1996, the European Command deployed the two brigades of the 1st Armored Division to Bosnia as part of the Dayton Accords Implementation Force and the Air Force flew over Bosnia in Operation Deny Flight and later in support of the peace operations. The lack of a peace settlement in the Middle East forestalled a new peacekeeper deployment there, but the continuous presence of Task Force Able Sentry in Macedonia absorbed the forces planned for the Golan Heights. In 1996,

> the U.S. responded to the humanitarian crisis in Rwanda with a large airlift operation (requiring force protection) and evacuated several African embassies. Typically, the U.S. European Command oversaw six Joint Task Forces in 1996, a pattern that continues today.
>
> That said, what the planners did not and could not foresee was the open-ended nature of many of these commitments. The U.S. European force structure, capped near 100,000, is not deep enough to support *rotational* deployments for these deployments. As a result, force planners are reviewing ETS to decide again how much force structure the U.S. needs to maintain in Europe to maintain its leadership of NATO and continue shaping the security environment, i.e., can European Troop Strength be safely reduced if we assign crisis response to CONUS-based general purpose forces? Do we still need a corps-equivalent in Europe? How much more benefit and influence with our Allies do we derive from permanently stationed versus deployed forces?

Risk and Uncertainty Profiles

Because of the inherent differences between objective and subjective probabilities, we approach each differently to minimize its detrimental effect on our decision making. First, however, we must carefully assess which parts of the decision contain risk and uncertainty, whether there is more knowledge we can gain, and what the consequences of those risks and uncertainties are. To do this, we develop a risk and uncertainty profile. It consists of the answers to these three questions:

- What precisely do we not know that we need to know to make a decision?
- How much more knowledge can we gain about them?
- What are the consequences of these risks and uncertainties and are they important?

By answering these questions—which parallel our considerations of validity, reliability, and practicality for evaluating criteria—we categorize the unknowns and decide which are worth our further attention and whether that attention will pay off. Usually, these answers center on another examination of our criteria, an extension of our earlier validity, reliability, and practicality evaluations.

The first question above is about validity: what do we need to know vice what information is at hand, regardless of how easily we can obtain information? Left to their own devices, many analysts will provide us with that which they can expeditiously collect; we seek instead to identify an ideal—what we need if perfect and limitless information were available. After we identify what we want to know, we examine how much is knowable about each criteria and how difficult it will be to collect more knowledge—another look at reliability and the quality of our measurements and data. For that which is knowable we seek to build objective probabilities. For the other unknowns, we will or must settle for subjective probabilities.

To decide how important an item of risk or uncertainty is to our decision, we can use a technique called sensitivity analysis (explained more fully in Chapter 7). Sensitivity analysis allows us to assess the potential impact of each risk and uncertainty on the outcome of a decision by examining the results of each alternative in isolation. In sensitivity analysis, we vary the effect of a single risk or uncertainty over what we believe is their plausible range of values while holding everything else constant, and then we examine the various results. If the overall outcomes do not vary greatly from one another, the decision is not sensitive to that risk or uncertainty; if the outcome does vary greatly, then the decision is sensitive to that unknown. In other words, we ask ourselves: How bad or well could a risky or uncertain event turn out to be, and does that matter to our decision? Knowing this determines whether we need or desire further investigation, a measure of practicality we apply to our analysis.

For example, suppose we are trying to decide how to arm strike aircraft for a mission with multiple targets of varying importance. Targeteering data tells us which weapons perform best against which targets in terms of a probability of kill. We also have statistics on weapons reliability. In short, we do not know how each weapon will actually perform, but we know the objective probabilities of each weapon against each target and therefore the risk of failure of any particular weapon. We have other uncertainties: how many aircraft will reach their weapons launch points? How many weapons on how many planes shall we dedicate to the highest priority target? Too few and the target may survive the strike; too many and we may have to fly another strike against the secondary targets that we could have destroyed during the first mission. By assessing the objective probabilities associated with munitions, we decide how many weapons we need to destroy the target; then, after assessing the subjective probabilities, we decide how many aircraft and how many more weapons beyond the earlier number we will assign to the strike.

Strike planners do this analysis by combining the databases and their experience. What they seek to uncover is whether, throughout its plausible range, any particular unknown matters greatly in this decision. They focus their energy upon those risks and uncertainties that are important with regard to the objective. In a decision where even relatively small amounts of error may matter a good deal, a risk or uncertainty deserves great attention. In a decision where being generally correct is good enough, only unknowns with large impacts are of further interest to us.

Additionally, creating the risk and uncertainty profile is necessary because acquiring more knowledge about important risks and uncertainties consumes resources. This raises cost-to-benefit and practicality issues. Do the resources we dedicate yield enough new knowledge about risk or uncertainty to improve our decision and therefore justify their expense? By culling the important unknowns from those less so, we avoid wasting resources on issues with little impact, and, when resources are limited, we can prioritize intelligently.

A RISK AND UNCERTAINTY PROFILE FOR NATIONAL MISSILE DEFENSE

The risks and uncertainties currently surrounding National Missile Defense (NMD) are being hotly debated. What would a risk and uncertainty profile include when applied to NMD and how is the Department of Defense likely to manage the attendant unknowns?

The risk of failure includes virtually all the design and engineering aspects of NMD – especially technological risk. The probabilities of detection of incoming missiles, especially discrimination from decoys by the ground-based radars and the kill vehicles' infrared sensors are physically knowable and we will discover them through simulation, testing, and experiments. We will also determine the probability of kill by the interceptor once an incoming warhead is detected and tracked. In the same way, the objective probabilities of various types of failures can be assessed and specified. The program managers will likely manage these risks by modifying designs for maximum cost-effectiveness and then buying out the risks that remain by procuring enough interceptors to ensure that, considering the objective probabilities of failure along each step from detect to engagement, the aggregate risk of overall mission failure is reduced to an acceptable level.

But there is another category of factors for which our knowledge is uncertain and not quantifiable. For example, how large an attack should we prepare for? How much warning are we likely to have? Will the attacker use sophisticated tactics? Will the attacker use penetration aids like decoys and chaff? These are unknowns for which we cannot calculate objective probabilities. Their likelihood depends upon future choices by our adversaries. For this reason, they are uncertainties with critical implications for the NMD force structure. We likely will attack these unknowns by first learning as much as we can to narrow the areas of our ignorance.

Good intelligence is key. How many weapons of what kind do our potential enemies possess? What kind of command and control doctrine does each adversary use and what is their likely salvo doctrine? What steps must they take prior to launching missiles and what are the signatures from those steps that will provide us with warnings? This intelligence should narrow the range of possible attack sizes and sequencing to subjective probabilities.

Once intelligence has told us all it can, we can design a National Missile Defense force structure on the basis of further scenario analysis. We will construct a range of hypothetical attacks to determine the NMD force structure that best defeats the aggregate. The scenarios should include the worst possible case, the worst plausible case, and middle-of-the-road cases, e.g., attacks based on a terrorist (vice nation-state) attack. We can then compare the resulting national missile defense force structures on the basis of cost and effectiveness, allowing senior leaders to make informed decisions about the level of capability that they wish to fund.

Dealing With Risk and Uncertainty

There are three straightforward and popular ways of dealing with risk and uncertainty. The first is simply to continue to solve more of the unknowns, thereby reducing uncertainty. Next, we can acknowledge that there will always be some risk of failure for any alternative, and we can attempt to buy out some, or all, of this risk. Finally, if the risk cannot be bought out, we can compensate for it by adjusting the attractiveness of an alternative by incorporating risk into our calculations.

REDUCING UNCERTAINTY

Generally, when we have uncertainty we desire to calculate objective probabilities. If we can create them, then we will likely understand enough about cause and effect to know whether or not we can change those probabilities if we so desire. Improving the reliability of our objective probabilities to predict outcomes may simply require more research if the necessary data already exists but is not at hand. If the information does not exist, then we may need to conduct experiments in a laboratory or at a test range.

Whether or not we choose to invest the resources to define objective probabilities for an unknown returns us to the practicality issue: is the knowledge gained worth the resources consumed? The Navy sometimes shock tests a new ship to gauge the quality of its construction and its resilience to battle damage by detonating a large explosive charge near it underwater. The shock test is expensive and the hull flexing decreases the strength of the ship tested by making the hull more brittle (for the same reason one should buy a new motorcycle helmet after an accident). Is the knowledge gained worth the cost? Sometimes. Therefore, the Navy shock tests the lead ship of every large class and additional ships if they make major design changes. We improve our understanding by uncovering objective probabilities when we decide it is practical to gather the additional information we need. If not, then we continue to deal with unknown probabilities as uncertainties. Of course, the same logic prevails in the use of improving subjective probabilities, which will be discussed later in the chapter.

BUYING OUT RISK

Suppose we find that not all of our alternatives provide enough certainty in their outcomes. We may be able to use our second technique of risk reduction: we buy out some of the risk. We ask ourselves (or, more likely, our analysts) whether additional resources could reduce or eliminate the risk of failure and, if so, what is required? The answer is usually more money, time, or equipment. For example, once we know the objective probabilities, we can reduce the risk of a failed air strike by increasing the number of aircraft assigned or the number of weapons they expend. Both involve increasing resources. In effect, we convert risk into something else, in this case, weapons systems. Note that we have not changed the objective probabilities that any particular aircraft or weapon will accomplish the mission. But we have reduced, or bought out, the risk that the mission will fail by increasing the resources devoted to it based on our knowledge of those objective probabilities. This approach to managing risk is quite common and intuitive. We encounter engineering redundancy, another good example, all the time in our professional and personal lives.

Again, note that we buy out the risk by transforming risks into resource consumption. We can sometimes buy out the risk to compare different alternatives on a common basis. For example, suppose we need to choose one of three designs for a new aircraft program. Each alternative has families of risk, benefit, and cost criteria. For this illustration, we selected one representative criterion from each of these categories to illustrate buying out risk: probability of a major mechanical malfunction during a mission (risk), maximum speed (benefit), and unit price in constant dollars (cost).

Alternatives	Risk (mechanical failure)	Benefit (maximum speed)	Cost (production price)
A	0.20	Mach 1.5	$20 million
B	0.15	Mach 1.0	$25 million
C	0.05	Mach 0.8	$30 million

Table 5-1. Three Aircraft Alternatives with Unequal Risk

Based upon table 5-1, which aircraft is the most prudent purchase? We can eliminate some of the complexity of this decision by buying out the risk of a major malfunction associated with Design A and B and adding the resources needed to do that to each of their costs.

Alternatives	Risk (mechanical failure)	Benefit (maximum speed)	Cost (production price)
A	0.05	Mach 1.5	$32 million
B	0.05	Mach 1.0	$30 million
C	0.05	Mach 0.8	$30 million

Table 5-2. Three Aircraft Alternatives with Equal Risk

Table 5-2 enables us to compare designs more simply on the basis of cost and effectiveness (only) by translating risk into cost. In so doing, the least expensive alternative can become the most costly while an expensive alternative can, in the light of risk, become a bargain. In this example, we have elevated the mechanical reliability of aircraft A and B to the same level as aircraft

C. Note that their benefits—top speed—remain unchanged. Now our decision is reduced to whether we value the higher speed of aircraft A enough to pay $2 million more per copy than we would for aircraft B.

EXPECTED VALUES

This leads to the third procedure we use to deal with risk. Suppose all the alternatives involve significant risk, and that the cost of buying out the risks to equal levels is too high. Or, suppose our circumstances call for simply accepting risk and choosing the best overall alternative, even though their benefits vary and each carries a differing level of risk, e.g., one has the highest level of risk and the greatest effectiveness. For situations like this we use the Expected Value approach. Like buying out risk, it is a way of adjusting the attractiveness of an alternative to reflect its probability of success.

Expected value computations can become complex and we need not become fully conversant with the mathematics involved. But, because we may have to compare alternatives based on expected values (calculated by someone else), we will familiarize you with the basic concepts here. Expected value computations use the concept of utility that we discuss more fully in Chapter 6, "Combining Criteria." Utility provides a way of translating the different attributes of alternatives into a common unit of measure that reflects their usefulness or value with respect to the decision objective. We can quantify these values for each alternative under each criterion and sum them to make direct comparisons. To obtain the expected value of an alternative, we multiply its utility (cost or benefit) by its probability of occurrence (risk).

For example, suppose a lottery prize is worth one million dollars, and one has a 100 percent chance of winning the lottery—only one ticket will be sold. The expected value of this ticket is one million dollars ($1M times 1.0) minus the cost of the ticket. Anyone who paid more than $1M for this winning ticket was unwise. Now suppose each ticket has one chance in two million to win the one-million dollar prize and tickets cost one dollar apiece. Is a ticket a cost-effective purchase? Because the chance of winning is one in two million, the expected value of a ticket is fifty cents: the benefit ($1M) times the probability of winning (0.0000005). The lottery makes $.50 on every ticket sold. To be cost-effective for the ticket purchaser, the ticket would have had to cost less than $.50, but, of course, no lottery could stay in business on that basis.[1]

We apply the same principle of expected value in defense decision making. For example, suppose we must select a weapon system alternative. One has a utility of 50 if all the subsystems perform as specified, but there is a 30 percent risk that they will not. The other weapon system has a utility of 40, but with a risk of only 10 percent of subsystem failure. The first system has an expected value of 35 (50 times 0.7) and the second system has an expected value of 36 (40 times 0.9). The higher risk reduces the expected value of the more effective alternative below that of the less effective one. This is as far as we need go in understanding expected value. Be mindful that there is nothing magical about expected values—they simply combine benefit (or cost) and risk into a single convenient number.

Improving Subjective Probabilities

Uncertainties that we can express only in terms of subjective probabilities are more problematic to decision makers than those with objective probabilities. The fact that we cannot predict the

1. That is why "Prairie Home Companion" radio host Garrison Keillor refers to state lotteries as a tax on people who did not do well in high school mathematics.

alternatives' outcomes with certainty or that we cannot assign probabilities to those outcomes based on statistical knowledge has important implications. It is harder to think about buying out this kind of uncertainty because we do not know when we have committed enough resources to eliminate it. Nor can we use the expected value approach since no objective probability exists on which to base an expected value calculation. A variety of other approaches do exist to accommodate uncertainty and risk, but there is no escaping that uncertainty and risk limited to subjective probabilities is among the most difficult aspects of defense decision making.

BETTER INFORMATION

The first approach to reduce uncertainty is to acquire more information. Perhaps the uncertainty we face is due, at least to some extent, to ignorance that we can dispel, if not to the point of objective probabilities. Perhaps we have not discovered everything we can. We should review our information about the problem and consider additional sources. This may be as simple as going to the library, searching the Internet, or making a telephone call; or it may require expensive research. As always, the issue is practicality: cost versus benefit. Do we have good reason to believe that more information has a reasonable chance of reducing the uncertainty? Is the decision deadline looming such that by the time we obtain the information, the decision will have become a moot point?

REFINE SUBJECTIVE PROBABILITIES

We often have genuine uncertainties that cannot be resolved by obtaining more information because what we need is not knowable. We cannot know the objective probability that there will be two overlapping major theater wars. We cannot know how close we actually came to nuclear war with the Soviets during the Cuban Missile Crisis. We cannot know statistically whether an assault on a hill will succeed. We cannot develop an objective probability that host-nation support for strategic mobility will be available as we plan for it. We cannot know how a particular unit will perform in combat as a function of its training. How then can we proceed to make rational defense decisions?

We can assign subjectively-derived probabilities to uncertainty as expressions of confidence based on our experience and the analysis we have done thus far. We think it likely that an Army unit will succeed better in combat if they complete a rotation at the National Training Center. That improvement we may describe subjectively in a change of readiness from C2 to C1 in unit readiness reporting or by increasing the unit firepower scores in a wargame. Subjective probability may be valuable because it provides a way to treat uncertainties somewhat like risk and it provides a frame of reference for discussion. This is advantageous because we have seen that risks can be relatively easy to incorporate into a decision. But we must keep in mind that subjective probabilities are prone to various kinds of errors that objective probabilities are not. We can compensate for these errors and guard against them, just as we do with the flaws in our memories, but we cannot be sure that our subjective probability estimates will be accurate. The most common errors are:

- ***Wishful thinking:*** We may estimate the subjective probabilities of various outcomes based on how desirable we regard those outcomes. But, of course, the likelihood of an event has no connection to how desirable we think it is. For example, defense decision makers who plan and execute an operation tend to be more optimistic about its chances of success than individuals uninvolved. The decision making prior to the Bay of Pigs invasion of Cuba in 1961 was distorted by this phenomenon. So was the decision making prior to the 1980 rescue attempt of the U.S. hostages held in Iran.

- ***Selective perception:*** We may not include all the factors that matter when we estimate subjective probability. In a similar vein to wishful thinking, we may include only those factors that we regard as special or are otherwise notable. For example, aviators may overestimate the impact of air strikes achieving a campaign's objectives, or dismiss the use of missiles or Special Forces as viable options for neutralizing a target.
- ***Experience:*** We may bias our subjective probability estimates of all outcomes based upon our memories of similar events. The more recent or powerful our memory is of similar events, the higher we will estimate the probability of the outcome that seems important or dramatic to us. Most people estimate that the chances of an airplane crash are higher than they would be otherwise if an airplane has crashed recently. For instance, in the wake of the September 11, 2001 terrorist attacks, many Americans are traveling by car instead of plane, even though historic experience shows air travel to be far safer. Perhaps this is also one reason why it is so often said that militaries prepare for the last war. Military veterans of the Vietnam War tend to be more opposed to peace operations than other constituencies.
- ***Framing effects:*** The way we define the question may significantly influence our subjective probability estimates. For example, we may change our estimate of the outcome of an operation depending upon whether we are asked to predict the probability of success or the probability of failure, even though the two estimates should be complementary.
- ***Overconfidence:*** Perhaps the most dangerous and prevalent influence is our sense of infallibility. We are generally far too confident of our ability to personally estimate the probability of an outcome with great accuracy. In a wide variety of scenarios, individuals are repeatedly much worse at making estimates of probability than they think they are. A great deal of sound research has repeatedly confirmed this disturbing problem.

There are two fundamental methods for improving our subjective estimates. The first is to be aware that virtually all of us are prone to making one or more of the above errors when we try to estimate subjective probabilities. By being aware of them, we can compensate for them.

Second, and usually more successfully, we can involve other people in our problem solving. While virtually all of us are prone to these perceptive errors, the forms they take in each of us are likely to be different. By involving several people to estimate the subjective probabilities of the same events, the weaknesses of one participant may be offset by the strengths of another. At some point, too many participants become unproductive. Our recommendation to involve others in your decision making may be difficult for those who prefer solitary reflection and have a low regard for group problem-solving activities. Despite these common and understandable sentiments, effectively including the military judgment of others is an important part of executive decision making. That said, some ways of obtaining group views are better than others, depending upon how much time is available and how much trouble we wish to take.

DELPHI METHODS

At the very least, our choice of decision-making participants should be based on their background, availability, reputation, and often the organization they represent. There are a variety of ways to bring their ideas together usefully. The most common, popular, and quickest approach is a BOGSAT: a Bunch of Guys Sitting Around Talking. There is not much more to add. We bring the right people together, usually after providing some read-ahead material, and moderate a discussion. It is helpful to have an agenda to guide the discussion and ensure that the essential issues receive attention. A recorder is a good idea, as well.

More elaborate is the Delphi[2] method. It begins by having the participants vote anonymously on a set of proposed subjective probabilities. While sometimes there will be tight convergence of opinions, usually there is a wide difference. Next, we discuss why each participant agrees with some estimates and disagrees with others. This exposes our assumptions and arguments to the group's assessment. Informed discussion will highlight when and how a participant may be making one of the misjudgments discussed above. Others can detect and correct those errors, and estimates will change in the process.

We follow the expository discussion with another vote. Usually, we see some convergence of opinions. Depending upon how much, we may have another conversation and another vote. At some point the group's estimates will stabilize to the point where any remaining possibility of convergence is not worth the effort to obtain it. Ultimately, we may obtain a consensus estimate or we may get two or three clusters of estimates. Occasionally, no convergence happens at all. In any case, we have important information about what people whose expertise and judgment we trust think about a critical risk or uncertainty. We learn what they believe are the cause and effect relationships that shape uncertainty, what assumptions they think carry important weight, and the direction the uncertain outcomes may take. In most cases, this information is more valuable than what you could have developed ruminating alone. Naturally, we can attain much of this knowledge, somewhat degraded, on a less formal basis.

Now that we have obtained subjective probabilities that are as informed as possible, we can begin to treat the unknown as if it had objective probabilities. For example, we can assess whether one alternative involves much less uncertainty (as expressed by subjective probability) than the others. If so, we may select that alternative to avoid uncertainty, if, at the same time, we can satisfy our minimum requirements for effectiveness and cost. Similarly, we can consider buying out the uncertainty. In this case, we can convene our group and ask them to make subjective probability estimates of how additional resources will affect the risks and uncertainties. Again, keep in mind that all these judgments are completely subjective. We cannot have the same confidence in our outcomes as we do when working with objective probabilities.

WORST-CASE SCENARIOS

Armed with subjective probabilities, we can buy a hedge against the most plausible and important range of outcomes, an investment against undue loss due to the uncertainties. One way to do this is by choosing alternatives with the flexibility to cover the range of outcomes that matter to us. But, like all capabilities, flexibility is not free. It may make wonderful sense to select an alternative that allows us to achieve our goals in a variety of circumstances, and the financial cost for this flexibility may be straightforward, but other costs may be subtler. For example, by preparing to respond to many situations, e.g., to achieve *Joint Vision 2020*'s Full Spectrum Dominance, we may not be particularly well trained for any. This is our concern for general-purpose forces like infantry battalions, multi-role fighters, and ships. Because we train them for many eventualities, our training costs go up while our readiness for warfighting may simultaneously

2. Delphi was the location of the Greek god Apollo's shrine where oracles tried to predict the future.

degrade. Although these kinds of cost may not be readily apparent, we must include them. We pay all costs, whether we know it or not.

Another classic way of using subjective probabilities is to select the alternative that is most effective against the worst plausible outcome, the worst-case approach. The crucial assumption we have to make, if we choose a worst-case alternative, is that all other plausible and important outcomes can be subsumed in the worst one. That is, if we can handle the worst case, then we must be able to handle less dire circumstances by definition. We know, however, by logic and experience that this assumption may not be true or may be only half-true.

During the Cold War, U.S. conventional forces focused on stopping a Warsaw Pact thrust across the inter-German border. The scenario that force planners envisioned required large-scale, high-intensity operations against the Warsaw Pact. Forces and concepts designed to stop an armored juggernaut in Europe were not well suited for the other applications that were originally thought to be less stressing. Vietnam is a case, or several cases, in point. Before the U.S. military's frustrating involvement, the French discovered this painful truth. The Soviets, of course, were effectively defeated in Afghanistan and later in Chechnya using forces meant for a war against NATO.

The lesson here is that the worst-case approach to handling uncertainty can be sensible and in some ways efficient, but we must take great care to be sure the worst-case assumptions are realistic and acknowledge when they do not transfer to other circumstances. When we have limited resources, the worst-case approach may form the basis for our force structure, but we need to include other capabilities when we know we will confront other circumstances. Our limited resources in the 1990s and our emphasis on the worst case (two overlapping major theater wars) has forced DoD planners to make just that kind of difficult decision and has resulted in today's High Demand/Low Density units, e.g., civil affairs, military police, tactical control elements, reconnaissance and air surveillance aircraft. Because of uncertainty about the future, they had to decide whether to accept more risk of failure in major theater wars by building more active duty combat support and service support units to support peace operations or whether to maintain the more traditional focus on warfighting. They chose the latter because the consequences of failure were so much higher even though the likelihood of peace operations was much higher.

EXPECTED VALUE WITH SUBJECTIVE PROBABILITIES

Finally, we can apply the expected value approach we used earlier to refine objective probabilities, now using subjective probabilities—cautiously. For example, most planners would say that the chance of nuclear war with the Russians is very low. That is a subjective probability. There is no objective or statistical way to know this since we have not had a series of nuclear war preconditions to tabulate. If we deem the probability so low, why do we spend all the resources that we do on nuclear forces? The answer is that, although the probability seems low, the consequences of being wrong are astronomical. This is an expression of expected value using a subjective probability. When we multiply the huge negative value (utility) of nuclear war by the small probability of its occurrence, the negative expected value is still far too large to ignore, so we continue spending resources on nuclear deterrence. Depending upon how confident we feel about that negative expected value, we may even gain some sense of how many resources are worth devoting to this mission.

We can follow a similar process with any other decision involving subjective probability. Our confidence in the resulting expected value depends upon how confident we feel in the judgments behind it and if the analyst does not tell us, we need to ask. Using these procedures, we can begin to unpack the problem of how many resources we should expend to be ready for a second overlapping major theater war. In a similar fashion, we can address how much effort we should expend to offset the uncertainty that host-nation mobility support will not be available.

ASSESSING RISK IN PREPARATION FOR THE QUADRENNIAL DEFENSE REVIEW

Since 1996, each new Presidential Administration during its first year in office has been required to report to Congress its defense strategy and the force structure and programs it requires to execute that strategy. The 1997 Quadrennial Defense Review (QDR) was roundly criticized for having a serious disconnect between its strategy and the force structure it identified to execute it. The latter was largely constrained by budget considerations and has been noticeably frayed trying to execute the former with a $50B per year funding shortfall.[3]

To facilitate rapid execution of the 2001QDR, the Chairman of the Joint Chiefs of Staff commissioned a study group of four field grade officers, one from each service, led by Michele A. Flourney at the National Defense University (NDU) to identify issues and options for the new administration as it conducts the QDR.[4] The group observed that within DoD, we do not have a commonly accepted risk management framework.

They proposed a structured methodology for examining risk. The NDU team proposed that decision makers evaluate what they called the strategic military risk of each force structure alternative; in other words, they proposed examining each force structure alternative's efficacy executing a national military strategy (vice a national security strategy which employs diplomatic and economic tools as well). Risk, in their terms, was the risk of failure.

Ms. Flourney's group broke strategic military risk into three categories[5], each having two subcategories as shown below:

Operational risk is how well (or poorly) a force structure alternative achieved the current military strategy. Force performance is the U.S. military's ability to achieve military objectives in support of war plans and peace operations; force sustainability is how well the military maintains its readiness over time across the spectrum of conflict, from engagement and presence to humanitarian operations, for peace operations and crisis response, and through major theater wars.

Force preparation risk is how successfully (or poorly) the military prepares for future operations, based primarily on future force structure and doctrinal choices and procurement strategies. Transformation risk refers to force structure designs for the most likely scenarios while hedging risk concerns less likely but still possible scenarios like a resurgent, expansionist Russia that threatens NATO and forces a return to Cold War practices.

Affordability risk evaluates whether the force planning choices DoD makes are affordable, first concerning the allocation of resources within DoD and then considering DoD's portion of the over-

3. Congressional Budget Office, *Budgeting for Defense: Maintaining Today's Forces* (Washington, D.C.: Government Printing Office, Sep. 2000).
4. Michele A. Flourney, Report of the National Defense University Quadrennial Defense Review 2001 Working Group (Washington, D.C.: National Defense University, Nov. 2000), 8. The remainder of the descriptions of the groups work, including illustrations, are derived from pp. 31-32, 49-52.
5. The actual Quadrennial Defense Review Report (30 September 2001, Washington D.C.) ultimately settled on four dimensions of risk: Force Management Risk, Operational Risk, Future Challenges Risk, and Institutional Risk. However, it (in essense) dealt with risk in the same fashion as the Flournoy Group.

all federal budget. By evaluating and then aggregating all six elements, one can derive an overall sense of how much strategic military risk a strategy funded at a given level of resources entails.

The study group makes four points about this process: (1) it is compatible with many different models; (2) sets aside contentious considerations about national will (casualties) and leaves that to decision makers outside the force planning realm; (3) the lack of reliable peace operations models hampers the analysis of force structure alternatives and their efficacy in various low-end scenarios, and; (4) the lack of reliable models for peace operations will hamper the QDR, but nonetheless the QDR must set a general strategic direction for peace operations to close the strategy-to-resources gap.

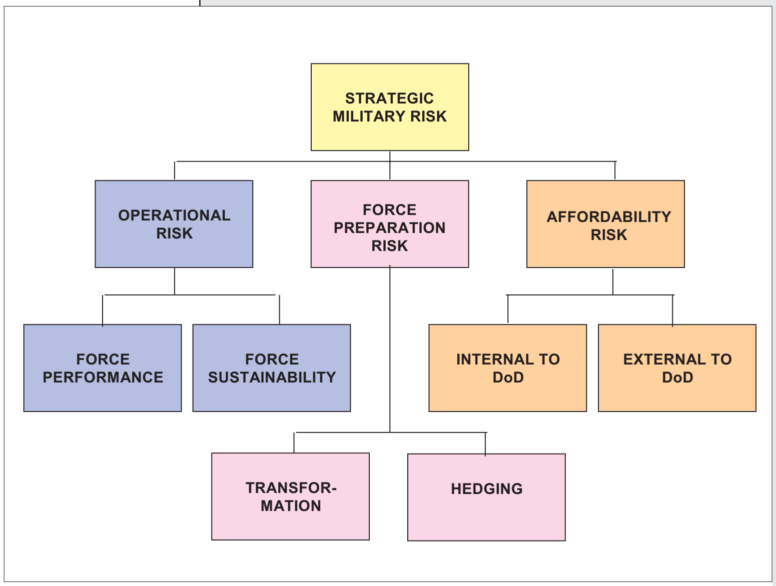

Figure 5-1. Measures of Risk Wiring Diagram Example

Each of the first four types of risk must be explored, they write, using the force-on-force analysis methods we describe in Chapter 8. After the force structure required to meet each strategy-driven situation is identified, planners can derive a force structure and its cost. Presumably, a zero-risk force would be fully funded and be able to accomplish all its goals simultaneously and almost instantaneously. Since that level of resources is very unlikely, they conclude, some forces will multi-tasked (as are general purpose forces today) and we may accept decreased response times by moving some capability to the reserve components. Funding may not materialize at the level DoD identifies (our current situation).

Resolving those tensions introduces risk, some of which we are living with already: the risk that there may not be enough forces simultaneously available for the maximum number of operations; the risk forces may not be available quickly enough; the risk that a high operations tempo will degrade sustainability; or the risk that force structure will not be funded adequately to transform and hedge. By studying and gaming the impacts of accepting risk in different areas, we can more intelligently decide where to accept it and plan for it, rather than watch it happen.

When applied to the QDR, the NDU group emphasizes that this risk analysis must be highly iterative before it will yield a worthwhile set of force structure alternatives. They provide a step-by-step process for each assessment in an appendix to their report; but before the process can begin, it needs a strategy with prioritized objectives. The prioritization implies where to take the risk. The study group also defines four levels of assessing risk:

- Low: failure is unlikely and the resources and time to achieve objectives is acceptable.
- Moderate: failure is unlikely, but achieving objectives will take longer and consume more resources.
- High: failure is possible but unlikely and more resources and time will be required.
- Unacceptable: failure is likely despite using high levels of resources and a lengthy timeframe.

The authors would have planners apply these ratings across each of the six types of risks above and they specify up to four criteria for each, e.g., analysts would assess force performance regarding protection of key terrain, time to achieve objectives, effects upon the enemy, and number of friendly casualties in scenarios across the spectrum of conflict. The worst level of risk assigned to a subordinate category or objective becomes the assessment for that risk, i.e., if a force structure's performance in a major theater war is low risk in the terrain, effects on the enemy, and casualties categories but high in terms of the length of the campaign, the risk under force performance for that force structure is high. Therefore, it is also high for the overall strategic military risk.

How useful is this methodology? It certainly captures the essential elements of force planning. Used as the NDU group did, risk is the counterpart of capability (the likelihood of success) and this process is very much like a bottom-up or capabilities-based force planning method. Defining risk as low, moderate, high, and unacceptable is helpful if they become standard in DoD. As the authors point out, the terms too often mean different things to different audiences. What is less clear to us is whether it is necessary or desirable to aggregate the different kinds of risk beyond the three—or even six—categories. Since they are calculated in different ways, using different criteria, evaluating a force structure as "moderate in strategic military risk" is not particularly meaningful compared to knowing it has moderate risk in affordability and low operational and force preparation risk. Likewise, two force structure alternatives may have equal strategic military risk for starkly different and meaningful reasons.

Notice, too, the subjective nature of many of the assessments: how much longer need a force conducting an operation take to slide its force performance risk from low to moderate? We can create some objective probabilities for some of these risk assessments, particularly affordability, but certainly not all. Here we are (in our lexicon) dealing with uncertainty and it will by necessity figure largely in this kind of "risk" assessment.

Summary

Uncertainty and risk are important aspects of defense decision making. Uncertainty results from our doubts about how much we know and risk stems from the possibility of failure that results from uncertainty. We evaluate them during decision making based upon how much we know about our alternatives. We create objective probabilities when we know, or can know, the statistical likelihood of an outcome. When we cannot arrive at definite probabilities for outcomes, we create subjective probabilities based on our best judgment.

The first method for coping with uncertainty and risk is to acquire more information about the alternatives by doing more measurement. We may improve our comparisons among alternatives by buying out risk—expending more resources to translate the risk into something else, like cost—if we have confidence in our objective or subjective probabilities. We can also address uncertainty and risk by using an expected values approach wherein we adjust the attractiveness of an alternative's cost and benefits by tying them to its probability of success. Finally, we may reduce uncertainty further by involving other people's expertise in our decision making to get their views on the information we have available.

Risk and uncertainty are ever-present in defense decision making. Indeed, they often dominate it. Our senior leaders make many major, high-level decisions despite distressing levels of uncertainty. As we get closer to procurement and operational matters, we tend to deal increasingly more with objective probabilities; strategic choices are invariably clouded by differing evaluations of subjective probabilities. Executive decision makers must understand both as they proceed through the Analysis Phase.