Simple Risk

Ted Lewis
Center for Homeland Defense and Security
Dept. of National Security Affairs
Naval Postgraduate School

Slide 1

Opening Credits

Music

Slide 2

Simple Risk

Ted Lewis

Slide 3

Topics

In this lecture I define simple risk as expected loss, which is computed by multiplying likelihood by consequence. This definition is applied to the telecommunications sector using historical data on customer-minutes lost due to telephone outages.

While simplicity is a virtue, this definition is shown to have limitations, especially when applied to human-made incidents that do not occur strictly by random accident. In addition, estimates of likelihood and consequence depend on historical data, which is often lacking.

I expand this definition to include human-made incidents, such as terrorism and crime. Likelihood is further broken down into threat and vulnerability. Threat is essentially a measure of intent; vulnerability is essentially a measure of asset weakness; and consequence is essentially a measure of damage.

Human-made risk is illustrated by application of the formulas to the Oklahoma Federal Building Bombing carried out by Timothy McVeigh and Terry Nichols in 1995.

An extension of risk, return-on-investment (ROI) incorporates a cost-benefit factor. Since the nation does not have infinite resources, it is important to consider the return on risk-reducing investments. Risk ranking is inadequate. If cost is incorporated into the formulas, risk minimization is preferred to risk ranking, because minimization considers return-on-investment. Risk ranking typically does not.

Risk-informed decision-making can be done in two ways: either by calculating risk and then ranking assets according to risk, or by calculating ROI and minimizing risk by careful allocation of investments across collections of assets. The difference is illustrated by examining the debate over retrofitting the San Francisco Bay Bridge versus replacing it entirely.

Slide 4

Definition of Simple Risk

What is risk? Perhaps the simplest definition is "expected loss". Mathematically, expected loss is the product of likelihood times consequence. That is, likelihood is the probability of an event, incident, attack, or accident, which I call a 'fault'. Likelihood is denoted by the Greek symbol, gamma. Assuming there are many assets that may be damaged by some kind of attack or incident, each asset is numbered, shown here by its numerical index, 'i'. Therefore, gamma sub-i symbolically represents the likelihood that asset i will suffer damages from some hazard.

Likelihood and consequence have little meaning without a target and a threat. A bridge, for example, may suffer $1 million in damage because of a bomb, but only $50,000 in damage because of a wind storm. Estimates of likelihood and consequences are relative to the asset and threat. This is why I often refer to the asset-threat pair, as shown here by an ASSET box above a THREAT box, connected by a line.

Once the asset and threat are known, estimates of likelihood and consequence make sense.

Slide 5

Multiple Asset-Threat Pairs

When more than one Asset-Threat pair exists…. Which is the most typical case …. Aggregate risk is the sum of risks of the individual asset-threat pairs.

Asset #1 might be a bridge that is threatened by a car bomb; asset #2 might be a building that is threatened by a hurricane, and so on. Assets differ, threats differ, and so the corresponding likelihoods and consequences differ. Each asset-threat pair has a corresponding risk. Aggregate risk is the sum across all asset-threat pairs.

This model also assumes threats, vulnerabilities, and consequences are independent of one another. This is not always the case, but for now we simplify. Later on I will describe a fault tree model that permits more complex interactions among asset-threat pairs.

Slide 6

Multiple Threats

Risk may also be the sum of risks due to multiple threats to a single asset as shown here. Assuming each threat is independent of all others, and only one incident at a time occurs, we can aggregate simple risks across all threats by summation, as before.

This model is very common, but it has limitations which I address later in the course. For example, how do we represent the possibility that two attacks occur at the same time? A telephone pole can fall down at the same time that a cell tower or fiber optical cable fails, which is not represented by this model.

A second deficiency of this model is the assumption that threats are independent. In other words, does the likelihood of a Bomb attack affect the likelihood of a hurricane, cyber exploit, or other hazard? This model assumes there is no correlation between attack types.

Slide 7

Illustration: Telecom Outages

Aggregated simple risk is easy to calculate. Consider the real-world problem of estimating risk due to a variety of hazards to the telephone system of the United States during the 1990s. Consequence is measured in customer-minutes lost due to an outage of some sort. Outages are enumerated in the Table shown on the right hand side of the screen. These numbers were provided by Richard Kuhn in 1997.

Given the frequency of each type of outage and consequences in terms of customer-minutes lost, we can estimate risk through a series of data reduction steps. Risk is calculated by summing likelihood times consequence as given by the table. Aggregate risk can then be calculated as follows.

Slide 8

Telecom: Expected Loss

What is the risk to the Telecom sector? In this chart each hazard or threat is numbered (see column one). The number of outages caused by each hazard or threat is given by Kuhn (see column two). I obtain an estimate of the likelihood of each outage type by dividing the number of each outage type by the sum, which in this case comes to 303 incidents (see the bottom of the second column).

Expected loss due to each outage type is obtained by multiplying likelihood calculated in column three by consequences listed in column four. For example, the first entry in the last column is the product of .043 times 2.8 million customer-minutes. The result, 0.12, is rounded off and shown as 0.1 in the last column.

Aggregate simple risk is the sum of entries in the last column, shown here as 56.4. Because consequence is given in millions of customer-minutes, the aggregate risk is 56.4 million customer-minutes. This amounts to 4.65% of the total customer-minutes lost, so the risk is slight. The total number of customer-minutes delivered is not shown, but assuming it is many times larger than the minutes lost, we conclude that the telephone system was highly reliable throughout the 1990s.

Slide 9

Limitations of Simple Risk

Simple risk is very useful and easy to apply, but it has obvious limitations. First and foremost, it relies on historical data to come up with estimates of likelihood. What if this data is unavailable? In many cases we depend on subject matter experts, simulation models, or red-team exercises to estimate likelihood and consequence.

Perhaps the biggest limitation in this approach is the assumption that threats occur at random. Terrorists and criminals perpetrate incidents rather than leave them to chance. Although it can be shown that terrorism is statistically similar to other natural incidents, a better estimate of risk is obtained if we know the criminal's intent. If we have intelligence information about a possible attack, likelihood can be estimated by a combination of attack probability and asset vulnerability.

Finally, simple risk ignores prevention and mitigation tactics that may be in place to deter or prevent a natural hazard or human-made attack. There may be prevention tactics in place that reduce consequences, for example.

These and other factors are considered in great detail as the simple model is expanded and elaborated.

Slide 10

Elements of Human-Made Risk

Consider the special case of a human-made incident. What is the risk formula when intent is one factor in the equation?

Let threat, t, be the probability of an attack. Specifically, threat is a probability that estimates the likelihood that an attack will be attempted, but it does not estimate the likelihood of success. This is the job of vulnerability, v.

Vulnerability is the conditional probability that an attack or hazard will succeed, if attempted. It describes a condition of the asset, while threat describes the intent of the threat.

Consequence is defined as before. It can be measured in units of lives, dollars, time, customer-minutes, etc. Risk is in the same units. Therefore, if consequence is in dollars, risk is as well.

Each of these elements are difficult to estimate. Threat may be estimated by experts or deduced by intelligence analysis. Vulnerability and consequence may be similarly obtained from experts or historical data. Nonetheless, estimating probabilities and damages is a major challenge when applying simple risk analysis.

Later on I describe other methods of estimation: modeling and simulation are used to obtain estimates of these elements by simulating the infrastructure system or situation; game theory is also used to estimate threats and vulnerabilities by assuming the attacker tries to maximize risk while the defender tries to minimize risk. More on this, later.

Slide 11

Calculating Simple Risk

In this model, simple risk is the product of threat, vulnerability, and consequence. Essentially, I have replaced likelihood with threat times vulnerability.

As I said before, threat is a probability that estimates the attacker's intent and capability. Vulnerability is a property of the asset. If the asset is relatively unprotected, its vulnerability is high. If it is defended, its vulnerability is low.

Threat and vulnerability are numbers between zero and one, because they are probabilities.

Consequence remains as before: it is measured in terms of damages: lives, financial loss, time, etc.

Aggregate simple risk is as before: the sum of individual asset-threat pairs, or perhaps the sum of risks due to multiple threats.

Slide 12

Example: Car Bomb

As an example, consider the attack on the Murrah Federal Building in Oklahoma City, April 19, 1995. Timothy McVeigh and Terry Nichols packed a Ryder Truck with nearly 5,000 pounds of fertilizer, nitro-methane, and diesel fuel, drove it to the Murrah Building at about 9am in the morning of April 19, 1995, and parked the rented truck under the day-care center portion of the building.

The explosion took out 1/3 of the building, killed 168 people, including children, injured another 680, damaged 324 other buildings over a 16 block radius, and destroyed 86 cars. 665 rescue workers were required to respond and rescue survivors. Estimated damages totaled $652 million, including loss of life and limb. It was the largest terrorist incident in U.S. history prior to the attacks of 9/11, and the largest FBI criminal investigation up to that time. Nearly one billion pieces of evidence were collected and 28,000 interviews conducted.

Applying the simple risk formula for incidents of this size requires estimating threat, t, vulnerability, v, and consequence, c. In hindsight, the threat was high – perhaps 100%, because McVeigh was retaliating for the federal government's handling of the Waco Siege (1993) and the Ruby Ridge incident (1992). Indeed, McVeigh's attack was timed to coincide with the second anniversary of the Waco Siege.

Prior to the attack, this building was considered relatively secure: the WBDG -- Whole Building Design Guide -- gave it a LEVEL 4 security rating out of 5, where 5 is the most secure level. Assuming levels are evenly distributed, it would be reasonable to assign a vulnerability of 20% prior to the attack. Finally, if we use the reported damage of $652 million as consequence, simple risk prior to the attack was 20% of $652 million, or $125 million.

One way to think of this number is as the fractional loss when a bomb destroys 20% of the buildings, annually.

Another way is to think of this risk as a gamble that the entire $652 million will be lost with 1-in-five odds.

Thus, expected lost may be construed in two different ways: one as average annual loss, and the other as 1-in-5 odds that $652 million will be lost each year.

Regardless of your interpretation, risk is still the same: it is the expected loss due to a certain threat represented as an asset-threat pair.

Slide 13

Return on Investment

Simple risk requires only two or three numbers; likelihood and consequence, or threat, vulnerability and consequence. But it lacks an important element: a measure of effectiveness.

If a threat or hazard is mitigated by some protective or responsive measure, then risk can be reduced. But this costs money. Suppose the cost of reducing or eliminating risk is known, and designated E for elimination cost. The effectiveness of investment E can now be computed as ROI: return-on-investment.

ROI is simply the difference between risk before investing E and risk after investing E, divided by investment E. This is designated as delta-R over investment, representing the benefits derived by a reduction in risk. This measure is intuitive – it is a measure of "bank-for-the-buck".

Risk reduction can be achieved by reducing threat, vulnerability, or consequence…. Or all three.

Threat, vulnerability, and consequence reduction is achieved through a variety of methods. Political, military, and other methods reduce intent, which reduces threat.

Target hardening, redundancy, and resiliency all reduce vulnerability. The details of "how much" and "how expensive" are left for later.

For example, target hardening, faster and more capable rescue and response capabilities, and better planning all reduce consequences.

In any case, reductions typically cost money, E. The effectiveness of this expenditure is expressed as ROI. Therefore, it makes sense to invest E in measures that increase ROI as much as possible. This is the key idea underlying risk minimization.

Slide 14

ROI Example

Consider the earthquake retrofit problem facing commuters using the East Span of the San Francisco Bay Bridge, pictured here in the upper right. After the 1989 Loma Prieta earthquake Caltrans proposed the replacement of the East Span with a modern bridge costing $2.6 billion,

shown here in the lower right. This proposal was controversial because 65% of the cost would be paid for by toll increases and taxes.

Opponents claimed that retrofitting of the existing span would cost only $50 million….. Considerably less than the multi-billion replacement estimate made by Caltrans. Retrofitting against earthquakes was the strategy being applied to other bridges in the state, so why not this bridge?

Experts say the annual likelihood of a 7.0 earthquake taking down this span is 10%, because 7.0 or larger earthquakes occur roughly every 1,500 years, and if the life of the bridge is 150 years, there is a 10% chance of destruction by an earthquake in each year of the bridge's life. Thus, likelihood is 10%.

To estimate risk, multiply likelihood times consequence and get $260 million. In other words, the community is risking $260 million each year of the 150 year life of the bridge.

However, if vulnerability can be reduced to zero with an investment of $50 million, return on investment is $5.2 million. In this example, return on investment was obtained by vulnerability reduction: the bridge was hardened against the earthquake. The community's return on investment is $5.2 million each year of the bridge's life.

Furthermore, if ROI for some other bridge is lower, say $3.0 million, then it makes more sense to invest in this bridge than the other bridge. If we get a bigger bang for the buck, then it only makes sense to invest in the other bridge. This is the idea of risk minimization. When funds are limited, risk minimization means using ROI to make decisions. The Department of Homeland Security calls this "risk-informed decision-making", and the strategy is called a "risk-informed strategy."

SLIDE 15

Risk-Informed Decisions

How do we reduce risk when resources are limited? One risk-informed decision-making strategy advises us to estimate risks among all assets threatened by human or natural hazards, rank them according to highest-to-lowest order, and invest in the highest-risk assets, first. Work your way down the ranked list and when you run out of money, stop. This strategy applies to both prevention and response. Target hardening typically means reducing vulnerability by erecting barriers, putting up surveillance cameras, checking backgrounds, and posting guards. Response typically means increasing capabilities for rescuing people, catching criminals, etc.

An alternate strategy calls for minimization of risk subject to budgetary constraints. Instead of ranking assets according to risk, we optimally allocate resources across all assets so that the sum total of risk is reduced to its minimum. This is computationally more difficult, so a computer program is used to perform the calculations.

If the objective is to reduce risk across a collection of assets, say bridges, telecommunications equipment, cities, or counties, then the ROI-informed strategy is preferable, because it gives the best-possible return on investment. That is, it allocates limited funds to return the greatest benefit for the amount invested. It also takes the politics out of the equation, because cost-effective spending is hard to argue against.

ROI-informed prevention essentially reduces vulnerability across a collection of asset-threat pairs such that aggregate risk is minimized. When applied to response, consequence reduction reduces risk. Recall that risk is threat times vulnerability times consequence, so reduction of any one or more of these elements of risk also reduces risk.

The risk-informed strategy uses ranking of assets, while the ROI-informed strategy uses optimization techniques. We only need two numbers to rank order risk: likelihood and consequence. However, we need three numbers to apply the ROI strategy: likelihood, consequence, and an estimate of what it costs to reduce either vulnerability or consequence. A computer is also useful for doing the minimization calculations.

SLIDE 16

Strategy Options

What is the best strategy for protecting critical infrastructure? At this point I am unable to give a definitive answer. However, here are some considerations.

Risk ranking is a rudimentary strategy for reducing the worst-case scenario, but it may not be the best use of funds, because it may lead to non-optimal allocation of funds. A high risk asset-threat pair may receive most of the funding even though risk reduction costs are also very high. Some asset-threat pairs may not be worth the investment, because their ROI is low. Remember that the objective is to reduce risk in the aggregate, which must consider all allocation combinations.

ROI-informed risk reduction is optimal in terms of minimizing expected loss across all asset-threat pairs. Inexpensive remedies are given higher priority than more expensive remedies, leading to optimal allocation of limited resources. High risk asset-threat pairs may cost too much to reduce, which tips the balance in favor of less expensive medium risk asset-threat pairs.

Because resources are limited, it may be necessary to spend more on low risk assets than high risk assets in order to minimize aggregate risk.

SLIDE 17

Strategy Alternatives

Risk is only one of several metrics for making decisions regarding critical infrastructure protection. I will introduce you to more approaches as the course evolves, but for now consider two alternatives.

Resilience hasn't been defined yet, but intuitively it involves hardening of the asset-threat pair, recovery after an attack or incident, and adaptive behavior, which is an entire subject in itself. Resilience may be achieved in a variety of ways. One novel approach is restructuring, whereby an infrastructure such as the electric power grid is restructured over time to make it more resilient. Another approach is redundancy, whereby backup assets stand ready to replace failed primary assets. A third approach is deception, whereby decoys are used to distract an adversary.

Investment tranches offer an alternative to the ‘all-or-nothing’ approach implied by risk reduction. An investment tranche is an installment payment. Risk might be reduced incrementally and across the board by periodic investments as shown here in terms of a bar chart. For example, a city might invest $10 million per year for each of 5 years, instead of $50 million in the first year. This approach is more adaptable to changing threats because targeted funding is spread over a number years, and allocated differently, depending upon the threat.

These and other strategies are explored in subsequent modules.

Slide 18

Closing Credits

Music