Baskerville, R. (1991). "Risk analysis: an interpretive feasibility tool in justifying information systems security." *European Journal of Information Systems* **1**(2): 121-130.

Risk Analysis: An Interpretive Feasibility Tool in Justifying Information Systems Security

Richard Baskerville
School of Management
State University of New York
Binghamton, New York 13902-6000

Abstract

Risk analysis is the predominant technique used by information security professionals to establish the feasibility of information systems controls. Yet it fails an essential test of scientific method, it lacks statistical rigour and is subject to social misuse. Adoption of alternatives from other disciplines, however, proves even more implausible. Indeed, even improved rigour in risk analysis may limit its usefulness. Perhaps risk analysis is misconceived: its ostensible value as a predictive technique is less relevant than its value as an effective communications link between the security and management professionals who must make decisions concerning capital investments in information systems security.

Table of Contents

Introduction

Risk analysis is the predominant technique used by information security professionals to establish the feasibility of information systems controls. Yet the technique is a disturbing mixture of quantitative techniques applied to interpretive data. The purpose of this paper is to explain why risk analysis is so widely accepted, despite its conflicting foundations. Unfortunately, its importance as a communication tool may have been subjugated to its importance as a predictive technique.

The explanation will require a brief overview of risk analysis, a critical commentary on risk analysis and a critical consideration of various alternatives. Many critiques on risk analysis focus on its contribution to `scientific practice' in the professional security community. The work below seeks to establish its more essential value as a means of accurately communicating the security analysis to management.

Virtually all authors of information system security methodologies adopt risk analysis as a technique for justifying the cost of controls acquisition to management. Robert Courtney [1977] and Jerry FitzGerald [1978] pioneered the technique. Many published methodologies adopt or adapt their approach (*cf.* Royal Fisher [1984], Donn Parker [1981] or Badenhorst and Eloff [1990]), and large organisations practice variations (see Saltmarsh and Browne [1983] for an excellent review). The U.S. government accepts risk analysis as one of its Federal Information Processing Standards [NBS, 1979].

Later authors, particularly from the auditing community (*e.g.*, Saltmarsh and Brown or Gallegos *et al* [1987]), sometimes differentiate between risk analysis and risk assessment. To these authorities, *risk analysis* is the process of identifying risks and their characteristics and *risk assessment* is the determination of the exposure to risk. When this is distilled to monetary units, this becomes *quantitative risk assessment.* However, this paper will agree with the early risk analysis work (*e.g.*, Courtney and the FIPS standard) that does not observe these delineations and assumes that the concept of *risk analysis* will include quantitative risk assessment.

Risk analysis is used in the establishment of system controls, *i.e.*, the selection of controls for computer-based information systems. Thus, it is a component of *meta-control*. Early management theorists found that control activities operate in a thousand different ways, and demand `a good deal of art'. Fayol defined control thus:

> `Control consists of verifying whether everything occurs in conformity with the plan adopted, the instructions issued, and the principles established, and then taking the appropriate corrective actions.' [Fayol, 1987]

Control is sometimes confused with planning or information (for verification), which are *prerequisite components* of control. Risk analysis is most commonly viewed as a tool for planning security by statistically selecting controls for implementation.

Risk Analysis

There are two dimensions to a critical description of risk analysis. First, there is the traditional role of risk analysis as a technique or tool in the design of information systems (IS). Second, there are the details of the technique itself.

The Role of Risk Analysis

Management will often view the acquisition of information systems as one of several competing alternatives for capital investment. For this reason, most system proposals will include an economic feasibility study or cost-benefit analysis. One role of these studies or analyses is that of a proof of the expectation that the system will produce benefits that equal or exceed the system's cost. A more important role, however, is to aid management in comparing the IS investments to their opportunity costs. Since the IS proposal's opportunity costs are alternative capital investments, the feasibility study must prove the expectation that the system's return on its investment costs will equal or exceed the competing capital opportunities.

The measurement of IS benefits can be problematic. The inability to specify benefits in monetary terms created a class of benefits called `intangible benefits.' Such intangibles may fare poorly in the capital investment arena, so quantification in any form is desirable [Emery, 1971]. Information economics [Kleijnen, 1980] and its predecessor, software engineering economics [Boehm, 1981], advocate approaches such as statistical decision theory as quantification solutions.

Like information economics, risk analysis serves as a benefits quantification technique in the context of information systems security. It too, is a creature of statistical decision theory. It provides a means by which intangible system attributes, such as probable compromise, are translated into monetary terms. The desirability of

such attributes can then be directly compared to their costs and their opportunity costs as part of the capital investment decision.

Risk Analysis Overview

This brief overview approximates the version of risk analysis proposed by Courtney and adopted by the FIPS standard. It is necessarily brief and will lack many of the attractive nuances of the technique. The interested reader should obtain the references for more lengthy discussions.

Risk analysis defines two major elements of risk (*R*): *P*, the probability of an exposure's occurring a given number of times per year, and *C*, the cost or loss attributed to such an exposure. Risk is calculated as

$$R = P \times C$$

The probability *P* is determined with the aid of a Probability Range Table (see Table I), which provides various subjective frequency times, and an equivalent annualised `loss multiplier'. For example, a subjective estimate of frequency of loss due to the risk exposure of an information system element might be `once in three years'. This yields an annualised loss multiplier of .3333 (one-third of a loss per year).

The cost/loss value *C* is estimated with the aid of a Cost/Loss Range Table (see Table II), which provides the subjective cost of a single loss in decimal exponential steps; i.e. £10, £100, £1000, £10000 and so forth. For example, an estimated loss of £80,000 would yield a cost/loss value of 100,000 (the next highest exponential increment of 10).

The examples discussed above would result in a total quantitative risk of £33,333. To summarise this example, the subjective estimate of frequency of loss as `once in three years' yields an annualised loss multiplier of .3333, and an estimated loss of £80,000 would yield a cost/loss value of 100,000. The risk resulting from these examples of *P* and *C* above would be calculated to be £33,333. This figure could then be used in assigning implementation priorities for new controls and cost justifying any control changes in the system.

The equation below summarises the risk formula for use with *P* (probability factor) and *C* (cost/loss value) in the tables:

$$R = \frac{10^{(P-3)}}{3} \times 10^{C}$$

A Critical Commentary on Risk Analysis

Risk analysis is a very important information security tool. It adds requisite order and classification to the process of determining which subset of the set of all possible controls should be selected for implementation. It is indispensable in gaining management support for the construction of safe and reliable IS. It is, perhaps, the primary component for meta-control. There are few alternatives. Yet, risk analysis fares poorly under critical thought. Considered as a scientific approach, as a statistical approach, or even as a practical approach, risk analysis emerges as a shallow and inappropriate predictive technique.

The Scientific Philosophy of Risk Analysis.

Risk analysis seeks to create professional knowledge about an existing or proposed information system. Such professional knowledge shares epistemological foundations with its basic science [Schön, 1983]. Risk analysis, however, assumes a strongly pluralist philosophy. That is, it is both positivistic *and* interpretivistic in its epistemology. Its basic data values (risk probabilities and loss estimates) are highly interpretive -- often gleaned whole by the professional from an unstructured study of the complex multivariate organisational landscape. These values are then manipulated with very positivistic formal and logical mathematical operations. (Clements developed an English equivalent to such statistical metrics with fuzzy set theory, translating probabilities to such modifiers as *very* and *moreorless* [Clements, 1977].)

As a scientific method, however, risk analysis is severely inadequate. Primarily, it lacks the ability to establish feedback regarding the effectiveness of the controls design, specification and implementation. Because of the effects of risk and uncertainty (luck), it lacks any `proof of performance' component. Typically, computer security is a low priority of management until something awful occurs [O'Mara, 1985]. Then, appropriately a highly visible effort is made to improve the controls: risk analysis, security consultants, security hardware, security software, etc. Slowly, however, the problem moves from central focus to the periphery and frequently, further to oblivion. What remains, like the echo from the abrupt conclusion of a military band, are some controls, some security hardware and software, an operationally more complex system and some uninteresting statistical analyses. In most cases the original system functioned perfectly for years without these `improvements.' Unless the disaster reoccurs, management never really receives effective feedback to indicate that the effort has actually provided real security. Indeed, *Providence* may be merely supplying `good luck.'

The effectiveness of techniques for designing computer security is a bit more difficult to establish than that of techniques for designing computer systems. The result may be that ineffective control designs are encumbering existing information systems with unnecessary (and perhaps ruinous) modifications. However, the cloudy feedback situation prevents effective measurement of the success or *necessity* of the controls.

Risk Analysis As Applied Statistical Decision Theory

Statistical decision theory has developed a sophisticated body of thought that regards decisions about future (predicted) events. For example, their language carefully delineates the element of probability discussed above from risk (the random variation of the estimate from the actual -- Providence) and from uncertainty (the error-induced variation of the estimate from the actual -- bad guesses) [DeGarmo, *et al*, 1979]. The ensueing Bayesian statistics can be exquisite, and decision-taking in the presence of uncertainty can become heuristic (*e.g.*, maximin, maximax and Laplace rules).

While a detailed review of Bayesian decision theory is beyond the scope of this paper, it might be useful to briefly mention its character in the context of risk analysis. Bayes theorem addresses the statistical `revision' of *prior* probabilities yielding *posterior* probabilities. It arises in decision theory when imperfect information is added to a decision situation. In risk analysis, improvements in the quantitative validity -- *e.g.*, considering probability of riots $P(S_1)$ and the probability of no riots $P(S_2)$ -- may be related to additional data (*e.g.*, the political stability of a regime). This requires the intersection of probabilities (*viz.*, the probability of riot and the probability of a politically stable regime). The model regards the possibility that any information (*e.g.*, `the regime is politically stable') could be incorrect (hence, additional data about the state of nature is probabilistic). Bayes' rule for calculating posterior probabilities is expressed as follows:

$$P(S_1|I) = \frac{P(I \cap S_1)}{P(I)} = \frac{P(I|S_1)P(S_1)}{P(I|S_1)P(S_1) + P(I|S_2)P(S_2)}$$

where $S_1 \cap I$ represents the intersection of events S_1 (a state of nature) and I (information predicting the state of nature); and the expression $P(S_1|I)$ is read as the `probability of state S_1 given information message I.' Notice the analysis involves the general probability of having the prediction $P(I)$; and importantly the probability of error $P(I|S_2)$ -- false or misleading information. There is also the plain misfortune of becoming the unhappy member of a statistical minority $P(S_1)$ regardless of the information.

In comparison, the risk analysis practiced by security professionals is highly simplistic, and ignores such advances in applied statistical decision theory. But, while avoiding the complexity that accompanies statistical modelling of probability, risk and uncertainty, risk analysis ignores the effects of *luck* and *guesswork* on its accuracy.

Risk analysis thus becomes an exercise in simple guesswork. The approach does attempt to provide the best scientific structure in which to couch the guesswork. However, almost all variations of the technique have been challenged on their approach to the original quantification of threat probabilities and costs. For example in Saltmarsh and Browne's extensive review, advantages and disadvantages of the various approaches are listed. Comments on comparative disadvantages include:

Courtney's Method:

> `Insofar as orders of magnitude are utilised, little credence can be placed in the far ends of the scales; in other words, a great difference exists between $100,000 and $1,000,000, for example.'

CITIBANK Method:

> `Results are inexact at best, if not merely guesses.'

Relative-Impact Measure (RIM) Method:

> `RIM is only useful in a relative sense. it can only be compared with other RIM values.'

Jerry Fitzgerald and Associates Method:

> `Demands much guesswork. Is susceptible to "garbage in, garbage out".'

Perhaps the last comment summarises the danger most clearly. If the original estimations are invalid, then the probability arithmetic which follows these is complete nonsense. Yet, little work has been attempted in validating the use of estimated `probability of threat occurrence' or `cost of threat occurrence'. The use of experts, committees or Delphi techniques is sometimes advocated (*cf.* Parker), however, it would be extraordinarily difficult to establish the usual accuracy of the estimations. Too many variables exist, such as duration of consideration, decision environment, skill and experience of estimators, and number of estimators.

Therefore, risk analysis is a technique which, while widely accepted, is commonly applied in an environment devoid of feedback as to the success of its use. Knowledge as to the effectiveness of this approach is too easily buried beneath its ignorance of Bayesian risk and uncertainty, and thus the oppressing dependence on good or bad `luck'. For example, the probability of a particular data center's damage owing to a major civil disruption (e.g., rioting or war) could be estimated to be `once in 250 years'. But, as riot damage to data centers is a dramatically uncommon threat, such an estimation is arbitrary. However, suppose controls were rejected under this evaluation. If the disaster should strike the following month, the inclination would be to shrug off the loss due to lack of substantial controls as `bad luck'. Rather than admit that the controls selection process is flawed, `luck' allows risk analysis to retain its credibility as a design technique for properly reaching a highly professional or scientific decision without supporting evidence.

Thus, the impact of strong positive or negative feedback as to the history of effectiveness of this technique is deftly eliminated. In the Popperian sense, the accuracy of risk analysis cannot be refuted or falsified. Risk analysis, as practiced, thus fails a primary hallmark of the scientific method.

Risk Analysis As Social Practice

Circumstances may motivate the *professional systems developer* to ignore or avoid security in information systems. The highly subjective nature of risk analysis permits its abuse in the development of unsafe information systems.

Cost justification of new information systems is often a difficult process. DeMarco, for example, offers a design methodology which deftly skirts this crucial element by declaring it to be subordinate to the particular organisation's customary tactics [DeMarco, 1979]. The `return on investment' for improved information systems sometimes involves quantified estimates at least as dubious as those involved in risk analysis. Company directors are appropriately guided by their obligation to shareholder profitability. Security controls, such as access control software, may fare badly in the comparative light of `return on investment', particularly when presented as an `option' on the proposed system which may be added later.

When a proposed information system is dramatically infeasible if controls were in place, consideration of the system without the required controls automatically ensues. Rather than recognising that the system is ultimately infeasible, it is considered feasible in light of certain risks. Risk analysis can be applied in statistically reducing the importance of these controls.

The ultimate result may be that management will proceed with the creation of an unsafe information system without truly understanding the implications of the risks involved in the investment. John Newton points out that many organisations are becoming operationally dependent on complex, interdependent information systems; e.g., automated tellers (ATMs) or point-of-sale (POS) terminals. This dependence renders the continual availability of the system elements as essential to the survival of the organisation. For example, the loss of a POS system for a week might bankrupt the organisation:

`Risk analysis techniques (financial costs of event multiplied by probability of event equals exposure) are not appropriate where business survival is at issue. This point can be illustrated by looking at the financial cost of

losing an entire enterprise valued at, say, $1 billion as the result of an event whose probability of occurrence is 10 million to 1 -- an exposure of $100.' [Newton, 1985]

Circumstances may also motivate the *professional systems manager* to ignore or avoid security in information systems. The shift in the importance of information technology to a position of prominence with regard to the survival of the organisation may have escaped the notice of management. But it may also have escaped the notice of our society as a whole. This could be far more important as an issue.

Organisational management usually recognises its responsibilities to `public welfare' in conducting its affairs. For example, building construction must not only be profitable, but the public can expect the resulting building to be *safe*. However, management may not have yet projected this responsibility into its information system design projects.

One argument notes that our social fabric has become more critically dependent on the thread of technology as its binder. If this technology is fragile, it is an obvious target for those attacking the society, not just the owning organisation (*cf.* Hoffman [1982] or Menkus [1983]). Circumstances such as war or political unrest can subject unsafe systems to unexpectedly strong attacks [Pollak, 1983].

Management is also treating privacy as an economic decision. In public data bases, the astonishing threat of such data aggregations has become evident. For example, an error in the U.S. FBI data base containing suspected criminal information precipitated the arrest, strip-search and imprisoning of an innocent New Jersey woman [Babcock, 1985]. This is not surprising, since a study of criminal history systems discovered three-quarters of the records in one major FBI file were significantly inaccurate [Laudon, 1986]. Little wonder the trend has been to wrench control of these data bases from management hands and place it into those of legislators. Government systems may be currently more visible, but private systems may be next: Schmitt [1982] cites a survey conducted by the Sentry Insurance Company in which 52% of their sample of computer executives expressed the opinion that privacy was inadequately safeguarded by computer systems.

Presently, however, management is unmotivated to treat the safety of new information system projects as a matter of `public responsibility'. Martin Hellman [1984], comparing efforts in obtaining computer controls to the efforts in securing nuclear weapons treaties, finds that the principals seem unwilling to work to obtain the desired result on *reasonable* terms. In the case of computer security, this commodity is only desired strongly enough to warrant acquisition when it is available at virtually no cost. Hellman singles out the security of many electronic funds transfer systems as examples. These, he finds, seem to discover the implementation of adequate controls to be fraught with impossible costs and barriers. Hellman contends that the organisation needs to reorient its values to account for the inevitable importance of security.

Summary of The Critique

The subjective nature of risk analysis, under guise of its appearance as a statistical predictor, is subject to misuse. By overrating its scientific qualities, it may cause the implementation of costly, unnecessary controls. Perhaps worse, it may also allow the deployment of inherently unsafe, fragile information systems.

Risk analysis has major social, scientific, and practical flaws. It crosses epistemological boundaries subversely. As its findings cannot be refuted, it fails at least one essential test of scientific method (refutation). Further, it lacks statistical rigour. Finally, it is subject to misuse as a social technique.

Alternatives to Risk Analysis

The literature on computing and information security does offer alternatives to risk analysis as a tool for justifying the capital investment in security of information systems. However, security professionals have not widely adopted these. Possibly, this is because the alternative meta-control tools are just as undesirable (or even less desirable) than risk analysis.

Risk analysis is a tool for selecting the minimum set of controls necessary for the safety of an information system. This presumes that the selected control set is a subset of all possible controls for that information system (*i.e.*, all technically feasible controls). Figure 1 presents a framework for the search of alternative paradigms for selecting such a subset of controls. The foundations of risk analysis are in engineering, and perhaps this is the most natural profession in which to seek systems design and specification techniques. However, the framework considers other professions, such as medicine, accounting, and law, that might offer improved meta-control, *i.e.*, a controls selection processes.

Improved Statistical Decisions.

Risk analysis evolved from industrial engineering thought, which pioneered statistical decision theory in management. The primary reference discipline is physics, where natural science seeks to accurately predict natural events. Probability arithmetic can be used to predict complex (even biological or social) systems. Increased complexity in the system under study requires increased complexity in the statistical model (or systematic problem reduction).

The effect would be the general improvement of risk analysis to take full advantage of the complex power of statistical decision theory. There is evidence, however, such statistical models have not been as successful in mainstream IS practice as risk analysis has been in security practice. For example, Cramer and Smith [1964] found that a `risk preference factor' had to be introduced into the model in order to account for non-linear risk preferences by management (*i.e.*, managers are often averse to substantial losses, even at low probability). Canada [1971] found such preferences to vary almost capriciously with mood and intangible considerations. Gilbreath [1986] coined the term *risk paralysis* for situations in which the obsessive Bayesian risk study defeats an otherwise worthy project.

Corr studied the capital investment decisions of twenty-six North American firms and found that none regularly developed the probability profiles or simulation models described by statistical decision theory. The summary included:

> `The results show that, although the techniques have been available for about two decades, still not much is done to recognise risk in evaluating capital investments. . . . The notion that management does not consider techniques for measuring risk to be important surfaced frequently in the interviews.' [Corr, 1983]

It seems that, in practice, management decision-makers must struggle to apply Bayesian statistics, even in projects where the outcomes are more predictable than information systems security risks. Thus it is likely that the further mathematical complications that evolve from improvements in risk analysis would only reduce its use.

Certified Professional Opinion

For centuries, the medical profession has relied on professional opinion rendered by a qualified individual. The profession seeks to control the preparation of the professional, and then relies heavily on that individual's structured or unstructured evaluation. Tough cases may call for second or third opinions. Thus, meta-control is exercised in the establishment of an accepted, common body of evolving knowledge.

In the systems design profession, a growing number of professional certification programs have presented themselves to the industry. In Great Britain, the British Computer Society has assumed the equivalent role, only accepting as full members (MBCS), those individuals with a minimum level of professional experience and whose proficiency has been demonstrated though examinations, academic achievement or an attested portfolio. The Society also participates in the certification of chartered systems engineers (CEng). North American certifications include Certified Computer Programmer (CCP), Certified Systems Programmer (CSP), Certified Information Systems Auditor (CISA) and Certified Office Automation Professional (COAP). Perhaps the most well known is the CDP (Certified Data Professional), which is accredited by the Institute for The Certification of Computer Professionals (ICCP). The ICCP, like The International Federation of Information Processing (IFIP), is a society of societies, and its membership includes the Association For Computing Machinery (ACM) and The Data Processing Management Association (DPMA). The CDP is granted after a series of comprehensive examinations in computer science, management, numerical methods and accounting; and also mandates a minimum level of professional experience. In addition, a recent requirement was instituted which demands continuing education or reexamination of certified professionals.

However, meaningful certification requires general agreement on what constitutes an acceptable systems professional. It is plain from the plethora of competing credentials above that this agreement is unachieved. An accepted systems professional would command an accepted body of knowledge. The computer science community may accept a common body of knowledge [Denning, *et al*, 1989], however, in information systems, there are divergent views of the constituent body of knowledge [Banville and Landry, 1989].

An underlying reason for the divergent views found in information systems may be its foundation in the social sciences. Computer science has a foothold in physics and medicine has its foothold in chemistry. Thus positivist repeatability enforces the accepted body of knowledge. Information systems, however, generally must rely on probabalistic social laws instead of natural laws. The science of information systems security therefore must carry all of the disputational baggage that accompanies the philosophy of social science. Wide acceptance of a single unchallenged body of facts and laws is unlikely.

In addition, information systems is an rapidly evolving social field. Inexperience accompanies the application of fast-pace technological developments, and time delays accompany any wide acceptance of the knowledge evolving from such developments. Thus, even if the social sciences could produce a body of accepted information systems knowledge, its shifting technological foundations will delay its feasibility for decades to come.

Standards and Attestation.

Another alternative to risk analysis might be modelled on the accounting profession's meta-controls. With its foothold in the mathematical model of the enterprise as maintained by the bookkeeper, this profession has managed to employ accounting standards and professional attestation to govern behaviour in the complexities of managerial and financial accounting.

Presently, there are only peripheral standards for the design of information systems. `Peripheral' refers to such standards as accounting Standards of Practice [Weber, 1988], or encryption and communication standards such as the Data Encryption Algorithm (ANSI X3.92) or the Financial Institution Message Authentication standard (ANSI X9.9) [Davies and Price, 1984]. In addition, several de facto standards can be found, such as the U. S. Defence Department's `Trusted Computer System Evaluation Criteria' [Schell, 1984].

For meaningful systems standards the information systems field must mature to an age where standards of practice in information systems analysis and design are both feasible and essential. Such standards might address tools (such as data flow diagrams) and system attributes (such as back-up and recovery). In addition, and pertinent to the present subject, the standards could define attestation.

Attestation would entail a practice in which a recognisedly qualified individual (such the CDP or CEng mentioned above) must certify that a new information system design meets minimum industry standards for performance and safety. The same principle is applied when a chartered civil engineer certifies a civil structure design, or a chartered accountant certifies an external audit.

However, if the profession cannot establish an accepted body of knowledge about information systems, it cannot say what constitutes an acceptable information system or a certified individual. It is equally plain that rapidly developing technology will continue to keep various forms of IS too new for an assessment of what constitutes `acceptability' in present-day information systems. Perhaps the acceptable qualities of an information system will always elude the standards setters. Unlike accounting, which can find a footing on models from bookkeeping and math (including statistics), information systems remain primarily socio-technical systems (social systems that employ technology). Truex and Klein [1991] argue that information systems, like grammars, are *emergent*, a continual process of structuration -- always in transit and never arriving. Structure in information systems is marked by temporary regularities captured and imposed by system designers; structures are doomed to fail eventually as the organisational system continues to emerge and the regularities dissolve.

From this view, standards as a meta-control for information systems are unachievable. This is because standard-makers must rely on information systems structures that have proven successful. By the time such success is recognised, the underlying temporary regularities will have disappeared as the referenced system continually emerges onto a different set of regularities.

Rules

The legal profession looks to history and ethics in its selection of controls. A complex set of rules is codefied, often developed to insure that certain undesirable occurrences are not repeated. Control selection is deductively reasoned from experience. Rules thus provide support for meta-control. Knowledge engineers have found that: `Legal rules prescribe social behaviour and have the force of the state behind them.' [Althaus and Backhause, 1989] While this view may ignore the socially established mix of precedent, interpretation, consent and argument upon which the rules of Law are founded; it does recognise that rules are the clearest means for universal predetermination of categorical issues (*e.g.*, right -- wrong, guilt -- innocence).

There are already parallels to this model in the area of controls selection. These are in the form of expert systems that assist in reducing overall system vulnerability in the face of a complex set of interrelated risks. An example is the knowledge base system (KBS) system for the security certification of computer facilities by Carroll and MacIver [1984]. This KBS is based on a two dimensional model contrasting *Components* (*e.g.*, personnel, hardware, and information) against *Attributes* (*e.g.*, reliability, integrity, and authorisation). Another example is the suggestion decision support system (DSS) for security plan analysis by the Naval Postgraduate School [Zviran, *et al*, 1990]. The DSS database details the interaction of resources, threats, risks and countermeasures.

In addition to expert systems, other authorities have discovered rule-based models with foundations in history, rather than decision statistics. These researchers seek universal characteristics that distinguish risky systems from less risky systems. For example, Perrow [1984] uses component interaction and coupling to distinguish between safe and dangerous systems. Baskerville [1988] identified common trends in internal and external threats in order to distinguish security problems in IS automation projects.

Yet these models and expert systems have not been widely adopted by practitioners of security design. The essential fault in the practicality of these rule-based approaches may be shared with the standards approach above. If systems are emergent, and subject to revolutions in technology, the historical basis of such rules may be without merit in predicting future successes in systems. That is, if new, successful systems are only loosely founded on previous experience, then *inductive* rather than *deductive* reasoning must prevail. The deductive value of rule-based systems is less important for emergent systems.

Interaction of The Alternatives

Clearly the alternative paradigms to risk analysis discussed above should not be considered to be orthogonal. These alternatives overlap considerably in practice. For example, professional opinion becomes less subjective when offered in reference to standards (consider the relationship between the opinion of a chartered accountant and standards of accounting practice). Rather than claim orthogonality, the taxonomy has recognised the *dominant* control selection technique in each paradigm. This is plainly identifiable, and adequate in supporting such a review of the alternatives to risk analysis. (For example, diagnostic rules are found in medicine that physicians can apply when seeking to establish their opinions. Certainly there is an overlap, but the professional opinion would outweigh the rules should a conflict arise.)

Still, in each case, this dominant control selection technique presents unattractive characteristics for the security professional. The inevitable mix of these techniques due to this interaction only serves to make the problem more intricate.

For example, consider the interaction of certified professional opinion and standards attestation. In one model, there is a problem in identifying the common body of professional knowledge that underlies the decision to attest. In the other, a problem arises in locating the appropriate standard to which the evolving system must adhere.

Risk Analysis As A Communication Technique

A critical evaluation of risk analysis reveals that it has social, scientific, and practical flaws. It fails an essential test of scientific method, it lacks statistical rigour and is subject to social abuse. Why then, has it survived in practice so strongly?

One argument surfaced in the above with a critical evaluation of the alternatives to risk analysis. This evaluation is a call for the *status quo*. The nature of information systems development renders meta-control adoption from medical, legal or accounting professions implausible.

Another argument may be found by searching for other innate benefits of risk analysis. Perhaps risk analysis is serving an underlying purpose that is not well recognised in the literature. Rather than being important for its statistical foundations, its true value may lie in its usefulness as a communication channel between the designer and management. From this perspective, risk analysis provides an essential communication link between the security and management professionals who must make decisions concerning investments in information systems security. Its simple probability arithmetic allows the security problem to be expressed in a calculus that is familiar to management and in terms (monetary) that permit comparison with capital opportunity costs.

Such a view recognises the importance of the interpretive, subjective contribution of the designer in estimating the costs and probabilities. Perhaps more importantly, the contribution of the designer lies in the identification of *potential* controls for implementation. That is, the designer must find the basic set of possible controls for the system such that an economically feasible subset then can be selected.

From this perspective, an improvement in the statistical rigour might even damage the usefulness of the technique. First, such rigour would reduce the subjectiveness of the designers prediction. Since the technique, on the whole, fails to meet the test of scientism, reducing the interpretive power of the observer would diminish, rather than amplify the non-positivist scientific power of the technique. Second, increased statistical complexity may add only noise to the communication channel between management and the designer. The communication bandwidth, *i.e.* its information content, would be reduced. If this indeed is its primary usefulness, an improvement in statistical rigour would be doubly harmful. That is, improved rigour in risk analysis would lack appeal to the practitioner, and might even damage the accuracy of the results by diminishing the interpretive contribution of the security professional.

Thus, the important, wide-spread predominance of risk analysis may be due to its attributes for conveying security implications uncovered by designers to the management decision makers who must authorise the controls implementation. The probability arithmetic is the language for expressing a subjective, but well-founded, professional opinion.

Conclusion

Perhaps surprisingly, the contribution of this paper thus lies in its well-argued defence of the risk analysis technique in practice. While it lacks attributes that support its ostensible function as a controls selection technique, it provides an excellent communications channel and is certainly preferable to the alternatives.

Alternatives to risk analysis do not promise much improvement. Models developed in other professions for meta-control fail to meet the needs of systems designers. Information systems does not have the concrete body of knowledge on which to base certification, like the medical community. Future information systems are not predictable enough to allow design standards to develop, like the accounting community. Finally, the emergent nature of information systems defeat historical rules, like those on which the legal community is founded.

Thus, risk analysis remains our strongest technique of meta-control. Its chief benefit is not found in its power for predictive modelling, rather it provides an excellent means of communication of the interpretive knowledge of the security designer to the investment decision-makers in management. That is, it superbly provides management with the professional opinion of the designer in capital investment terms.

Risk analysis predominates because it can be both subjective and less rigourous, precisely the two most commonly criticised
qualities. From a broader perspective, this is consistent with claims that, in general, practitioners are autonomous in the construction of epistemological frameworks which are appropriate for the circumstances at-hand [Baskerville, 1990]. Some professional situations call for knowledge that has been acquired in a very positivist manner -- those, for example, with few critical variables. Other situations may require more interpretive techniques in developing professional knowledge -- perhaps those with large constellations of critical variables.

Thus, practitioners require tools which are valid across a spectrum of philosophical frameworks, *i.e.*, tools that are both positivist and interpretivist in nature. In a sense, risk analysis is one of those tools which is *philosophically versatile*, and thus of emmense value in the security professional's methodological toolkit. Perhaps future research might consider whether the majority of popular systems methods share this distinguishing attribute.

The present study, however, reveals that both enlightened managers and elightened security design professionals should be aware of the essential nature of risk analysis as a tool in meta-control. The validity of the controls feasibility study rests in the subjective experience of the designer, not in the objective prediction of the statistics. Risk analysis is a means for expressing this experience in monetary units. If the designer and management both recognise the potential for social bias, interpretive error and Providence, risk analysis prevails as the most serious design technique for justifying the components of information systems security.

References

Althaus, K. and Backhause, J. An expert system for the modelling of legal norms, in G. Doukidis, F. Land, and G. Miller (eds.), *Knowledge-Based Management Support Systems*. Chichester: Horwood, 1989.

Babcock, C. Online crime suspect system implicated in false arrest, *Computerworld* 19 (19 August 1985): p. 12.

Badenhorst, K. and Eloff, J. Computer security methodology: risk analysis and project definition., *Computers & Security 9*, (June, 1990): pp. 339-346.

Banville, C. and Landry, M. Can the field of MIS be disciplined? *Communications of the ACM 32* (Jan 1989), pp. 48-61.

Baskerville, R. Practitioner autonomy and the bias of methods and tools, in Nissen, H-E; Klein, H. and Hirschheim, R. (eds.) *The Information Systems Research Arena of the 90's Vol II*. Preceedings of the IFIP TC8 WG 8.2 Working Conference, Copenhangen, Dec, 1990, pp. 313-337.

Baskerville, R. *Designing Information Systems Security*, Chichester: Wiley, 1988.

Boehm, B. *Software Engineering Economics*. Englewood Cliffs: Prentice-Hall, 1981.

Canada, J. *Intermediate Economic Analysis for Management and Engineering*. Englewood Cliffs: Prentice-Hall, 1971.

Carroll, J. and MacIver, W. Towards an expert system for computer facility certification, in J. Finch and E. Dougall (eds)., *Computer Security: A Global Challenge*. Amsterdam, North-Holland, 1984, pp. 293-306.

Clements, D. Fuzzy models for computer security system metrics. PhD thesis, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, 1977.

Corr, A. *The Capital Expenditure Decision*. Montvale, New Jersey and Hamilton, Ontario: National Association of Accountants and The Society of Management Accountants of Canada, 1983.

Courtney, R. Security risk assessment in electronic data processing. *AFIPS Conference Proceedings NCC* 46, 1977.
Cramer, R. and Smith, B. Decision models for the selection of research projects, *The Engineering Economist 9*, (Winter, 1964).
Davies, D. and Price, W. *Security for Computer Networks*. Chichester: John Wiley, 1984.
DeGarmo, P.; Canada, J. and Sullivan, W. *Engineering Economy* (6th ed.). New York: Macmillan, 1979.
DeMarco, T. *Structured Analysis and System Specification*. New York: Yourdon, 1979.
Denning, P.; Comer, D.; Gries, D.; Mulder, M.; Tucker, A.; Turner, A. and Young, P. Computing as a discipline, *Communications of the ACM 32* (Jan 1989), pp. 9-23.
Emery, J. *Cost/Benefit Analysis of Information Systems*. The Society of Management Information Systems, 1971 (reprinted in J. Couger and R. Knapp, *Systems Analysis Techniques*, New York: Wiley, 1974).
Fayol, H. *General and Industrial Management*. Belmont: Lake Books, 1987.
Fisher, R. *Information Systems Security*. Englewood Cliffs: Prentice-Hall, 1984.
FitzGerald, J. EDP risk analysis for contingency planning. *EDP Audit Control and Security Newsletter* I (Aug, 1978).
Gallegos, F.; Richardson, D. and Borthick, F. *Audit and Control of Information Systems*. Cincinnati: South-Western, 1987.
Gilbreath, R. *Winning at Project Management: What Works, What Fails and Why*. New York: Wiley, 1986.
Hellman, M. Beyond War: implications for computer security and encryption, in *Computer Security: A Global Challenge* edited by J. Finch and E. Dougall. Amsterdam: North-Holland, 1984, pp. 41-47.
Hoffman, L. Impacts of information system vulnerabilities on society, *1982 NCC Conference Proceedings*. Arlington, Va: AFIPS Press, 1982.
Kleijnen, J. *Computers and Profits: Quantifying Financial Benefits of Information*. Reading: Addison-Wesley, 1980.
Laudon, K. Data quality and due process in large interorganizational record systems, *Communications of the ACM* 29 (Jan, 1986): pp 5-11.
Menkus, B. Notes on terrorism and data processing *Computers and Security* 2 (Jan 1983), pp. 11-15.
NBS, (National Bureau of Standards). Guideline for Automatic Data Processing Risk Analysis, Federal Information Processing Standards Publication FIPS 65, August, 1979.
Newton, J. Strategies for problem prevention, *IBM Systems Journal* 24(Nos 3/4, 1985): pp. 248-263.
O'Mara, J. Computer security, a management blindspot, *Computer Security Handbook*. Northborough, Mass: Computer Security Institute, 1985.
Parker, D. *Computer Security Management*. Reston: Reston, 1981.
Perrow, C. *Normal Accidents: Living With High-Risk Technologies*. New York: Basic Books, 1984.
Pollak, R. Implications of international terrorism on security of information systems, *Proceedings of IEEE INFOCOM 83*. New York: IEEE, 1983, pp 270-276.
Saltmarsh, T. and Browne, P. Data processing - risk assessment, in *Advances in Computer Security Management* Vol. 2, edited by M. Wofsey. Chichester: J. Wiley, 1983.
Schell, R. The future of trusted computer systems, in *Computer Security: A Global Challenge* edited by J. Finch and E. Dougall. Amsterdam: North-Holland, 1984, pp. 55-67.
Schmitt, W. Data security program development: an overview, *Computer Security Journal* (Winter, 1982).
Schön, D. *The Reflective Practitioner: How Professionals Think in Action*. New York: Basic, 1983.
Truex, D. and Klein, H. A rejection of structure as a basis for information systems development, *Proceedings of The Working Conference on Collaborative Work, Social Communications and Information Systems*, Helsinki, 1991.
Weber, R. *EDP Auditing: Conceptual Foundations and Practice 2nd ed*. New York: McGraw-Hill, 1988.
Zviran, M., Hoge, J. and Micucci, V. SPAN -- a DSS for security plan analysis, *Computers & Security 9*, (April, 1990), pp. 153-160.

Table II. Cost/Loss Range Table.

Subjective Cost ($)			**Constant Value (C)**
0	-	10	1
10	-	100	2
100	-	1K	3
1K	-	10K	4
10K	-	100K	5
100K	-	1M	6
1M	-	10M	7
10M	-	100M	8

Table IIII. Probability Range Table.

Subjective Frequency (Time)	**Value (P)**	**Annualized per year**	**Loss Multiplier**
Once in 300 years	1	1/300	.00333
Once in 30 years	2	1/30	.03333
Once in 3 years	3	1/3	.33333
Once in 100 days	4	365/100	3.6500
Once in 10 days	5	365/10	36.500
Once per day	6	365/1	365.00
10 times per day	7	365/.1	3650.0
100 times per day	8	365/.01	36500.

Paradigm	**Essential Reference Discipline**	**Primary Control Selection Technique**
Engineering	Physics	Statistical Decision Theory
Medicine	Biology	Certified Professional Opinion
Accounting	Math	Attestation to Standards
Law	History	Rules

Figure 1. Alternative paradigms for controls selection.

Short Author Biography

Baskerville is an Assistant Professor of Management. He holds Ph.D. and M.Sc. degrees from The University of London, and is a member of the IFIP Working Group on *Information Systems and Organizations*. He is a former security systems engineer, and author of *Designing Information Systems Security*.