Advanced Sciences and Technologies for Security Applications

Carl S. Young

Risk and the Theory of Security Risk Assessment

Springer

Advanced Sciences and Technologies for Security Applications

The series Advanced Sciences and Technologies for Security Applications comprises interdisciplinary research covering the theory, foundations and domain-specific topics pertaining to security. Publications within the series are peer-reviewed monographs and edited works in the areas of:

- biological and chemical threat recognition and detection (e.g., biosensors, aerosols, forensics)
- crisis and disaster management
- terrorism
- cyber security and secure information systems (e.g., encryption, optical and photonic systems)
- traditional and non-traditional security
- energy, food and resource security
- economic security and securitization (including associated infrastructures)
- transnational crime
- human security and health security
- social, political and psychological aspects of security
- recognition and identification (e.g., optical imaging, biometrics, authentication and verification)
- smart surveillance systems
- applications of theoretical frameworks and methodologies (e.g., grounded theory, complexity, network sciences, modelling and simulation)

Together, the high-quality contributions to this series provide a cross-disciplinary overview of forefront research endeavours aiming to make the world a safer place.

The editors encourage prospective authors to correspond with them in advance of submitting a manuscript. Submission of manuscripts should be made to the Editor-in-Chief or one of the Editors.

More information about this series at http://www.springer.com/series/5540

Carl S. Young

Risk and the Theory of Security Risk Assessment

Springer

Carl S. Young
New York, NY, USA

ISSN 1613-5113 ISSN 2363-9466 (electronic)
Advanced Sciences and Technologies for Security Applications
ISBN 978-3-030-30599-4 ISBN 978-3-030-30600-7 (eBook)
https://doi.org/10.1007/978-3-030-30600-7

This Springer imprint is published by the registered company Springer Nature Switzerland AG.
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

To Irving Young, MD
August 15, 1922 – October 3, 2016

Foreword

Organizations of all types face significant security challenges these days. Threats to both people and assets are increasing as the pressure to contain costs intensifies. In institutions of higher learning, the demand for successful and cost-effective security is coming from trustees, faculty, staff, students, and parents.

Therefore, the need for rigor in assessing and managing security risk is greater than ever. However, rigor requires pedagogy grounded in theory. If security practitioners lack such a foundation it is difficult to gauge whether their decisions are truly risk-based and the resulting security strategies are indeed cost-effective.

Carl S. Young has written a timely book that provides the theoretical basis for security risk assessments. Fundamentally, this book enables problem solving by teaching how to reason about security risk from first principles. I believe it fills a longstanding gap in the risk management literature.

Carl is an accomplished risk theorist as well as an experienced practitioner. He has tested his theories over the course of a long career that includes senior-level positions in government, industry, and consulting, and most recently at The Juilliard School as both Chief Information Officer and Chief Security Officer.

As president emeritus of one of the world's leading performing arts conservatories, I am encouraged by this fresh and long overdue approach. This book has the potential to become a standard reference for students, professionals, and academics in the field of security risk management. It should be required reading for any individual interested in the theory of security risk as well as anyone required to translate theory into practice.

President Emeritus, The Juilliard School
New York City, NY, USA

Joseph W. Polisi

Preface

Since the events of 9/11 there has been an intense focus on security risk management. Many organizations have been professionalizing their physical and information security programs by hiring staff and implementing security technologies. The overarching objective of these enhancements is to ensure the security and safety of people, property, and information.

Given the assets at stake and the investments made in protecting those assets, it may be surprising to learn that security-related decisions are often based largely on intuition. Although intuition can be a valuable by-product of experience, it is no substitute for rigorous analysis.

Perhaps even more surprising might be the persistent misconceptions about basic security risk management that pervade the industry. For example, even security professionals regularly conflate the terms "threat" and "risk." Such distinctions might seem unimportant if not pedantic. However, the misuse of basic terminology by experts suggests that the processes used to design security strategies and/or assess their effectiveness might also be flawed. The situation begs the question of why such confusion is indeed so pervasive.

In this author's view, the absence of pedagogy is a contributing factor. Formal instruction on security is generally missing from academic curricula, which is the natural place to learn the conceptual foundations of any discipline. In addition, reasoning about risk often takes a back seat to methods and technology. The reality is that the foundations of risk-based thinking have not been formalized and/or effectively communicated within the security community.

A truly rigorous, i.e., risk-based, approach to security risk assessment is especially important in disciplines where relevant data are in short supply and confirming the efficacy of solutions via experiments is not practical. In these circumstances even the most basic questions can be difficult to answer. For example, has the lack of historical threat incidents been due to the effectiveness of security controls, adversary disinterest, or just dumb luck? The answers to such questions are never obvious, but gaining the necessary insight is impossible in the absence of rigor.

That said, theory might be considered extravagant if not downright irrelevant to security professionals who must address real threats every day. Learning theory will always be a lower priority than satisfying operational requirements. Also, a professional might reasonably question why learning theory is worth the effort given that significant threat incidents are relatively rare. For example, why not just evaluate each threat scenario on a case-by-case basis and forgo the formalism?

Ironically, a low number of threat incidents actually *increases* the requirement for rigor. One phrase is particularly apropos, and succinctly captures a fundamental aspect of security risk assessment and management:

An absence of threat incidents does not imply an absence of risk.

In other risk management fields, plentiful threat incidents enable the application of standard statistical methods to calculate the probability of future incidents. Furthermore, experiments can be conducted to confirm the efficacy of risk management. Unfortunately, security problems are often constrained by a lack of data and conducting security-related experiments is impractical. However, a risk-based approach is especially needed in such circumstances, and this constrained condition is what motivates much of the material in this book.

Importantly, the motivation for a theoretical treatment is not to fill classrooms but to enhance security operations. The objectives are to assess security risk more accurately and apply security controls more effectively, efficiently, and in proportion to the assessed risk. In the end, a risk-based assessment reduces the dependence on luck, which is the ultimate objective of all risk management.

The appropriate place to both formulate and promulgate theory is academia. Yet most universities lack even an introductory course in security risk management let alone an entire curriculum.[1] It is certainly true that numerous professional certifications are available in this area. Although such courses have tactical value, they typically do not teach students how to *think* about assessing risk.

In addition, science and engineering concepts and methods are conspicuously absent from security pedagogy. This situation exists despite the proliferation of security technologies whose performance is governed by the laws of nature. Attracting more scientists and engineers to security risk management should certainly be a priority.

Finally, the contention is that a good grasp of the theory will actually help address real-world security problems by providing a logical basis for decision-making. That logic must be based on risk and apply to any threat scenario. It is the incontrovertible logic and general applicability of risk-based analyses that bridge the abstract and the practical, and thereby promote confidence in the effectiveness of security solutions so identified.

New York, NY, USA Carl S. Young

[1] A notable exception is the John Jay College of Criminal Justice of the City University of New York (CUNY).

Acknowledgments

Like many obsessions, writing a book is characterized by intermittent periods of pleasure and pain. In this case, the pain substantially increased when my father died in 2016. I wrote portions of this book in his hospital room, and it was difficult to regain momentum after his passing.

I ultimately did return to writing partly because my father would have been distraught if I abandoned this project because of him or, more precisely, because of his absence. Although the book took longer than planned, I eventually completed it in part to honor his memory.

Certain individuals helped me to regain my bearings. Specifically, my mother, Geraldine Young, has been a constant source of love and inspiration throughout my life. Her mere presence is enough to motivate my best effort. My sisters, Diane Uniman and Nancy Young, have never faltered in their support for anything I've undertaken, and this book is no exception.

My lifelong friends, Fran Davis and Jim Weinstein, David and Lisa Maass, Ruth Steinberg, and Peter Rocheleau, deserve special mention. Suffice it to say, I am extremely fortunate to have had such supportive friends for many decades. Recent events merely add to my accruing indebtedness.

Some of my parents' closest friends fall into the same category, and we have become that much closer in recent years. In particular, Bill and Vivian Seltzer have become like parents, a burden they likely didn't anticipate nor welcome at this stage in life. They and another of my family's closest friends, Sora Landes, provided companionship, food, and much-needed support during regular trips to Philadelphia. I hope these impromptu appearances haven't been too onerous, as they have meant the world to me.

I must acknowledge two cardiologists, who likely give new meaning to the term "heartfelt thanks." The first is Donald C. Haas, co-director of the Comprehensive Heart Failure Program at the Abington Memorial Hospital in Abington, Pennsylvania. Don was my father's cardiologist and, I am proud to say, has become my friend. His expertise is exceeded only by his dedication and compassion. I will never forget

the day he drove to the hospital from New Jersey *on his day off* to check on my father's condition and perhaps mine as well.

The second is Erica C. Jones, formerly of the Weill Cornell Medical Center. She successfully diagnosed and treated my own health issue that arose while writing this book. Fortunately, she too is a superlative physician and risk manager who is devoted to her patients. William Osler, an icon at my father's beloved Johns Hopkins University Medical School, clearly had physicians like these in mind when he remarked, "A good doctor treats the disease. A great doctor treats the patient."

Curious, intelligent, and motivated stakeholders are essential to developing effective security risk management strategies. I have been fortunate to work with two such individuals in my work as a consultant. Doug Maynard and McLean Pena are attorneys at Akin Gump Strauss Hauer & Feld who are both proponents and practitioners of rigorous security risk management. In addition to being a pleasure to work with, their insights have been instrumental in translating theory into practice.

The Juilliard School is a world-renowned performing arts conservatory as well as a rewarding place to work. As with any institution, it has its share of idiosyncrasies that are ingrained in the culture. I am grateful to the vice president for Administration and General Counsel, Maurice Edelson, for helping me to navigate that culture and, most notably, for having faith in my abilities.

I must thank Zhaodong Zheng ("Z"), a highly skilled Web technologist, who contributed a number of the figures. Her contributions are much appreciated, and I apologize if my random pleas for assistance have been onerous or in any way inconvenient.

Finally, Annelies Kersbergen and Anitta Camilya of Springer deserve mention. Although we have never met in person, Annelies relieved much of the pressure by allowing me to miss (many) deadlines. I believe I was able to produce a worthwhile product as a result, but ultimately, that will be for others to decide. I appreciate her understanding and guidance throughout this process. Similarly, Anitta indulged my dilatory ways and undoubtedly devoted considerable time to addressing the proof corrections and other issues. Whatever this book may ultimately achieve it will in no small way be due to her and her team's considerable efforts.

Introduction

The objective of this text is to provide the conceptual foundation of security risk assessments. This task may appear straightforward, but it is complicated by the nature of risk itself, which might explain why theoretical treatments of security risk are relatively rare. Risk is multi-faceted, and often assessments are incomplete since only one of three components is evaluated. Moreover, risk actually describes a *relationship* between the elements of a threat scenario. Therefore, issues affecting its magnitude can be subtle, and assessment results can easily be misinterpreted.

A conceptual foundation sounds abstract, but it actually provides a practical framework for problem-solving. From this framework emerges an assessment *process* that can be applied to any threat scenario. Importantly, the theory enables generalizations about the magnitude of risk, which in turn facilitates comparisons of diverse threat scenarios and the prioritization of security controls.

Ironically, a fundamental problem in assessing security risk is a limited number of threat incidents. The situation is exacerbated by misconceptions about probability and statistics. The suggestion that an absence of incidents is a handicap seems analogous to claiming a lack of disease inhibits the practice of medicine. Although somewhat impolitic, such a statement would be technically correct. Today's medical patients benefit from the collective misery of their antecedents whose ailments have yielded valuable data for both researchers and clinicians.

Medicine and security risk management are similar in that they both assess the magnitude of risk albeit within dissimilar contexts. Although there are obvious differences in the two fields, the universal nature of risk ensures that the risk assessment *process* in each case is identical. In fact, at a high level *all* risk problems are equivalent. The differences and similarities of medicine and security risk management are worth exploring.

In medicine, the threats are disease or injury, and the entity affected by these threats is the human body. Human anatomy and physiology are fortunately relatively similar within any given population. This similarity is what enables medical practitioners and medications to be generally effective despite obvious differences in our respective phenotypes and genotypes.

Consider if human anatomy and physiology varied significantly from person to person. In that case, bespoke treatments would be required for every individual. There could be no standard medical references since each patient would be the subject of a unique textbook. Moreover, identifying risk factors for diseases, which is essential to identifying treatments, would be impossible since each data sample would consist of a single individual. A limited number of drugs would be available because drug manufacturing would be highly unprofitable.

Another contributor to success in assessing health risk is the ability to conduct controlled experiments. These can isolate the effect of individual risk factors and determine the effectiveness of treatments. Statistical results gleaned from experiments enable generalizations about the magnitude of risk for the entire population. Controlled human trials and animal experimentation are the sources of much of this data.

In contrast, threat scenarios are varied, and the specific effect of individual risk factors on the magnitude of risk is often difficult to ascertain. For example, there are numerous risk factors for terrorism, and their respective contributions to the likelihood of future terrorism incidents is often impossible to quantify.

Quantifying the likelihood of any future threat incident is impossible in the absence of statistics. Moreover, even if threat incidents have occurred, the conditions that spawned such incidents must be stable over relevant time scales in order to extrapolate to the future. However, the reality is there are typically few comparable threat incidents, so any probability distribution based on this small sample would likely have a large variance. The upshot is assessments of likelihood for many threat scenarios are inherently subjective.

However, subjective conclusions based on objective risk criteria are as valid as estimates based on a sample of historical threat incidents. The difference in each case is that a qualitative estimate of risk inevitably results from the former and a quantitative estimate is possible for the latter.

Importantly, theory provides the basis for security-related decisions, and therefore it has both theoretical and practical consequences. In particular, a common frame of reference for assessing risk evolves from the theory, which is grounded in a set of core principles. These principles specify the building blocks that are common to all threat scenarios as well as the nature of the connections that link each block. The implications are profound: all security risk problems are equivalent and the general approach to security risk assessments is always the same.

It might be difficult to appreciate theory in a field that often demands decisive and immediate action. Moreover, if the theory seems too disconnected from reality it will surely and perhaps justifiably be ignored. Therefore, explicit connections to the real world are required to demonstrate relevance as well as to facilitate comprehension.

In that vein, my undergraduate mathematics professor, the late Gian-Carlo Rota, urged his students to focus on examples rather than theorems. He believed it easier to extrapolate from the tangible to the abstract rather than the other way around. Life lessons gleaned from experience have since confirmed the wisdom of his insight.

This book attempts to explain theory by providing real-world examples in addition to some admittedly not-so-real ones. The latter are frequently quite

scenario-specific and therefore not particularly applicable in general. Nevertheless, they serve a purpose, which is often to demonstrate the power inherent in certain assumptions, and the applicability of various methods once such assumptions are accepted.

Concepts relating to basic probability and statistics must accompany any theoretical treatment of security risk assessment. This requirement is often driven by the need to assess the likelihood of a future incident of a specific type or possessing a particular feature should it occur. However, it is important to understand when such methods are actually applicable. Probability and statistics can provide quantitative insights that are unobtainable otherwise, but they are not applicable to every threat scenario.

Analogies with various physical phenomena are presented throughout the text. Examples such as constructive and destructive interference are useful because they provide visual representations of abstract concepts. However, it would be a mistake to interpret these analogies too literally. That said, their prevalence suggests that perhaps security and science have deeper connections, which should be explored further.

The book has three parts consisting of 12 chapters in total. Part I, "Security Risk Assessment Fundamentals," provides the building blocks of the theory of both security risk assessment and management. These fundamentals include definitions and concepts that are required to assess and manage security risk from first principles.

Part II, "Quantitative Concepts and Methods," describes the analytical machinery that is useful in estimating the magnitude of threat scenario risk. For readers disinclined to delve into the details, Part II can be skipped without severely compromising the key theoretical concepts. However, those readers can't be given a complete pass if an in-depth understanding of security risk is the ultimate objective.

Part III, "Security Risk Assessment and Management," explores topics intended to round out the theory and demonstrate its applicability. In particular, Part III specifies a model for threat scenario complexity that is derived from elementary principles of information theory. Metrics that point to systemic security risk issues are also presented in a later chapter. Part III segues from theory to practice by presenting a risk-based security risk management process that evolves directly from the theory. Descriptions of the individual chapters are provided next.

Chapter 1 is entitled "Definitions and Basic Concepts," and as its name suggests, it specifies the basic definitions and concepts of risk and security risk assessment. The pillars of the theory that include threat scenarios, the components of risk, and risk factors are also introduced in this chapter. The distinction between probability and potential is explained, which is a key facet of the theory and represents the predicate for many of the methods described in later chapters.

Chapter 2 is entitled "Risk Factors," which is arguably one of the most important concepts in any field of risk management. Risk factors determine the magnitude of risk and mediate the relationship between threats and affected entities. Five types of risk factors are identified, and the relevance of features specific to each type is explained.

"Threat Scenarios" is both the title of Chap. 3 and the focus of all security risk assessments. Five threat scenario categories are identified and explained in detail. Phenomena specific to each category plus a security risk assessment taxonomy are also presented in this chapter.

Chapter 4, "Risk, In Depth," discusses some of the most significant themes pertaining to the theory. Many of the discussions in later chapters build on the topics discussed here. Key topics in Chap. 4 include risk universality, threat scenario equivalence, uncertainty, quantifying security risk, the effect of time on risk, direct and indirect assessments of likelihood, and risk relevance.

Chapter 5, "The (Bare) Essentials of Probability and Statistics," is the first chapter of Part II. It provides some of the concepts and methods that enable quantitative assessments of the likelihood component of risk. Although only the most basic statistical concepts are included, these are sufficient for assessing the likelihood component of risk for most threat scenarios. Ultimately, a basic familiarity with the fundamentals of probability and statistics as well as their limitations is key to rigorous assessments of likelihood.

Perhaps most importantly, the content in this chapter gives an appreciation for the inherently statistical nature of the theory, and security-related examples demonstrate the relevance of specific methods.

"Identifying and/or Quantifying Risk-Relevance" is the title of Chap. 6. This chapter complements Chap. 5, where the focus is on methods that are applicable to quantifying security risk and resulting metrics. These topics are potentially relevant to all three components of risk, which include trends, time series, and correlations. The chapter introduces the calculus, which is ubiquitous in traditional science and engineering disciplines but also has relevance to security risk assessment. More esoteric topics such as the random walk are discussed along with examples that demonstrate their potential, if narrow, applicability.

Chapter 7 is entitled "Risk Factor Measurements." This chapter provides examples of analytic methods applied to all three components of risk. Sections of the chapter focus on specific threat scenarios, and are organized according to the risk factor categories specified in Chap. 2, and are further organized according to the components of risk. Chapter 7 includes analyses of the effect of risk factor changes, which is the basis for indirect assessments of the likelihood component of risk. It also introduces fundamental concepts such as the statistical uncertainty associated with multiple risk factors, the confluence of likelihood risk factors, and risk measurements in the time and frequency domains.

Chapter 8, "Elementary Stochastic Methods and Security Risk," focuses on probabilistic methods in estimating the likelihood component of risk. The applicability of these methods is predicated on the assumption that a threat incident behaves like a random variable. Additional details associated with probability and their applicability to threat scenarios are also discussed.

"Threat Scenario Complexity" is the subject and title of Chap. 9. Complexity affects most real-world threat scenarios, and is often a significant contributor to the likelihood component of risk. Notwithstanding its prevalence, complexity is not addressed in traditional security risk assessments, and is often excluded from

academic treatments of security risk. A model of threat scenario complexity derived from elementary information theory is presented, which leads to metrics that enable comparisons of its magnitude across diverse threat scenarios.

Chapter 10 is entitled "Systemic Security Risk." Five threat scenario metrics are identified, which relate to the spatial distribution and temporal history of risk factors. Each metric is indicative of the overall approach to security risk management, which exposes potential systemic issues and the need for cultural change.

Chapter 11, "General Theoretical Results," identifies, organizes, and summarizes some of the significant theoretical results. Most importantly, these results include the core principles that represent the crux of the theory. This chapter also specifies metrics and thresholds that could be incorporated into security policies and standards. The content is organized according to the threat scenario categories identified in Chap. 3.

Chapter 12, "The Theory, In Practice," is the final chapter. It synthesizes the material developed in the preceding chapters to reveal the logic and sequencing of security risk management efforts. The fundamentals of security risk assessments are reviewed, which dovetails with the security risk management process that naturally evolves from the assessment fundamentals. Two detailed examples are presented, which help explain how the theory is applied to actual threat scenarios.

Finally, we note that the terms "security risk" and "threat scenario risk" are used interchangeably throughout the text. This is admittedly less than ideal, but it should not cause confusion since their interchange has no effect on the theory. Threat scenarios are the focus of any security risk assessment so the terms "security" and "threat scenario" are closely related if not completely interchangeable. Security risk is part of the vernacular whereas the use of threat scenario risk is more consistent with a formal treatment. In general, we will use the latter term but recognize the two are functionally equivalent.

Contents

About the Author

Carl S. Young specializes in applying science to information and physical security risk management. He has held senior positions in government, the financial sector, consulting, and academia. He is the author of three previous textbooks in addition to numerous technical papers and has been an adjunct professor at the John Jay College of Criminal Justice (CUNY). Young earned undergraduate and graduate degrees in mathematics and physics from the Massachusetts Institute of Technology (MIT).

Part I
Security Risk Assessment Fundamentals

Chapter 1
Definitions and Basic Concepts

1.1 Introduction to Risk and Risk-Relevance

Remarkably, humans make many decisions with relatively little deliberation. Issues are routinely resolved with minimal effort, and include such diverse events as crossing the street, ordering from a menu, choosing a pair of socks, buying a house and selecting a spouse.

It is tempting to believe that humans are hard-wired for decision-making and to speculate that this capability has evolved over the millennia via natural selection. Decisions that were necessary to stay alive would have been critical to the survival of our ancestors. To put it bluntly, the outcome of many decisions likely meant the difference between eating and being eaten. The ability to consistently make good decisions might be one reason *Homo sapien* survived and other species perished.

It is difficult to separate decision-making from thinking itself given the role decisions play in converting thoughts into actions. It is plausible that a key feature of human thought is a robust decision-making capability, which seems like a logical adjunct if not inherent to enhanced cognition.

Biological and social scientists generally agree that the human brain and its cognitive agent, the mind, evolved via a combination of natural selection and cultural influences.[1] However, it appears no one completely understands the relative contributions of culture versus biology in shaping the mind. The issue was famously discussed by Descartes ("*Cogito ergo sum*") and contemplated by Plato. A passage from the preeminent biologist Edwin O. Wilson eloquently expresses the complex interplay between culture and the human brain[2]:

[1]Edwin O. Wilson, *Consilience, The Unity of Knowledge (Chapter 7: From Genes to Culture);* Random House/Vintage Books, New York, 1998.
[2]Ibid.

C. S. Young, *Risk and the Theory of Security Risk Assessment*, Advanced Sciences and Technologies for Security Applications,
https://doi.org/10.1007/978-3-030-30600-7_1

> The brain constantly searches for meaning, for connections between objects and qualities that crosscut the senses and provide information about external existence. We penetrate that world through the constraining portals of the epigenetic rules. As shown in the elementary cases of paralanguage and color vocabulary, culture has risen from the genes and forever bears their stamp. With the invention of metaphor and new meaning, it has at the same time acquired a life of its own. In order to grasp the human condition, both the genes and culture must be understood, not separately in the traditional manner of science and the humanities, but together, in recognition of the realities of human evolution.

To this author, who admittedly has zero expertise in the area of human cognition, the decision-making process seems instinctual in a way that resembles language acquisition.[3,4] Humans are not taught decision-making any more than they are taught how to walk or speak a language. The process is inherent to every sentient human and is apparent at an early age. Judgment and experience modulate decision-making as humans mature.

Perhaps humans developed a simple decision-making algorithm that was easily processed by a relatively primitive brain. Even a nascent decision-making capability might have marginally increased chances of survival. Alternatively, human neurological circuitry may have evolved so that our brains became capable of processing an effective decision-making capability. Whatever the explanation, a structured decision protocol might have given *Homo sapiens* an evolutionary edge, and thereby helped to avoid a catastrophic winnowing of the species.

Thankfully, most of today's decisions do not affect personal survival much less the future of human civilization. Bad decisions do not typically carry the same evolutionary penalty they once did. However, an ironic consequence of the intellectual ascent of humans is that the results of certain decisions could lead to the destruction of civilization. Their impact mandates intelligent, responsible and ethical decision-makers in addition to rigorous assessments of risk.

The previous discussion begs the questions of what actually constitutes a decision, and what are the individual decision criteria. Simply put, a decision is a choice between possible outcomes, where the relative "goodness" and "badness" of the outcomes are evaluated as part of the decision process.

Every choice has at least two possible outcomes where the consequences of each outcome vary. Therefore, inherent to the decision process are the concepts of relative goodness and badness. Implicit in the word threat is the notion of a bad outcome. But to reiterate, "bad" is a relative term, and a choice by definition involves a gradient of outcomes whose relative goodness or badness depends on one's perspective and the decision details. Therefore, any process yielding a spectrum of possible outcomes can be considered a form of "threat" in the most general sense since some of the outcomes are *relatively* bad.

Simply put, the objective of every decision is to optimize process outcomes, which is tantamount to reducing the likelihood, vulnerability and impact of a

[3] S. Pinker, *The Language Instinct; How the Mind Creates Language*, Harper Collins, 2007.

[4] The notion that language was instinctual was first posited by Noam Chomsky as part of the theory of generative grammar.

relatively bad outcome. We will see shortly that these criteria are the same ones used to assess risk. Therefore, we conclude that *every decision is a form of risk assessment.*

At a high level, the context for both risk assessments and decision-making consists of three elements: a threat, an entity affected by that threat and the environment where the threat and entity interact. The context for decisions and risk assessments is a *threat scenario*. As noted above, the criteria for decision-making and assessing risk are identical. We will now review those criteria in more detail.

The likelihood of a threat scenario outcome, the magnitude of the effect of a threat scenario outcome, and the significance of a threat scenario outcome are the three (universal) criteria used to both assess risk and make decisions. These criteria specify the nature of the relationship between a specific threat and a particular entity affected by that threat, which exists within the context of a given threat scenario.

In the absence of such a relationship a threat is irrelevant to an entity within a defined threat scenario. The magnitude of this relationship is precisely the *risk* associated with that threat scenario.

Deciding when to cross the street is illustrative of the equivalence of decisions and risk assessments, and provides an introduction to the three decision-making/ risk assessment criteria noted above.

Pedestrians who wish to enjoy continued good health focus on estimating the likelihood of a violent encounter with an approaching vehicle as they cross the street. That estimate is based on judgments regarding the approaching vehicle's distance and speed relative to the pedestrian's speed. The likelihood of a threat incident, which in this particular instance is a collision between the vehicle and the pedestrian, is a criterion common to both decision-making and security risk assessments.

Note that the purpose of a traffic signal is to eliminate the need for such estimates. Humans can be overly optimistic in determining their chances of a safe crossing. They are also susceptible to influences that cloud their judgment such as the prospect of being late for a dental appointment. When evaluated in a more rational light, the choice between death and the dentist seems obvious. A slight miscalculation in crossing the street could be life altering if not life ending, whereas a few minutes more or less in the dentist chair will be relatively inconsequential.[5]

Given the overwhelming number of urban distractions as well as the ongoing competition between drivers and pedestrians, orderly traffic flow requires an enforceable process to regulate their interactions. Otherwise, the situation can quickly devolve into chaos as anyone who has witnessed a broken traffic signal at a busy intersection can attest.

Even in New York City where a traffic signal is viewed as more of a helpful suggestion than an absolute requirement, most New Yorkers would agree that complete autonomy in deciding when to cross the street could have disastrous consequences. Given the potential outcome resulting from an inattentive or overly

[5] http://www.nyc.gov/html/dot/downloads/pdf/nycdot-pedestrian-fatalities-by-bike-motor-vehicle.pdf

aggressive driver, prudent pedestrians look both ways before stepping into the crosswalk even if a traffic signal is functioning properly.

The inevitability of a bad outcome in any violent encounter with a moving vehicle obviates the need to estimate the second decision criterion: vulnerability. Vulnerability is the magnitude of loss or damage resulting from a decision or process outcome. At a high level there are two possible outcomes resulting from crossing the street. The first is the vehicle strikes a pedestrian, which will almost certainly yield physical impairment or the loss of life. The second outcome is collision avoidance, which nearly guarantees pedestrian health for at least the time spent in the crosswalk. Rarely does life offer such unambiguous choices.

Therefore, an assessment of the vulnerability associated with the process of crossing the street reduces to a simple choice between two extremes, which is not much of a choice for anyone with a strong desire to live. Even the most impatient individual should be willing to sacrifice a few seconds in the interest of remaining injury-free. It is therefore a constant source of amazement to observe pedestrians regularly tempt fate by prematurely venturing into the crosswalk, sometimes preceded by strollers containing infants, for relatively little gain.

The third decision criterion is impact, which is a seriously unfortunate term in this context. Impact is defined as significance-per-threat incident. The term significance has subjective overtones. Even reasonable people might disagree on the significance of a particular decision outcome or threat incident. In many threat scenarios, significance equals the monetary value of a particular outcome. For example, the significance of a theft threat scenario might be characterized by the value of stolen items. In such instances the impact risk-decision criterion can be characterized as the vulnerability-per-threat incident.

Note that the vulnerability component of risk for a street-crossing threat scenario is always significant: injury or death awaits the careless pedestrian at a busy intersection. Therefore, the impact risk-decision criterion is also significant. Most pedestrians intuitively recognize this condition and therefore take the precaution of looking both ways before crossing.

In addition, there is no way to reduce the magnitude of the impact decision criterion in any confrontation between pedestrian and moving vehicle. For example, it is impossible for a pedestrian to dilute the effect of a single collision by spreading out the injuries over multiple street crossings.

Such an approach is obviously ridiculous. However, other types of threats allow for reductions in impact by reducing the loss-per-incident. For example, the financial sector blunts the effect of a dramatic downturn in one asset class by diversification. This phenomenon is colloquially referred to as “hedging one’s bet.” Unfortunately, the irrevocable outcome of a single violent encounter with a massive, fast-moving machine mandates that the same precautions be taken each time.

Other threat scenarios further illustrate how the risk-decision process revolves around ensuring favorable outcomes, or conversely, avoiding adverse ones. Selecting an item from a restaurant menu might not appear to qualify as a threat scenario, but we now know that any process qualifies as a threat because of the spectrum of outcomes. That is, each menu selection represents a decision with

relatively good and bad outcomes just like any other decision or assessment of risk. In fact, ordering from a menu entails evaluating the same criteria as those used to cross the street.

For example, the specter of disappointment looms over any restaurant patron confronted with a choice between lobster and steak. Disappointment is a form of loss, albeit not a material one. Ultimately, the patron's selection is partly based on an estimate of the magnitude of disappointment resulting from choosing one option over another. Avoiding disappointment is equivalent to seeking happiness in this zero-sum decision process.

In addition to estimating the magnitude of disappointment, the restaurant patron also assesses the *likelihood* of experiencing disappointment with respect to each option. Somehow taste, mood, hunger and other intangibles are evaluated to yield such an estimate. Previous experiences at this or similar eating establishments might also influence the decision.

Of course, a meaningful quantitative estimate of the magnitude of disappointment, or any feeling for that matter, is not possible. A number could be assigned to rank disappointment that is based on some arbitrary scale such as when a doctor requests a rating of pain on a scale from one to ten. This assessment is useful as a relative guide, but it is not equivalent to a proper measurement using a calibrated instrument.

Personal sensations and perceptions are subjective and therefore inherently qualitative, which is not necessarily an obstacle to decision-making. In fact, some of the most impactful decisions in life rely exclusively on sentiment, intuition and/or feelings of one kind or another. For example, most individuals do not base their choice of a spouse on quantitative metrics unless money somehow figures in the calculation. The challenge is to understand when quantitative methods are required and applicable, and to identify viable alternatives as necessary.

Returning to the restaurant threat scenario, the disappointment resulting from choosing steak or lobster is fortunately confined to a single dining experience assuming there are no leftovers. Therefore, the impact decision criterion should not be remotely life altering.

Anguishing over the choice between lobster and steak seems over-the-top precisely because the difference in the two outcomes is relatively trivial, especially in contrast with threat scenarios such as crossing the street. However, the underlying decision process is noteworthy irrespective of what is at stake (or *at steak*). The point is that the decision-making process is identical no matter how trivial or profound the consequences.

Decisions of all types also occur in professional settings. Experts are constantly asked to assess possible threat scenario outcomes related to their particular area of expertise. For example, medical doctors assess the threat of disease, meteorologists assess the threat of storms and security professionals assess the threat of crimes.

Professional certifications are sometimes required to ensure practitioners are properly trained, especially in professions where the loss associated with a single threat incident could be significant. Such certifications attest to a minimum level of competence that is affirmed by examination and/or relevant experience.

For example, no prudent individual would voluntarily fly on an airplane, undergo surgery or allow the only toilet on the premises to be fixed by anyone other than a qualified professional. A basic level of proficiency affirmed by objective criteria is required in professions where incompetence could have life-altering implications. Although airplane crashes, surgical mishaps and dysfunctional toilets might appear to have little in common, they all could result in significant damage or loss.

We now know that the criteria for decision-making and assessing risk are identical, and are used to assess the likelihood, vulnerability, and impact of the spectrum of possible outcomes. The three risk assessment-decision criteria apply to any threat scenario. For example, the criteria used to determine the relevance of a disease to a community are identical to those used to assess the relevance of a terrorist group to that same community. In the former scenario, an epidemiologist is required to assess the magnitude of the three criteria associated with epidemics, and a counterterrorism expert would evaluate the same criteria for terrorism.

To repeat for emphasis, notwithstanding the fact that the knowledge, skills and abilities required to assess various threat scenarios might be different, the assessment processes are identical. The brain surgeon and the plumber evaluate their respective threat scenarios in exactly the same way, which exemplifies the principle of threat scenario equivalence. Much of the theory of security risk assessment evolves from this principle.

Although not generally viewed in this light, physicians are quintessential risk managers. The good ones can effectively assess so-called risk-relevance, which in this context is the unique relationship between a specific disease (threat) and a particular patient (affected entity). Later we will see that this assessment is partly based on statistics pertaining to *other patients'* historical relationship to this disease.

Medical diagnoses and treatments are exercises in assessing risk-relevance. However, medical knowledge is so vast that a single individual cannot possibly know the relevance of every symptom to every disease nor be aware of the appropriate treatment. As a result, sub-specialties have emerged, and each sub-specialist has an increasingly granular view of the medical landscape.

The specialist is understandably more capable of diagnosing and treating diseases that relate to his or her sub-specialty. However, this specialization comes with a price. Difficulties sometimes arise when a disease affects multiple organ systems or when a patient suffers from multiple conditions. Such difficulties are amplified if no one is managing the patient at the enterprise level.

Determining the relevance of a specific threat to a particular entity is the essence of a security risk assessment. Of course, there are many types of threats, and not every threat is relevant to every entity. The magnitude of risk is highly contextual. For this reason, threats cannot be evaluated in a vacuum. In fact, threats are mere abstractions in the absence of risk. Conversely, risk is meaningless unless the specific threat scenario elements are specified.

In addition, risk-relevance could change with time or be affected by the environment in which the threat and entity interact. As noted above, the highly contextual nature of threat scenarios requires specific expertise to accurately assess risk-relevance and thereby identify appropriate remediation.

To illustrate the importance of context, consider a fire-related threat scenario. The magnitude of risk for a fire threat scenario depends on the environment where the fire is ignited. There are certainly general guidelines for fire safety, but a meaningful fire risk assessment requires an evaluation of a particular environment. The magnitude of the likelihood, vulnerability and impact criteria could vary significantly if the fire threat scenario is a forest versus a high-rise apartment building.

Fire can be either deadly or lifesaving, even within the same physical setting but displaced in time. A campfire provides warmth in a forest in January. However, a flame in the same forest could destroy life and property during the summer months.

The somewhat obvious conclusion is that assessing risk is impossible in the absence of context, which is represented by the threat scenario. In fact, risk is not meaningful without context since by definition it specifies the relationship between threat scenario elements.

1.2 Threat Scenarios and the Components of Risk

We first provide a high-level overview of the canonical threat scenario, which anticipates additional details provided later in this and subsequent chapters.

The threat scenario is the focus of a security risk assessment and always consists of the following three elements:

- Threats
- Entities affected by threats
- The environment in which threats and entities interact

Threats are the progenitors of risk. There is no risk without the presence of a threat, and conversely, a threat is meaningless without risk, i.e., a single component of risk is zero. We will investigate this statement more carefully because of its theoretical and practical implications. For now, we present a formal definition of a threat:

> A threat to an entity is anything that results in relative harm, loss or damage to that entity. By definition, an entity is always "worse off" after experiencing a threat.

As noted in the first section, a multi-faceted feature called "risk" describes the relationship between threats and affected entities within the context of a threat scenario. Identifying the specific features of the threat scenario that affect this relationship is the crux of a security risk assessment.

Figure 1.1 shows the canonical threat scenario structure.

Threat incidents are what results from threat scenarios, and their number and/or distribution are determined by the risk factors. A formal definition of a threat scenario risk factor is provided in Chap. 2: