

Is Risk-Based Security a Failed Concept?

A Reply to Donn Parker

By Jeff Lowder – ISSA member, Puget Sound (Seattle), USA Chapter

This paper considers Donn Parker’s various arguments against risk analysis playing a role in information security management. The author concludes that Parker has not succeeded in showing that diligence-based information security should replace risk-based approaches.

Abstract

This paper considers Donn Parker’s various arguments against risk analysis playing a role in information security management. I consider one-by-one Parker’s various reasons in support of his proposal, namely, that diligence-based methods should replace risk-based information security. I show that Parker’s proposal is based upon a false dichotomy and that his various arguments, *in their current form*, have no force whatsoever. I conclude that Parker has not succeeded in showing that diligence-based information security should replace risk-based approaches.

In a 2006 article in this journal,¹ Donn Parker claims that risk-based security is a “failed concept.” Parker thinks that there is too much uncertainty and complexity in the data regarding rare incidents to effectively apply the principles of decision theory to information security. Furthermore, even if these obstacles could be overcome, Parker believes “it is too easy for management to accept security risk rather than reducing it by increasing security that is inconvenient and interferes with business.”²

While I respect Parker’s considerable expertise on information security, I strongly disagree with Parker on this topic. In fact, if followed, I believe his advice would be harmful to security practitioners and their organizations. Given Parker’s well-earned influence in the information security community, the importance of the topic, and the lack of a critical response, I want to provide such a response.

Now the whole point of risk-based security is to help organizations make good decisions about security. Risk-based security just is one example of applied decision theory.³ And among scholars there is no doubt whatsoever regarding the correctness of the fundamental theorems of decision theory. Indeed, over the last 35 years, scholars have applied the principles of decision theory to a variety of academic disciplines in fields as diverse as aerospace, economics, epidemiology, environmental protection, engineering, food safety, transportation, management science, political science, and social science. Indeed, there are now four academic journals devoted to risk analysis, *Journal of Risk and Uncertainty*; *Law, Probability, and Risk*; *Decision Analysis*; and *Risk Analysis*, with the last two journals sponsored by professional associations for risk scholars, the Decision Analysis Society and the Society for Risk Analysis, respectively. At the time this article was being written, a new professional association focused on information risk analysis was being formed, the Society of Information Risk Analysts.⁴

1 Donn Parker, “Making the Case for Replacing Risk-Based Security,” *ISSA Journal*, May 2006, pp. 6-9. Cf. idem, *Enterprise Information Security and Privacy* (ed. C. Warren Axelrod, Jennifer L. Bayuk, and Daniel Schutzer), pp. 91-101; idem, “Toward a New Framework for Information Security” *Computer Security Handbook* (4th ed., ed. Seymour Bosworth, Michael E. Kabay, New York: Wiley, 2002), p. 5-19; idem, “Risks of Risk-Based Security,” *Communications of the ACM* 50:3 (March 2007), p. 120; and his presentation, “Failed Risk-Based Security and How to Fix it,” RSA Conference (2010), <https://365.rsaconference.com/blogs/podcast-series-rsa-conference-2010/2010/03>.

2 Parker 2006, p. 6.

3 For an accessible introduction to decision theory, see Robert T. Clemen, *Making Hard Decisions: An Introduction to Decision Analysis* (second ed., New York: Duxbury, 1996); and Michael D. Resnik, *Choices: An Introduction to Decision Theory* (Minneapolis: University of Minnesota, 1987).

4 See the Society for Information Risk Analysts (SIRA), <http://societyinforisk.org>.

In his 2006 article, Parker proposes that risk-based security "... must be replaced with practical, doable security management with the new objectives of due diligence, compliance consistency, and enablement"⁵ (hereafter, "diligence-based security"). Due diligence is necessary to avoid negligence; compliance to avoid penalties; and enablement to be competitive. According to Parker, "Reduction of security risk then becomes serendipitous." In his more recent writings, Parker has avoided any implications that his diligence method attempts to meet or is related to the legal concept of due diligence, by referring to *diligence* only.⁶

Here we need to be careful to distinguish three options regarding the foundation for information security management:

1. Information security management based upon diligence-based security only
2. Information security management based upon risk-based security only
3. Information security management based upon *both* risk- and diligence-based security

Parker endorses (1), whereas I shall argue for (3). What is striking about Parker's article is that he writes as if the only

options were (1) and (2), but that is a false dichotomy. It seems to me that diligence- and risk-based security are complementary in four ways.

First, compliance often requires risk-based security. As Parker himself now admits,⁷ some laws require information security risk analysis (ISRA). In the U.S., such laws include the Federal Trade Commission Act ("FTC Act"),⁸ Gramm-Leach-Bliley Act (GLBA),⁹ the Federal Information Security Management Act of 2002 (FISMA),¹⁰ etc. Outside the U.S., such laws include the EU Data Protection Directive¹¹ and Ja-

5 Parker 2006, p. 6.

6 Donn Parker to Jeff Lowder, November 26, 2010.

7 Donn Parker, comment posted at (<http://chuvakin.blogspot.com/2009/09/donn-parkers-risks-of-risk-based.html>), April 2010.

8 FTC Act, 15 U.S.C. §45(a).

9 See GLBA, Public Law 106-102, §§501 and 505(b), 15 U.S.C. §§ 6801, 6805. Cf. Federal Financial Institutions Examination Council, "Risk Identification and Assessment" Federal Financial Institutions Examination Council (n.d.), <http://www.ffiec.gov/ffiecinfobase/booklets/mang/07.html>.

10 44 U.S.C. §§ 3541, 3544; FIPS 200, "Minimum Security Requirements for Federal Information and Information Systems" (Gaithersburg: NIST, March 2006), <http://csrc.nist.gov/publications/fips/fips200/FIPS-200-final-march.pdf>; and NIST Special Publication 800-53, "Recommended Security Controls for Federal Information Systems and Organizations" (3rd rev., Gaithersburg: NIST, May 1, 2010), http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final_updated_errata_05-01-2010.pdf.

11 "Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data" Article 17, *European Commission* (1995), http://ec.europa.eu/justice_home/fsj/privacy/docs/95-46-ce/dir1995-46_part2_en.pdf.

Start Your Adventure and Save!

Secure your place at RSA® Conference 2011, and stay one step ahead of the challenges that face the information security industry.

- **Expand your knowledge** with 200+ key industry-focused sessions, including network security, identity management, cloud computing, Web 2.0 security and more
- **Learn how new regulations** and constantly evolving compliance rules affect both the industry and your job
- **Discover what's new and upcoming** from over 350 exhibiting security companies—including ISSA
- **Network with thousands of peers**, industry leaders and security luminaries

* **Register by January 14** with qualification code **1311ISSDL15** and receive a special ISSA member discount of \$150—saving you a total of \$550 off the Standard Registration Rate!

the adventures of

alice

bob

SAVE \$550*
Register by
January 14 to secure
your discount!



RSACONFERENCE2011 
FEBRUARY 14-18 | MOSCONE CENTER | SAN FRANCISCO

www.rsaconference.com/issa



©2010 Information Systems Security Association • www.issa.org • editor@issa.org • Permission for author use only.

©2010 EMC Corporation. All rights reserved. EMC, RSA, RSA Security, the RSA logo and the RSA Conference logo are registered trademarks of EMC Corporation in the United States and/or other countries. All other marks are trademarks of their respective companies.

pan's Personal Information Protection Act.¹² Furthermore, contractual obligations can require an organization to perform a formal risk assessment. For example, many organizations are contractually obligated to comply with the Payment Card Industry (PCI) Security Standards Council's Data Security Standard (DSS), which in turn requires an annual risk analysis (RA).¹³ Along the same lines, any organizations subject to any of the U.S. state laws mandating PCI DSS compliance¹⁴ arguably have a duty to perform an RA.

Second, business enablement should properly be taken into account in an effective RA, as failing to enable the business amounts to a failure to achieve business objectives, which is broadly equivalent to risk.¹⁵

What usually is done may be evidence of what ought to be done, but what ought to be done is fixed by a standard of reasonable prudence, whether it is usually complied with or not.

Third, new and emerging security threats are especially problematic if one eschews a risk-based approach to security. When a new threat is discovered, laws are unlikely to mandate specific security controls to deal with that threat. Moreover, diligence is unlikely to be helpful either, since there probably will not be any de facto industry standard for mitigating that threat. For example, laptop encryption is surely part of the standard of care today, but it was not five years ago.¹⁶ Additionally, compliance is often the lowest common denominator that protects entities outside of an organization (e.g., consumers, government, payment card brands) more than the organization itself. Compliance is rarely an efficient way of allocating resources, since compliance requirements are rarely (if ever) designed with a specific organization's nuances in mind.¹⁷

Fourth, there are often multiple options that can be used to avoid negligence, achieve legal compliance, and enable the business. Organizations have limited resources to invest in information security. Sometimes the resources required just to achieve compliance, diligence, and enablement exceed the resources that are available. Lawmakers and other organizations need a decision-making method for selecting one of those options. The methods of decision theory, including risk analysis, are empirically well-supported. The diligence-based method is not.

Furthermore, while Parker has modified his position by removing an appeal to the legal concept of due diligence, that

does not deny the fact due diligence to avoid negligence itself requires a risk-based approach. Let us define *due diligence* as the prudent person's fulfillment of the duty to use reasonable care; *negligence* is the failure to do so. In this context "reasonable care" means "such care as what a reasonable, prudent, and careful person would use under similar circumstances."¹⁸

In *United States v. Carroll Towing Co.*, Judge Learned Hand determined that a party has a duty to take adequate measures to prevent harm if the cost (B) of taking adequate measures to prevent harm is less than the monetary loss (L) multiplied by the probability (P) of its occurring, expressed by the equation " $B < PL$."¹⁹ Thus, the due diligence needed to determine a party's duty of care requires a cost-benefit analysis that weighs the risk ($P \times L$) against the cost (B) to mitigate that risk.

Moreover, as explained by U.S. Supreme Court Justice Holmes in *Texas & P.R. v. Behymer*, "[w]hat usually is done may be evidence of what ought to be done, but what ought to be done is fixed by a standard of reasonable prudence, whether it is usually complied with or not."²⁰ There is strong evidence that risk-based security is the standard. First, several high-profile U.S. organizations have publicly endorsed security risk management, including the Government Accountability Office,²¹ the Federal Trade Commission,²² the U.S. Marine Corps,²³ the U.S. Air Force,²⁴ and Microsoft.²⁵ Second, a number of professional societies and standards bodies advocate security risk management, including the Information Systems Security Association (ISSA),²⁶ the Information Systems Audit and Control Association (ISACA),²⁷ the American Society for Industrial Security (ASIS),²⁸ the Institute of Internal Auditors

18 Black's Law Dictionary, 6th ed., 1032.

19 Robert Braun and Stan Stahl, "An Emerging Information Security Minimum Standard of Due Care," http://www.citadel-information.com/library/1/emerging-infoSec-standard-of-care-Braun_Stahl_rev_050803.pdf

20 *Texas & P.R. v. Behymer*, 189 U.S. 468, 470, 1903, cited in Braun and Stahl n.d., p. 14.

21 Government Accountability Office, *GAO/AIMD-00-33 Information Security Risk Assessment: Practices of Leading Organization* (Washington, D.C.: GAO, 1999). <http://www.gao.gov/special.pubs/ai00033.pdf>

22 FTC consent decrees regarding information security and privacy practices contain certain recurring themes, including (i) the identification of reasonably foreseeable risks; and (ii) the design of reasonable and appropriate controls and safeguards. See, e.g., FTC File No. 052-3069, *United States of America (for the Federal Trade Commission) v. ChoicePoint Inc.* (United States District Court for the Northern District Court of Georgia, Atlanta Division), at III, p. 15. For a meta-analysis of FTC consent decrees, see Margaret P. Eisenhauer, *The IAPP Information Privacy Case Book: A Global Survey of Privacy and Security Enforcement Actions with Recommendations for Reducing Risks* (York, Maine: IAPP, 2008).

23 United States Marine Corps, Marine Corps Order 3500.27B: Operational Risk Management, *Marines.mil* (2004), <http://www.usmc.mil/news/publications/Documents/MCO%203500.27B%20W%20ERRATUM.pdf>.

24 United States Air Force, Air Force Instruction 90-901: Operational Risk Management, Air Force Link (2000), <http://www.e-publishing.af.mil/>.

25 Microsoft Corp., *The Security Risk Management Guide v1.1* (Redmond, Washington: Microsoft Corporation, 2004).

26 See version 0.3 of the draft Generally Accepted Information Security Principles (GAISP), principle 2.5.

27 As evidenced by ISACA's announcement to launch a new professional certification, Certified in Risk and Information Systems Control (CRISC). See *Information Systems Audit and Control Association* (2010), <http://www.isaca.org/Certification/CRISC-Certified-in-Risk-and-Information-Systems-Control/Pages/default.aspx>.

28 ASIS, *General Security Risk Assessment Guideline* (Alexandria, Virginia: ASIS International, 2003).

12 Act on the Protection of Personal Information, Law No. 57, 2003.

13 See PCI DSS version 1.2.1, requirement 12.1.2.

14 E.g., Minnesota's Plastic Card Security Act (H.F. 1758), Nevada's Security of Personal Information Law (SB 227), and Washington's PCI Law (HB 1149).

15 I owe this point to Aaron Weller.

16 I owe this example to Aaron Weller.

17 I owe this point to Aaron Weller.

(IIA),²⁹ Standards Australia and Standards New Zealand,³⁰ the British Standards Institute,³¹ the U.S. National Institute of Standards and Technology,³² and, most notably, the International Organization for Standardization.³³

In light of this widespread adoption, it seems likely that organizations with an information security function have a duty to implement risk-based security. Therefore, it appears an organization that does not adopt a risk-based approach to security is in danger of failing to exercise due diligence in its management of the information security function. It follows that due diligence requires that security be “risk-based.”

There is no doubt that the job of a CISO would be much easier if one could always sell management on security initiatives by just mechanically pointing to an unambiguous law or survey results about what other companies are doing. The situation is not that simple, however. ISRA can and should be part of the CISO’s tools for managing risk.

Information Security Risk Analysis as Applied Decision Theory

A brief overview of risk management terminology

Given the importance of the word *risk* to Parker’s argument, it is imperative that we work with a clear and rigorous definition. Following Bedford, Cooke, Kumamoto, and Henley, I begin with the word *hazard*. Bedford and Cooke define a *hazard* as an outcome that constitutes a “source of danger.”³⁴ Notice the simplicity of the definition. So defined, *hazard* merely describes a possible source of danger. It makes no claims about the relative probability or improbability of the hazard actually having an impact on anyone or anything.

With the term *hazard* defined, we are now in a position to define risk. Although some disciplines that use risk management do not require that risk involve a hazard, information security specialists (and probably Parker himself) generally agree that an information security risk involves at the least the possibility of a hazard.³⁵ Therefore, I shall define *risk* as a situation in which more than one outcome is possible (and hence, not certain), and at least one outcome involves a hazard. For example, a weather forecast such as, “There is

a thirty percent chance of cloudy skies tomorrow,” is a statement about risk.³⁶ In the context of information security, the relevant hazards would be those that lead to the total or partial loss of confidentiality, integrity, or availability of information. An extreme risk is a risk that concerns a hazard that has both a low probability and a high impact.³⁷

This definition of risk requires uncertainty about the probability of adverse events.³⁸ There is no risk of a hazard if there is a 100% probability of the hazard happening. For example, there is no risk of a Windows computer becoming infected with a worm or virus if the computer is connected to the Internet for sufficient time, not running anti-virus software, and does not have any security patches installed. Rather, the hazard’s occurrence is a certainty.

Statements about risk are probabilistic statements about causal relationships between events. A trigger is an event that partially or completely causes an outcome or hazard event, which in turn partially or completely causes a *consequence event*. Additionally, there may be one or more *preventive events*, which are events intended to prevent the hazard, and one or more *mitigating events*, events designed to reduce the probability or severity of the consequence event.³⁹

For a risk of system compromise due to lack of security patching, a relevant (if greatly oversimplified) influence diagram or risk map is shown in Figure 1.

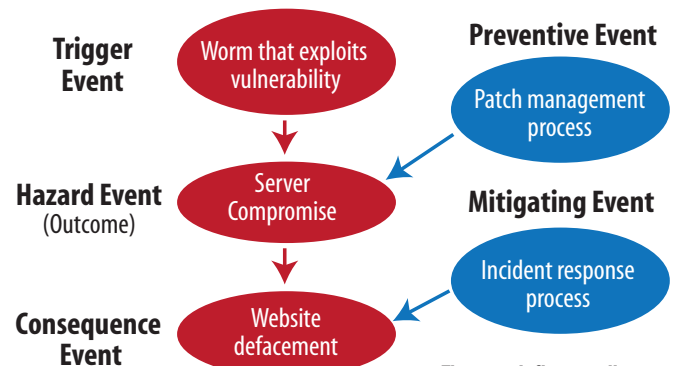


Figure 1: Influence diagram

36 My use of cloudy skies as an example contains an implicit value judgment, namely, that having cloudy skies is an adverse event. I live in the Pacific Northwest, which is notorious for the large number of cloudy days, so I tend to regard cloudy days as an adverse event. People in extremely sunny climates would probably have a different point of view. Nevertheless, I think the cloudy example is a good example for illustrating that risk statements refer to both hazards and their probabilities.

37 James Franklin, Scott A. Sisson, Mark A. Burgman, and Jennifer K. Martin, “Evaluating Extreme Risks in Invasion Ecology: Learning from Banking Compliance,” *Diversity and Distributions* (2008) 14, p. 581. Cf. Nassim Nicholas Taleb’s concept of a black swan event. See his *The Black Swan: The Impact of the Highly Improbable* (Random House, 2007).

38 Cf. Glyn Holton, “Defining Risk,” *Financial Analysts Journal* 60 (November/December 2004), p. 23; Hiromitsu Kumamoto and Ernest J. Henley, *Probabilistic Risk Assessment and Management for Engineers and Scientists* (IEEE, 1996), p. 4. In my opinion, the requirement that risks only involve uncertain outcomes does not make much practical difference. If an outcome is certain (i.e., has a 100% probability), one may plug the event into the same risk management methodology as used for uncertain events; the only decision that has to be made for certain events is whether to accept the outcome as is or to take steps to mitigate it (either prevent it from occurring or decrease its impact to the organization). Nevertheless, I shall follow the trend and define risk in a way that requires uncertainty.

39 Cf. Norman Fenton and Martin Neil, “Visualizing Your Risks: Making Sense of Risks by Letting Them Tell a Story” (N.P.: n.p., n.d.), p. 2. http://www.agenarisk.com/resources/white_papers/Visualising_Risks.pdf

29 The Institute of Internal Auditors, “GAIT for Business and IT Risk (GAIT-R)” (n.p.: Institute of Internal Auditors, March 2008).

30 Standards Australia and Standards New Zealand, AS/NZS 4360:2004 *Risk Management* (Sydney and Wellington: Standards Australia and Standards New Zealand, 2004). Cf. AS 13335-3:2003, *Techniques for the Management of IT Security* (Sydney, Australia: Standards Australia, 2003).

31 BS 7799-3:2006, *Guidelines for Information Security Risk Management* (n.p.: British Standards Institute, 2006).

32 Stoneburner, Gary, Alice Goguen, and Alexis Feringa, *Risk Management Guide for Information Technology Systems: Recommendations of the National Institute of Standards and Technology*, SP 800-30 (Gaithersburg, MD: NIST, 2002).

33 See ISO/IEC 27005, Information Technology – Security techniques – Information security risk management (ISO: Geneva, 2009); cf. *Code of Practice for Information Security Management*, clause 4; ISO/IEC 27001, *Information Security Management Systems – Requirements* (ISO: Geneva, 2005).

34 Tim Bedford and Roger Cooke, *Probabilistic Risk Analysis: Foundations and Methods* (New York: Cambridge University Press, 2001), p. 10.

35 Cf. ISO/IEC Guide 73:2002, p. 2.

Risk analysis (RA) is the identification and estimation of risks. Risk identification is the process whereby one identifies the sources of risk. In an ISRA, risk identification is the identification of hazards. Risk estimation is the process whereby one estimates the probability and (dis)utility of prospective risks. In an ISRA, the probabilities of threats are often measured conditionally – conditional upon the vulnerabilities present in the asset.

In other words, RA answers four questions:

1. What could go wrong (hazard)?
2. How likely is it (probability)?
3. If it happened, how bad could it be (impact)?
4. How much uncertainty is present in the answers to the first three questions (confidence level)?⁴⁰

Risk management is the continuous process of RA, risk treatment, risk acceptance, and risk communication.

The role of statistical syllogisms in risk estimation

As the above remarks make clear, the concept of probability is foundational to risk management. Risk managers use the principles and methods of inductive logic to estimate the probability of hazards. Inductive logic is the study of the strength of the evidential link between the premises and conclusions of arguments. Inductive logic is usually contrasted with deductive logic. If the truth of an argument's premises requires the truth of its conclusion, then the argument is deductively valid. Some invalid arguments, however, contain premises that provide evidence for the conclusion. Such arguments are inductively correct. Inductive strength is a measure of the degree of support the premises provide for the conclusion.

The *Statistical Syllogism* is an inductively correct argument that moves from a proposition about a population to a proposition about a sample: "what is generally, but not universally, true (or false) is also true (or false) for a particular case."⁴¹ Statistical syllogisms have the following form:

- *Z percent of F are G*
- *x is F*
- *Things that are F bear such-and-such relevance to property G*
- *Therefore, x is G*

F is called the *reference class*, the class of individuals or properties that x belongs to or is referred to. G is called the *attribute class*, the class that has the property attributed to x.⁴² In a statistical syllogism, Z can refer to either a single value (i.e., 65%) or a range of values (i.e., 90-95%). We often use

40 Will Ozier, "Risk Analysis and Assessment," *Information Security Management Handbook* (4th ed., Vol. 1, ed. Harold F. Tipton and Micki Krause, Boca Raton: CRC Press, 2004), p. 795.

41 Merrilee H. Salmon, *Introduction to Logic and Critical Thinking* (third ed., Harcourt Brace: New York, 1995), p. 99.

42 I owe these definitions to Salmon 1995, p. 100.

Inductive strength is a measure of the degree of support the premises provide for the conclusion.

fuzzy probabilities for Z to represent a range of values without providing actual numbers for the limits of the range.⁴³ Fuzzy probabilities are expressed with phrases like *most of, usually, probably, often, frequently, almost all, vast majority, high percentage*, and the like.

Inductively correct statistical syllogisms must obey two rules. First, Z must be greater than 50%; the closer Z is to 100%, the stronger the argument.⁴⁴ Second, the statistical syllogism, like all inductive arguments, must obey the Rule of Total Evidence, which is the requirement that the premises of an inductively correct argument must represent all of the available relevant evidence. "Relevant" here means something that can affect the probability of the conclusion (Z). In the context of the statistical syllogism, when selecting the reference class F, we must consider the class that is most relevant to the probability that x is a G. In practical terms, this translates into two requirements. First, the defining properties of F are relevant to x's being G, and, second, F is the most narrowly specified of such classes.⁴⁵

Here is an example. Suppose that all we know about Buster is that he is a dog. We can then use the following statistical syllogism to conclude that Buster likes to go for walks:

- *Virtually all living dogs like to go for walks*
- *Buster is a living dog*
- *Nothing else is known about Buster that is relevant to whether he likes to go for walks*
- *Therefore, it is highly probable that Buster likes to go for walks*

This conclusion is both justified and modest. The conclusion is fully justified, since (1) the only thing we know about Buster is that he is a living dog, and (2) that fact is relevant to whether he likes to go for walks. At the same time, however, our conclusion is not certain. It is possible, even if highly unlikely, that Buster does not like to go for walks. Nevertheless, that possibility alone does not undermine the argument, since the argument's conclusion is merely that it is highly probable that Buster likes to go for walks. The uncertainty is presupposed by the conclusion.

Or consider the following information security example. Let X represent any information security control:

43 Cf. L.A. Zadeh, *Fuzzy Sets, Fuzzy Logic, and Fuzzy Systems: Selected Papers* (River's Edge, NJ: World Scientific, 1996).

44 If Z were to equal 100%, then the generalization would be categorized as a universal generalization, not a statistical generalization. The argument would then become a deductive argument.

45 William Gustafson, *Reasoning from Evidence: Inductive Logic* (Macmillan: New York, 1994), p. 50.

- *The international community of information security practitioners (“community”) are a reliable authority on information security controls*
- *The community believes that X is an effective security control*
- *Therefore, X is an effective security control*

Is this argument inductively correct? Strong? It depends on X. If X is outside of the community’s area of expertise or if equally well-qualified authorities disagree with the community, then the argument is weak. It does not provide evidence for its conclusion.

Parker’s Arguments against Risk-Based Information Security

As I read him, Parker’s critique of risk-based security consists of eight supporting arguments. Those arguments may be divided into three categories: theoretical, empirical, and practical. Let us examine each of his supporting arguments in turn.

Theoretical arguments against risk-based information security

Let’s begin by considering Parker’s arguments against the possibility of actually doing risk-based security in the real world.

First supporting argument: Uncertainties involved in ISRA

Here is Parker:

The frequencies and impacts of future incidents are under the control of unknown and often irrational enemies with unknown skills, knowledge, resources, authority, motives, and objectives from unknown locations at unknown future times attacking known but untreated vulnerabilities and vulnerabilities that are known to the attackers but unknown to the defenders (a constant problem in our technologically complex environments).⁴⁶

This objection to risk-based security is multiply flawed.

First, many of the variables listed by Parker are simply not relevant to assessing the probability of an attack. One does not need to know the identity of an attacker, much less his “skills, knowledge, resources, authority, motives and objectives” (SKRAMO), in order to estimate the probability of an attack. There is no doubt that we often lack knowledge about the SKRAMO of our attackers, but that does not mean we cannot calculate the probability of an attack.

Suppose we have historical data about the frequency of occurrence of a particular type of security threat spanning mul-

iple years, across multiple organizations of varying size, geographical location, and so forth. For example, let the threat be workplace violence, a threat that involves “irrational, unknown humans.” Based on the historical data just mentioned, one can infer a statistical generalization about the frequency of the workplace violence threat overall. One can also make more specific generalizations about various subsets of the overall workplace violence threat. For example, one can make statistical generalizations about the rate of incidents of workplace violence in organizations that recently went through a round of layoffs, in individuals who were subject to one or more negative personnel actions, and so forth.

My claim is that, as a security professional, I can use that statistical data in a quantitative RA of the workplace violence threat for *my company*. Yes, there are differences between my company and the other organizations for which we have statistical data about incidents of workplace violence. But the mere existence of such differences does not automatically invalidate inductively correct or statistically valid inferences about the level of risk for *my company*, given what we know about the rate of occurrence of workplace violence in other companies. In order for such inferences to be inductively incorrect or statistically invalid, one would have to show that the differences between my company and other companies are *probabilistically relevant*. Suppose Microsoft announces tomorrow the existence of a previously unknown security vulnerability in one of its software products. There will not be any historical data tomorrow regarding the rate of occurrence of attempted exploits of that vulnerability. There will not be any historical data tomorrow regarding our enemies and whether they are planning on exploiting that vulnerability. Despite that lack of data, however, we can still make accurate calculations of the level of risk, based upon what we know about past security vulnerabilities and past enemy attacks.

Second, even in those situations where historical data about the statistical frequency of a specific threat is unavailable, historical data is often available for some larger class of events for which the specific threat is a member. Every time a new vulnerability is announced in a given piece of software, by definition there will be no historical data about that vulnerability. Yet we have a large amount of statistical data about information security vulnerabilities in general. We also have a wealth of historical information about vulnerabilities affecting specific types of software. For example, the next time a new vulnerability in the Apache web server software is announced, we will not have any statistical data regarding the frequency of exploits of *that* specific vulnerability, but we will have data about the frequency of past Apache vulnerabilities for which exploit code is publicly available. That latter data is relevant to determining the probability that exploit code will be made publicly available for the new vulnerability.

A similar point applies to the worry about threats stemming from the acts of individuals (as opposed to “acts of nature”). While we may not have any historical data about the prob-

⁴⁶ Parker 2006, p. 7.

ability of a specific enemy committing attack Y, we do have statistical data about attacks in general, specific types of attacks, and attacks against specific organizations. To be sure, the more specific the reference class, the more confidence we will have in our probability values. Just because our reference class is not identical to the event in question, however, it does not follow that we cannot have a reasonable or even high degree of confidence in our probability values. The fact that we cannot know something with *certainty* (i.e., probability = 100%) does not prevent us from knowing it with a high degree of *probability* (i.e., probability > 50%).

We are now in a position to identify the fatal flaw in Parker's objection from uncertainty. The objection reduces to a requirement that statistical syllogisms always appeal to a reference class that is as specific as the event itself, *even if evidence regarding the frequency of the event is not available*. Inductive logic, however, imposes no such requirement. As we saw earlier, an inductively correct statistical syllogism need only embody the total (available) evidence; it does not need to include unavailable evidence. If evidence regarding the frequency of the event is not available, it is perfectly acceptable (and logically correct) to select as the reference class a relevant, broader class of events for which the frequency is known. Parker's requirement that the reference class be identical to the event itself is an arbitrary, unjustified requirement. Indeed, if we were to require that the reference class be identical to the event itself, that would pretty much eliminate the statistical syllogism as a correct argument pattern. That is absurd, however, and so Parker's requirement is not genuine.

Parker also introduces the following related objection:

*In addition, when enemies fail in attacking one possible vulnerability, they often attempt attacks on other vulnerabilities to accomplish their goals. Therefore, risks may be related in unknown complex ways so that reducing one risk may increase or decrease other risks. This alone precludes the effective use of risk assessment methods.*⁴⁷

Parker is certainly correct that if an initial attempt to exploit a vulnerability fails, many attackers will try other attacks in order to accomplish their goals. It does not follow, however, that this fact "alone precludes the effective use of risk assessment methods" (emphasis mine). This is an incredibly strong claim that requires a supporting argument from Parker, but such an argument is not provided in his article.

In sum, far from disqualifying the use of RA, our uncertainty about attackers provides a strong reason for using probabilistic, risk-based methods. As Doug Hubbard writes, "We use probabilistic methods because we lack perfect data, not in spite of lacking it. If we had perfect data, probabilities would not be required."⁴⁸ Furthermore, "It is a fallacy that when a variable is highly uncertain, we need a lot of data to reduce the uncertainty. The fact is that when there is a lot of

uncertainty, less data is needed to yield a large reduction in uncertainty."⁴⁹

Second supporting argument: Complex, unknowable relationships between risks and security efforts

Parker's next argument claims that that the relationships between risks and security efforts are complex and often not completely knowable:

*Also, there is no one-to-one relationship between one risk and one security effort since many security efforts may affect one risk and one security effort may affect many risks, which is occurring now with powerful security packages. Thus the impacts are related in unknown ways as well.*⁵⁰

Parker is certainly correct that risks and security controls are often related in complex ways. On the other hand, just because RA is difficult does not mean it should not be done, or that there is no value in performing it.

Many business decisions are based on many-to-many relationships where cause and effect are not necessarily related in a simple way. This does not preclude organizations from making decisions where no one-to-one relationship exists between risks and actions.⁵¹

Furthermore, Bayesian Networks (BNs) are perfectly suited to modeling the many-to-many, probabilistic relationships between uncertain propositions as described by Parker. The advantages of BNs are well-summarized by Norman Fenton:

*BNs enable reasoning under uncertainty and combine the advantages of an intuitive visual representation with a sound mathematical basis in Bayesian probability. With BNs, it is possible to articulate expert beliefs about the dependencies between different variables and to propagate consistently the impact of evidence on the probabilities of uncertain outcomes. BNs allow an injection of scientific rigour when the probability distributions associated with individual nodes are simply "expert opinions".*⁵²

BNs (and their applicability to RA) have been carefully described, but unfortunately Parker's writings evince little awareness of this literature.⁵³

49 Hubbard 2009, Kindle location 3950-1.

50 Parker 2006, p. 7.

51 I owe this point to Aaron Weller.

52 Norman Fenton, "Using Bayesian Networks to model Expected and Unexpected Operational Losses" 30 November 2004, p. 3. http://www.agenarisk.com/resources/white_papers/Operational_Losses_Risk_Analysis_Journal.pdf

53 See Richard E. Neapolitan, *Learning Bayesian Networks* (Upper Saddle River, NJ: Prentice Hall, 2004); Cavusoglu, H., S. Raghunathan, W.T. Yue. 2008. "Decision-theoretic and Game-theoretic Approaches to IT Security Investment" *Journal of Management Information Systems* 25(2) 281-304; Fenz, S. and Neubauer, T. 2009. "How to Determine Threat Probabilities using Ontologies and Bayesian Networks" In *Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research: Cyber Security and information intelligence Challenges and Strategies* (Oak Ridge, Tennessee, April 13 - 15, 2009). F. Sheldon, G. Peterson, A. Krings, R. Abercrombie, and A. Mili, Eds. CSIRW '09. ACM, New York, NY, 1-3. DOI= <http://doi.acm.org/10.1145/1558607.1558686>; Neil, M., N. Fenton, et al., "Using Bayesian Networks and Simulation for Data Fusion and Risk Analysis," *NATO Science for Peace and Security Series: Information and Communication Security*. Skanata and D. M. Byrd, Amsterdam: IOS Press, 2007.

47 Parker 2006, p. 7.

48 Douglas W. Hubbard, *The Failure of Risk Management* (New York: Wiley, 2009), kindle reference: 2296.

Third supporting argument: Impossibility of measuring risk reduction

Parker's next argument against risk-based security is that it is impossible to measure the reduction of risk.

Investing in the reduction of security risk, on the other hand, results in some unknown possible reduction of loss that is not measurable since it hasn't happened yet and is unknown. ...

It is not possible to know the effectiveness of safeguards that have no detection capabilities associated with them.... Good security is when nothing bad happens, but with nothing bad happening, there is nothing to measure. Therefore, the goodness of security against rare incidents can't be measured since you don't know what you may have stopped.⁵⁴

Regarding Parker's claim that "[g]ood security is when nothing bad happens," I interpret "nothing bad happens" to mean perfect risk prevention, i.e., the non-occurrence of any hazards. I think that standard sets the bar way too high for "good security." Ever-increasing security investments obey the law of diminishing returns, such that any attempt to ensure "nothing bad happens" would be an ill-fated waste of money. An organization with "good security" may have security incidents; a lucky organization with "bad" security may have no incidents.

Furthermore, risk reduction is measurable. Recall that the definition of risk is the situation in which more than one outcome is possible (and hence not certain), and at least one outcome involves a hazard. *Reduction of risk*, therefore, means a reduction of a hazard's probability, impact, or both. Since both the probability and impact of hazards can be measured, it follows that the reduction of either quantity (and hence, risk reduction) can also be measured.⁵⁵

Now Parker charges that risk reduction cannot be measured for low probability risks because safeguards or other risk mitigations may have prevented the occurrence of a hazard entirely. Even if this were true, it wouldn't follow that diligence, compliance, and enablement would be the only way to justify said safeguards. At worst, this would simply represent a textbook example of a *decision under ignorance*. The information security practitioner could then use any of the standard methods for making a decision under ignorance, such as the *Maximin Rule*, the *Minimax Regret Rule*, and the *Optimism-Pessimism Rule*.⁵⁶

But in fact I think Parker has overestimated the number of cases where the probability is unknowable. When Parker writes, "the goodness of security against rare incidents can't be measured since you don't know what you may have stopped," he's presupposing a *frequency* interpretation of probability, which defines probability as the relative frequency with which an outcome appears in a long series of similar events. By defini-

⁵⁴ Parker 2006, p. 7.

⁵⁵ Cf. Douglas W. Hubbard, *How to Measure Anything: Finding the Value of Intangibles in Business* (Hoboken: John Wiley & Sons, 2007), pp. 45-47.

⁵⁶ Resnick 1987, pp. 21-44.

If evidence regarding the frequency of the event is not available, it is perfectly acceptable (and logically correct) to select as the reference class a relevant, broader class of events for which the frequency is known.

tion, it is difficult at best to measure the *relative frequency* of a rare hazard. And Parker is largely correct: if an organization implements a safeguard that has no detection capability associated with it, and if there are no other relevant detection capabilities, it will be impossible to *precisely* measure the relative frequency of averted incidents.

But that is where my agreement with Parker ends.⁵⁷ Douglas Hubbard put it best when he wrote, "*the lack of having an exact number is not the same as knowing nothing*."⁵⁸ There are other (admittedly less precise) methods for measuring a hazard's probability besides the relative frequency method. Such methods include eliminative induction,⁵⁹ arguments from analogy, Fault Tree Analysis (FTA), the calibration of subject matter experts,⁶⁰ BNs, and extreme value theory.⁶¹ These methods are not even mentioned, much less refuted, by Parker.

Consider the calibration of subject matter experts. Fully informed, calibrated experts can provide reliable estimates of their subjective estimate of a hazard's probability within a given confidence level.⁶² By "calibrated," I mean trained to remove one's personal estimating bias (towards either overconfidence or underconfidence).⁶³ If an expert is calibrated, then, "over the long run, for all propositions assigned a given probability, the proportion that is true equals the probability assigned."⁶⁴

Finally, BNs are uniquely effective in combining diverse sources of data – such as calibrated expert inputs, FTA, Statistical Syllogisms – to quantify risk, in a mathematically rigorous way.⁶⁵ Several international banks have used BNs to

⁵⁷ While it would be a hopeless task in the limited space of this article to provide a critical analysis of Parker's presuppositions, I think it is salutary to point out that his assumption of a frequency interpretation of probability is question-begging and highly controversial. For a discussion of the different interpretations of probability, see Donald Gillies, *Philosophical Theories of Probability* (New York: Routledge, 2000).

⁵⁸ Hubbard 2007, p. 64. Italics in original.

⁵⁹ L.J. Cohen, *The Probable and the Provable* (Clarendon: Oxford, 1977).

⁶⁰ Roger M. Cooke, *Experts in Uncertainty: Opinion and Subjective Probability in Science* (New York: Oxford University Press, 1991).

⁶¹ Emil Julius Gumbel, *Statistics of Extremes* (New York: Columbia University Press, 1966).

⁶² Sarah Lichtenstein, Baruch Fischhoff, and Lawrence D. Phillips, "Calibration of Probabilities: The State of the Art to 1980," *Judgment under Uncertainty: Heuristics and Biases* (ed. Daniel Kahneman, Paul Slovic, Amos Tversky, New York: Cambridge University Press, 1982), 306-334.

⁶³ I owe this definition to Hubbard 2007, p. 54.

⁶⁴ Lichtenstein et al 1982, p. 307.

⁶⁵ Cf. Norman Fenton and Martin Neil, "Combining Evidence in Risk Analysis Using Bayesian Networks" 23 July 2004, p. 1.

satisfy the Basel II Banking Accord's requirement that banks quantify operational risk, including the risk of fraud.⁶⁶

Empirical Arguments against Risk-Based Information Security

The second category of Parker's arguments against risk-based information security attempt to show that there is no good evidence for risk-based information security while there is good evidence against it. Let's consider each of his arguments in turn.

First supporting argument: Security risk assessments fail

Parker's first empirical argument is introduced by the opening sentence of his article: "IT trade publications are increasingly reporting the failings of risk management and risk assessments." According to Parker, risk-based security is a failed method of making security decisions because the data needed for risk management does not exist. Based upon statements made elsewhere in the article, we can put together the pieces of this first argument as follows:

- According to some IT trade journals, information security is under-funded and under-staffed
- According to various articles in IT trade publications, risk management is a failed method of making security decisions
- Therefore, risk management and risk assessment are failed strategies

This is an argument from authority. Since Parker is a security practitioner, not a philosopher, I will gloss over the authority-related problems with this argument.⁶⁷ Instead, I want to focus on the assumptions of his argument that are more central to risk-based security.

First, Parker has not supported or even identified a specific cause-and-effect relationship between the use of specific risk management methods and information security programs that are under-funded and under-staffed. Indeed, it's not even clear that Parker (or his sources) know which methodologies were being used at the organizations with such programs.⁶⁸

66 Norman Fenton and Martin Neil, "Managing Risk in the Modern World: Applications of Bayesian Networks", (London: London Mathematical Society, November 2007), p. 14. Cf. Shilpa Ramamurthy, Harpreet Arora, and Anirbid Ghosh, "Operational Risk and Probabilistic Networks: An Application to Corporate Actions Processing" (November 2005), <http://www.infosys.com/offerings/industries/banking-capital-markets/white-papers/documents/operational-risk-probabilistic-networks.pdf>

67 The problem is not that this argument is an appeal to authority, and all appeals to authority are incorrect. Appeals to authority are a special version of the Statistical Syllogism. As such, they can be inductively correct. Rather, the problem is that this is a weak argument from authority, since (1) the authorities cited don't make the claims attributed to them, and (2) there many equally well qualified authorities who believe that security needs to be risk-based. See my discussion of due diligence for a list of organizations that promote risk-based security.

68 I owe this point to Douglas W. Hubbard.

Second, it's far from obvious the failed risk management methods were representative examples of widely used risk management methodologies, much less evidence-based risk management methods. This is a case where the details matter; not all methodologies produce conclusions about risk that are inductively correct. Dismissing the entire discipline of information risk management on the basis of flawed methodologies is akin to arguing that we should not pursue chemistry because the alchemists have been proven to be frauds.⁶⁹

Second supporting argument: Specific attempts to use risk analysis methods have failed

Next, Parker turns his sights on specific RA methods, such as Annual Loss Expectancy (ALE).

*In my experience many efforts to use such methods (e.g. NIST Annual Loss Expectancy) such as those recommended in three recent articles in The ISSA Journal have faded away when the high cost, easily disputable results, and changing environment are realized. Many small automated risk assessment businesses quickly fail.*⁷⁰

Parker may well be correct that attempts to use the Annual Loss Expectancy (ALE) approach for ISRA fail. It does not follow, however, that there is no correct method for analyzing information security risks. Parker's objection may not unfairly be summarized as follows:

- There are many ISRA methodologies
- One of the methodologies, ALE, is a failure
- Therefore, ISRA is a failure

Before discussing this argument, I need to make one minor correction. While NIST endorsed the ALE approach at one time, the ALE approach has been withdrawn by the U.S. government; it is not the current (or even recent) NIST recommended method for ISRA.⁷¹ Turning to Parker's argument, ALE is not representative of ISRA methodologies – today or in 2006. Parker needs to provide reasons for rejecting the most common ISRA methods (e.g., FAIR,⁷² OCTAVE,⁷³ FRAP,⁷⁴ ISO/IEC 27005,⁷⁵ AS/NZS 4360,⁷⁶ SABSA⁷⁷, IRAM⁷⁸, CRAMM).

69 I owe this point to Douglas W. Hubbard.

70 Parker 2006, p. 7.

71 Cf. See Gary Stoneburner, Alice Goguen, and Alexis Feringa, *NIST Special Publication 800-30: Risk Management Guide for Information Technology Systems* (Gaithersburg, Maryland: NIST, July 2002).

72 Risk Management Insight, *FAIR™ Risk Assessment Guide, Risk Management Insight* (n.d.), http://www.riskmanagementinsight.com/media/docs/FAIR_brag.pdf.

73 Christopher Alberts and Audrey Dorofee, *Managing Information Security Risks: The OCTAVE™ Approach* (Boston: Addison Wesley, 2003).

74 Thoms R. Peltier, *Information Security Risk Analysis* (2nd ed., Boca Raton, Florida: Auerbach, 2005).

75 ISO/IEC 2009.

76 Standards Australia and Standards New Zealand 2004.

77 John Sherwood, Andrew Clark, and David Lynas, *Enterprise Security Architecture: A Business-Driven Approach* (San Francisco: CMP Books, 2005).

78 "IRAM -- Information Risk Assessment Methodology," *Information Security Foundation* (n.d.), <https://www.securityforum.org/services/publictools/publiciram>.

Third supporting argument: Lack of evidence supporting the validity of risk-based security

Parker's third supporting argument may be categorized as a "lack of evidence" argument. According to Parker, there is no study that demonstrates that security risk management actually works. In his words, "No study has ever been published to demonstrate the validity of information security risk assessment, measurement, and control based on real experience."⁷⁹

I agree with Parker's implicit assumption that we should require evidence that information security RA works.⁸⁰ And I suspect that Parker is probably correct that there has been no study published that demonstrates the validity of ISRA *specifically*. By itself, however, that fact hardly calls into question the validity of the ISRA discipline. There also has been no empirical study published that demonstrates the *invalidity* of ISRA.

Furthermore, Parker ignores the possibility of other types of support for ISRA. Contrary to Parker, I believe that support for ISRA may be found in the support for generic RA – such as the axioms of probability and utility, the evidence regarding the reliability of calibrated experts – combined with the lack of a good argument against applying RA to information security.

Practical Arguments against Risk-Based Information Security

Turning to the third and final category of Parker's arguments, Parker argues that risk-based security should be rejected on practical grounds.

First supporting argument: Ease of management risk acceptance

Parker's first practical argument highlights the ability of decision makers to choose to accept security risk, rather than to make the investment necessary to mitigate risk. He writes:

*Management's business is risk-taking, and when a security risk is presented to them, they are able to respond that they take risks every day. And they simply accept the risk presented to them and refuse to support the security that is claimed to reduce it, especially when they see the negative impact and inconvenience that security has on the organization and their business goals. When the reputed risk doesn't materialize into a loss event, or some other larger unpredicted risk materializes instead, there is a justified loss of trust and belief in the value of security risk assessment and those presenting it.*⁸¹

Elsewhere, Parker asserts that management is especially willing to accept risk when confronted with a risk-based case for safeguards against rare incidents.⁸² In contrast with Parker's

theoretical objections that risk cannot be *measured*, this objection targets the unreliability of *risk management* at guiding management to the "correct" decision. In other words, Parker's objection may be paraphrased as follows: "Even if the risk manager gets the math right and communicates it effectively, management can still make the wrong decision when presented with a risk-based argument for a voluntary security safeguard."

As a reason for *eliminating* RA from information security management, this argument is multiply flawed. First, with all due respect to Parker, he seems to have forgotten that management has the autonomy to make business decisions about security, however "wrong" those decisions may appear. Second, Parker ignores instances in which management has voluntarily approved security safeguards, solely on the basis of a risk-based approach. Third, the risk-based approach is hardly unique in giving management discretion regarding whether to implement a safeguard; that same discretion often exists in compliance settings (though admittedly much less frequently than in risk management discussions). Has he never had to deal with an ambiguous compliance requirement, one that reasonable people can disagree about whether an organization complies with? Fourth, as I argued earlier, diligence-based and risk-based security are complementary. In the case of new, emerging threats, RA may be able to justify a security safeguard when diligence-based methods are unable to do so. Fifth, contrary to Parker, when a reputed risk does not materialize into a loss event, it does not follow that "there is a justified loss of trust and belief in the value of security risk assessment and those presenting it."

Perhaps Parker has fallen prey to what Gigerenzer, et al have called "[t]he ambiguity of a single-event probability and the resulting possibility of miscommunication about risks."⁸³ A single-event or *single-case probability* is a probability value that does not specify a reference class. Thus, for example, when a weather forecaster says "There is a 30% chance of rain tomorrow," one may ask, "30% of what?" Gigerenzer and his colleagues found that a majority of people interpret that forecast to mean that it will rain in 30% of the *area* or 30% of the *time*. What forecasters actually meant, however, was that, "when the weather conditions are like today, at least a minimum amount of rain (such as .2mm or .01 in.) will fall the next day in 3 out of 10 cases."⁸⁴

Given the misunderstanding, it is easy to understand why people may lose trust in weather forecasting and the people presenting it. Along the same lines, then, when a risk manager describes a risk, saying there is, say, a 65-80% chance of a hazard event occurring and the hazard does not happen, it does not follow that the original probability estimate (or the

Please continue on page 49

79 Parker 2006, p. 7.

80 Hubbard 2009, Kindle location 449-62.

81 Parker 2006, p. 9.

82 Parker 2006, p. 6.

83 Gerd Gigerenzer, Ralph Hertwig, Eva van den Broek, Barbara Fasolo, and Konstantinos V. Katsikopoulos, "A 30% Chance of Rain Tomorrow: How Does the Public Understand Probabilistic Weather Forecasts?" *Risk Analysis* 25 (2005), p. 629.

84 Gigerenzer et al, p. 624.

Is Risk-Based Security a Failed Concept? Continued from page 31

method used to produce it) is incorrect. The occurrence or non-occurrence of a single event does not invalidate a probability estimate.

Second supporting argument: Abuses and misuses of ISRA

Parker's second practical argument appeals to abuses and misuses of RA by security professionals. For example, Parker tells the story of a CISO who admitted to performing security RA backwards.⁸⁵ This single anecdote is not of obvious relevance to risk-based security, however. There have been individuals who have misused RA, just as there have been individuals who have misused the software development life cycle (SDLC) or falsified financial records. The fact that such behavior has occurred does not invalidate RA, the SDLC, or generally accepted principles of financial accounting. Indeed, there have also been individuals who have misused the due care controls approach, but critics of RA do not suggest that we toss out the due care controls approach because of the incorrect actions of a few individuals. RA should not be held to a different standard.

Conclusion

I agree with Parker that compliance, diligence, and enablement are worthy goals for any information security program; what I disagree with is his claim that those goals should replace risk-based security. None of his arguments seem to be successful in showing that risk-based security should be replaced by compliance, diligence, and enablement. Of course, this in no way rules out the fact that there are problems with the current state of ISRA. For example, it may be the case that game theory is better suited than traditional decision theory

⁸⁵ Parker 2006, p. 8.

to model the risk associated with a malicious threat agent.⁸⁶ Additionally, I suspect many information security risk managers are unaware of where to obtain empirical data regarding the frequency of hazards, or how to calibrate expert opinion when such data is unavailable. I am also skeptical of traditional risk matrices for management reporting of risks measured by qualitative methodologies.⁸⁷ These challenges, however, do not invalidate risk-based security. As I have argued, ISRA is complementary with the goals of compliance, diligence, and enablement. Furthermore, RA can help to put compliance in its proper place as a by-product of an effective security program, not as a primary goal in itself.⁸⁸

About the Author

Jeff Lowder, CISSP, is an independent consultant who helps organizations balance information security with business agility using evidence-based governance, risk, and compliance (GRC) methods. His previous roles include infosec leadership positions at the U.S. Air Force Academy, United Online, and Disney. He may be reached at jefflowder@gmail.com or visit www.jefflowder.com.



⁸⁶ See Huseyin Cavusoglu, Raghunathan, Srinivasan and Yue, Wei T., "Decision-Theoretic and Game-Theoretic Approaches to IT Security Investment," *Journal of Management Information Systems* 25 (2008): 281-304; J. Boon and B. M. Gorman, "Applying Game Theory to Balance Risk and Cost for Security Inspection Systems," *Managing Critical Infrastructure Risks: Decision Tools and Applications for Port Security* (ed. Igor Linkov, Richard J. Wenning, and Gregory A. Kiker, Dordrecht: Springer, 2007), 309-326.

⁸⁷ See Louis Anthony Cox, Jr., *Risk Analysis of Complex and Uncertain Systems* (New York: Springer, 2009), pp. 101-24.

⁸⁸ I am grateful to Aaron Weller for lengthy discussions about this article. I am also grateful to Thom Barrie, Douglas W. Hubbard, John Christiansen, Tony Cox, Jim Lippard, Darek Milewski, Dennis Opacki, Donn Parker, Jared Pfost, Joel Scambray, and Hal Tipton for helpful comments on an earlier draft of this article.