Information Assurance Techniques: Perceived Cost Effectiveness

Jose M. Such[a], Antonios Gouglidis[a], William Knowles[a,*], Gaurav Misra[a], Awais Rashid[a]

[a]*Security Lancaster, School of Computing and Communications, Lancaster University, Lancaster, LA1 4WA, United Kingdom*

Abstract

The assurance technique is a fundamental component of the assurance ecosystem; it is the mechanism by which we assess security to derive a measure of assurance. Despite this importance, the characteristics of these assurance techniques have not been comprehensively explored within academic research from the perspective of industry stakeholders. Here, a framework of 20 "assurance techniques" is defined along with their interdependencies. A survey was conducted which received 153 responses from industry stakeholders, in order to determine perceptions of the characteristics of these assurance techniques. These characteristics include the expertise required, number of people required, time required for completion, effectiveness and cost. The extent to which perceptions differ between those in practitioner and management roles is considered. The findings were then used to compute a measure of cost-effectiveness for each assurance technique. Survey respondents were also asked about their perceptions of complementary assurance techniques. These findings were used to establish 15 combinations, of which the combined effectiveness and cost-effectiveness was assessed.

Keywords: Security, Assurance Techniques, Perceptions, Security Assessment

1. Introduction

At the heart of the information assurance process lie the "assurance techniques" that are used to evaluate and measure security. Despite this, and against the backdrop of the trend of year-on-year annual increases of security expenditures for organisations of all sizes [1, 2], the characteristics of assurance techniques remain largely unstudied. This leaves a lingering question unanswered: how do we ensure that the increasing number of trained professionals, products, and services in the information assurance space are deployed and utilised in a cost-effective manner? The necessity of such knowledge increases through the growing number of certifications and legal regulations for organisations of all sizes that mandate a "level" of assurance that must be met.

This study intends to address this gap through a large-scale study on the perceptions of industry practitioners on the value of such assurance techniques. This work is intended to facilitate the economic use and procurement of assurance techniques by entities seeking to evaluate their security posture, inform the design of future assurance schemes which mandate particular assurance techniques, and provide a resource for academic research on cost-effective approaches to assessing security. The key contributions of this paper are:

1. A consistent and coherent assurance terminology to clearly define assurance schemes, targets, techniques, and evidence along with their relationships.
2. The definition of an assurance technique framework consisting of 20 assurance techniques classified across 5 categories, along with the relationships between them.
3. An analysis of the perceptions of 153 industry practitioners about the characteristics (e.g., the effectiveness) of the assurance techniques defined within the framework, both as individual entities and as combinations, along with how perceptions differ between practitioner and managerial roles.
4. The synthesis of perceptions to derive measures of assurance technique cost-effectiveness.

The remainder of this publication is organised as follows. Related literature is introduced in Section 2. Section 3 describes the methodology used within this study. Terminology for the assurance ecosystem is then defined in Section 4, along with the framework of 20 assurance techniques across 5 categories in Section 5. Data on the survey and composition of respondents is presented in Section 6. Section 7.2 then examines the perceptions for individual assurance technique characteristics. A metric for cost-effectiveness is introduced in Section 7.3 along with the results of the analysis. Combinations of assurance techniques are then established, and analysed for their effectiveness and cost-effectiveness in Section 7.4. Section 8 concludes the paper.

2. Related Work

Despite the extensive body of research for information assurance, the techniques with which we measure security

*Corresponding author

Email address: w.knowles@lancaster.ac.uk (William Knowles)

have largely escaped rigorous analysis. Two dimensions of existing literature are explored below: the effectiveness of assurance techniques themselves and the economics of effectiveness.

The discussion of assurance techniques within existing literature has largely fallen on their role within software assurance. In particular, assurance techniques and their use within the Software Development Life Cycle (SDLC) (e.g., [3, 4, 5, 6, 7]), or in rare cases, their use within specific product-focused assurance schemes (e.g., the classification of assurance techniques for use within Common Criteria [8]). The predominant body of work in this area has been instigated by the National Institute of Standards and Technology (NIST) project, Software Assurance Metrics And Tool Evaluation (SAMATE[1]), which is sponsored by the U.S. Department of Homeland Security (DHS). An abundance of publications have been produced under this umbrella[2]; in particular around the topic of source code analysis, with the predominant focus on static analysis (e.g., [9, 10]). SAMATE also performs comparative analyses of static analysis tools as part of its Static Analysis Tool Exposition (SATE) project. The fourth iteration is published as NIST Special Publication 500-279 [11]. Beyond SAMATE, static analysis is notable for receiving wider interest as a topic of academic security research (e.g., [12, 13]), along with its counterpart, dynamic analysis (e.g., [14]). More broadly, a comprehensive review of existing software security assessment tools is presented in [15], focusing on when they can be used, their required skills, and their benefits and drawbacks. One assurance technique that has seen research that includes but spans beyond software assurance, is that of penetration testing (which is also frequently used as a misnomer to describe other assurance techniques, such as vulnerability assessments). Little of this research has looked at measuring the effectiveness of penetration testing; however, the core themes have centred on its potential effectiveness to organisations and the motivations for procuring them (e.g., [16, 17, 18]), ensuring those who conduct penetration tests are appropriately skilled which has a direct relationship with the resulting effectiveness (e.g., [19, 20]), and the methodologies for conducting a successful penetration test (e.g., [21]).

The cost-effectiveness of assurance technique usage is one component within the larger domain of research surrounding the economics of information assurance. Although a marked increase in research activity has been seen here over the past five years (see [22] for an early survey), the emphasis has predominantly fallen on topics such as incentives (e.g., [23]), the related topic of cyber insurance (e.g., [24]), and cyber crime (e.g., [25, 26]), while limited attention has been paid to the economic aspects of assurance techniques – in particular, their cost-effectiveness. Where this exists, the focus has again fallen on software assurance. For instance, [27] investigated the economic impact of inadequate infrastructure for software testing and [28] elaborated on existing approaches to model and assess the cost and value of software.

The scope of assurance techniques falls beyond software assurance, however, and it is in this broader application that this study is concerned: the multitude of assurance techniques, both non-technical (e.g., interviews and observation) and technical (e.g., penetration tests), which can be used in the assessment of security controls (be they technical, organisational or physical). To the authors' knowledge, existing literature has not yet covered such a comprehensive analysis.

3. Methodology

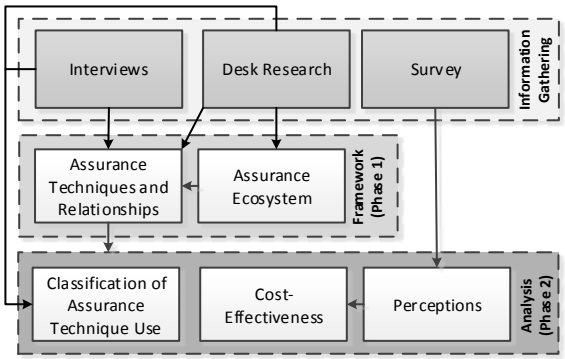

Figure 1: Methodology

This study presents the first comprehensive study of the characteristics of assurance techniques from the perspective of industry stakeholders. The methodology is illustrated in Figure 1. It can be seen to span two phases, with information gathered from three sources: First, desk research examined existing literature and the definition and usage of assurance techniques within 17 assurance schemes (e.g., within standards). Second, 14 targeted interviews (i.e., for particular assurance schemes or scenarios) to understand the role of assurance techniques in practice. Third, an online survey that received responses from 153 industry stakeholders.

This study's first phase defined an assurance technique framework to ensure the consistent and reliable data collection during phase two. Through the desk research the components involved in conducting assurance activities in practice were dissected, and re-constructed here in the form of a terminology model of the assurance ecosystem. Desk research was then further combined with the experiences and clarifications of subject matter experts that were gained through interviews, in order to establish a consolidated but comprehensive set of assurance technique definitions and their classifications. The outputs of phase one can be found in Section 4 and Section 5.

[1] `http://samate.nist.gov/Main_Page.html`

[2] A comprehensive list of SAMATE publications can be found at: `http://samate.nist.gov/index.php/SAMATE_Publications.html`

The second phase involved a stakeholder-supported analysis of this framework. This was primarily achieved through the gathering of perceptions about assurance technique characteristics from 153 stakeholders through an online survey, and their subsequent analysis. Although caution must be pursued in interpreting the subjective construct of perception, it is pursued here with the intent of providing insight into a key group of individuals within the information assurance process: those charged with using and procuring assurance techniques as part of their day-to-day roles. Prior to the public release of the survey a three stakeholder pilot study was conducted; the feedback from this process (provided either through email or in-person) allowed the clarification and refinement of questions that were perceived as unclear or misleading. The assurance technique definitions were provided to stakeholders during the survey to facilitate consistency of understanding. The primary output of the survey was the collection of perceptions about five characteristics of assurance techniques (number of people required, time taken, expertise required, cost, and effectiveness), along with perceptions about which assurance techniques are complimentary when pursuing cost-effective approaches to security assessments. This data is presented and analysed here; this analysis further considered the extent to which perceptions differ between practitioner and management roles, with the intention of providing insights into the drivers and barriers for decision making regarding information assurance. The perceptions of cost and effectiveness were then aggregated into a measure of cost-effectiveness at both an individual and group level (based on perceptions of complementarity). This analysis enabled further analyses of assurance technique use in practice within assurance schemes [3]. The outputs of phase two can be found in Section 6 and Section 7.

4. Terminology

The use of consistent terminology aids comprehension of meaning and facilitates the process of collecting reliable data within the study. However, this study detected, through a review of related literature and publicly available information about assurance schemes, that there were inconsistencies and incoherences in the names and ways assurance techniques are referred to between sources. Therefore, the first contribution of this study is a terminology to describe four basic components of assurance. Each component is described below, and their relationships collectively illustrated in Figure 2.

Assurance Scheme. This encompasses both standards (e.g., ISO/IEC 27001) and qualifications (e.g., CISSP). For both, at least one assurance target is set. In some assurance schemes, there are explicitly defined assurance techniques that should be used to assess targets (represented in Figure 2 as a dashed line). For others, these are set and enforced through an external body (e.g., an accreditation body).

Assurance Target. An assurance target may be either a security control (e.g., asset management) or the competence requirements to assess such security controls (e.g., an individual must possess a certain qualification).

Assurance Technique (also known as an Assurance Activity). A method of assessing an assurance target. There are two types of assurance techniques. Those which assess security controls (e.g., penetration testing) and those that assess the competence requirements for using those assurance techniques (e.g., a multiple choice or lab-based exam).

Audit and Assessment Evidence. The use of an assurance technique to assess an assurance target generates audit or assessment evidence. Such evidence is used to assess compliance to an assurance scheme. This component is produced through the interaction of the three aforementioned components, and can be seen as the component which integrates the ecosystem together.

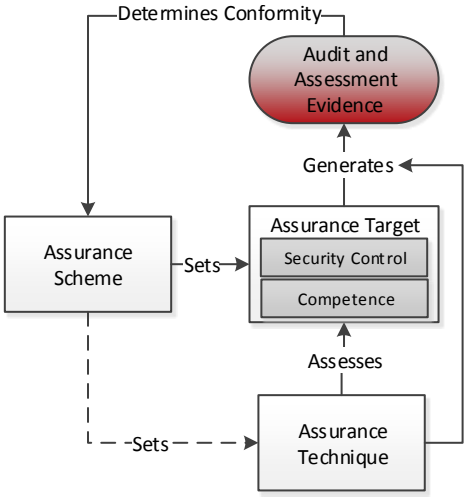

Figure 2: The Assurance Ecosystem

5. Assurance Techniques

Potential variations of assurance techniques are abundant. Therefore, the definition of a consolidated set of assurance techniques is paramount to allow for consistency within the survey and ensuing analysis. This study defines 20 high-level assurance techniques, which are split over 5 categories. Four of these categories represent the broad techniques for assessing assurance targets, in the traditional sense of a security control: Review; Interview; Observe and Test. This is supplemented by a fifth category, Independent Validation, which represents third-party assessment. As highlighted in Section 4 assurance techniques can also exist to assess individual competency for using other assurance techniques. For this a further category

[3] The results of this output can be found in [29].

and 5 assurance techniques were defined, but will not be discussed here; the focus instead, being placed on assurance techniques to assess security controls[4].

The following set of assurance techniques must be distinguished from two meta-techniques.

The first of these is the *audit*, which is more appropriately defined as a process in which other assurance techniques are used to determine conformance to a specification. Assurance techniques in this context generate *audit evidence*. Such assurance techniques may be used directly by *auditors* (i.e., one or more individuals conducting an audit), although equally, an auditee (i.e., the client undergoing the audit) may also use assurance techniques, or procure services that use them (e.g., penetration tests), for which the audit evidence may be used by an auditor.

The second is *risk assessment*, which can be reduced to five consolidated steps [30]: asset identification; threat assessment; vulnerability assessment; risk evaluation (i.e., computing a measure of "risk"); and the recommendation of countermeasures. The assurance techniques that we have defined here are predominantly concerned with that of vulnerability assessment, although some assurance techniques contribute in full or part to the two prior steps (e.g., asset identification is a fundamental step of architectural reviews of operational systems, while threat assessment is explicitly defined here). The appropriate choice of assurance techniques here is paramount, as it is the outputs of these techniques that provide the variables for risk computation, which ultimately influences choices surrounding risk treatment (e.g., the implementation of new security controls). This importance for appropriate assurance technique choice can be extended when examining their role in *risk management*, which goes beyond the scope of a single risk assessment through monitoring and reviewing organisational risk over time. Controls may be implemented as part of the risk assessment process; the level of risk, pre and post-treatment, will then influence the choice of assurance techniques that are used within subsequent iterations of risk assessments. Therefore, inappropriate assurance techniques use can have an impact on the wider risk management process.

The definition of the 20 high-level assurance techniques organised in 5 categories is provided below. Figure 3 visualises assurance technique categorisation and relationships.

5.1. Review

Review of Documented Policies, Procedures, and Processes - The process of analysing the documented specifications (e.g., procedures and security properties) and processes (e.g., managerial) for a component or system under assessment.

[4]The definition of competency assessing assurance techniques and their relationships to those defined here can be found in [29]

Review of Client-Completed Self-Assessment Form - An analysis of a client submitted review of their implementation of assurance targets as set out within an assurance scheme. Self-assessment forms typically consist of a multitude of questions that a client must answer is multiple choice or narrative form.

Threat Assessment - A multi-stage process used to identify and rank the threats to computer software, a component, or IT system. Threat analysis builds upon the analysis of sub-processes such as asset identification and architectural reviews against a security policy.

Architectural Review - An analysis of the components (type, quantity, configuration, etc.) and their relationships within a piece of software, component, or system to determine if their implementation meets a desired security policy.

Configuration Review - A review of the way a system or its software has been configured to see if this leads to known vulnerabilities. Configuration reviews can be passive (e.g., manually checking software versions for known vulnerabilities) or active (e.g., automated build review scanners).

Source Code Review - The examination of source code to discover faults that were introduced during the software development process. Source code reviews are predominantly manual; however, they may be supplemented with automated techniques (e.g., using static analysis tools).

5.2. Observe

Observe - The process of watching a live, operational system to identify real-world deviations from documented assurance targets.

5.3. Interview

Interview - The process of questioning one or more individuals about security-related matters within the organisation being assessed through any medium (e.g., in person or virtually).

5.4. Test

Vulnerability Scan - The process of using an automated scanner on a web application or network to identify vulnerabilities. Discovered vulnerabilities are not exploited.

Penetration Test - A simulated attack on a component or system using similar techniques to that of a real-world malicious attacker. A penetration test may build upon a vulnerability assessment; however, it differs in having an implicit or explicit goal that the assessment attempts to realise (e.g., compromise sensitive data or obtain a certain level of network access). Typically this requires vulnerabilities to be exploited, which would not be undertaken within a vulnerability assessment.

Red Team Exercise - A simulated attack on a system that is given more freedom than is available during a penetration test, in order to more realistically simulate a real-world malicious attacker. This freedom is given in terms

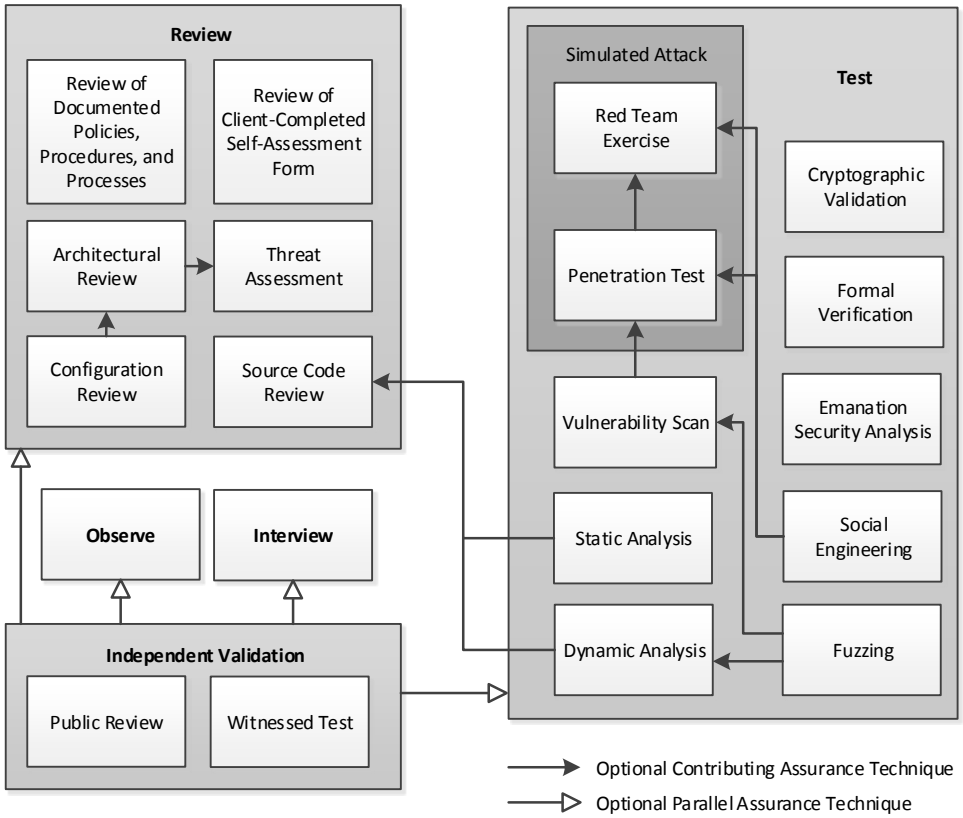

Figure 3: Assurance Techniques

of the engagement's duration (e.g., often months in duration), available human resources (e.g., large teams built around individuals with different specialisms), allowed use of tools (e.g., a heavy use of social engineering is common), and restriction of defender knowledge to test their day-to-day responses to cyber threats.

Social Engineering - An attempt to manipulate one or more human users into performing an action that does not conform to operational procedures. This can be conducted in a manner that is goal-based (e.g., access data) or audit-based (e.g., the percentage of a department vulnerable to a phishing attack).

Static Analysis - Without executing computer software, static analysis attempts to debug and identify potential software vulnerabilities through an analysis of its source code. Static analyses are predominantly automated; however, they may contain some elements of manual interaction (e.g., in order to understand the context and implications of the results). Human-led analyses fall under source code review.

Dynamic Analysis - Once computer software has been executed, this technique attempts to debug and identify potential software vulnerabilities through active methods (e.g., inputting unexpected data through fuzzing) and passive methods (e.g., memory analysis).

Fuzzing - The process of injecting erroneous and unexpected data into an input field in order to trigger faults (e.g., crashes and exceptions) that could be leveraged to discover software vulnerabilities. Fuzzing may be dumb (i.e., random) or intelligent (i.e., with a knowledge of the protocol being tested).

Formal Verification - The use of mathematical techniques for assessing functional properties of information and communication systems.

Cryptographic Validation - A method used to analyse a cryptographic algorithm and/or its implementation within a component or system (e.g., entropy testing).

Emanation Security Analysis - One or more methods used to assess device emanations (e.g., electromagnetic or sound emanations) for the unintentional leakage and disclosure of information.

5.5. Independent Validation

Independent validation occurs when a third party is used to verify the assessment methodology of an assurance technique, or otherwise validate the results of its assessment of assurance targets.

Witnessed Test - The use of an independent witness to provide a second level of verification that the results of an assurance technique are as described.

Public Review - The process of opening a technology, component, or system to wider review by the public. Public reviews may be of documents (e.g., drafts of future cryptographic algorithms) or live systems (e.g., bug bounties).

6. Survey

Expert knowledge was gathered from 153 security practitioners who responded to the survey. Through this we sought to understand: First, the requirements to conduct each assurance technique, including the expertise required, the number of people required, and the time required; Second, the cost of conducting each assurance technique; Third, the effectiveness of each assurance technique; Fourth, which assurance techniques are complementary when pursuing cost-effective security assessments. Such variables are dynamic and change according to the assurance target under assessment. Notably, with regard to the size and nature of the target organisation, and the environment and terms of evaluation. In order to enable meaningful comparisons across assurance techniques and to maximise the fairness of any such comparison, survey respondents were suggested to consider a medium-size commercial scenario when providing their perceptions of assurance techniques. The scenario presented to respondents was defined as follows:

"For each assurance technique, assume a commercial target of medium size. Examples: company with 250 employees; infrastructure with 16 external IPs or 150 internal IPs; web application with one database and 100 static or dynamic pages; product like a Firewall, Router or Switch."

The definitions of assurance techniques presented within Section 5 were also provided to stakeholders to encourage consistency in interpretation and understanding between stakeholders. The composition of stakeholders who responded to the survey are presented in Section 6.1, and the findings in Section 7. In both cases, where percentages are provided, they have been rounded to the nearest whole number.

6.1. Stakeholder Composition

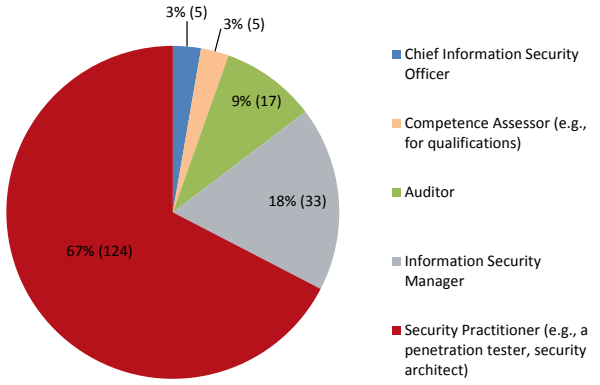

Figure 4: Primary Role of Survey Participants

Primary Role: The distribution of the day-to-day roles of the survey respondents can be seen in Figure 4. The focus of questioning was placed on primary roles; however, some stakeholders deemed their day-to-day roles to span multiple roles defined within this paper, and were reported here as such. The majority of respondents reported being security practitioners (124; 67%), which is advantageous for this research activity as such individuals are those who utilise assurance techniques on a frequent basis. Such reasoning applies further to the auditor role (17; 9%); as described in Section 5, auditing is a meta-technique involving assurance technique usage and evaluation. With the exception of competence assessors (5; 3%), the two remaining roles are managerial: Chief Information Security Officer (5; 3%) and Information Security Manager (33; 18%). While individuals in these roles may not utilise assurance techniques on a day-to-day basis, their perceptions do drive wider organisational security programmes.

Industry Experience: The distribution of the number of years the respondents have spent in the security industry is shown in Figure 5. Notably, 45% respondents have spent over 15 years in the security industry, and 81% over 5 years.

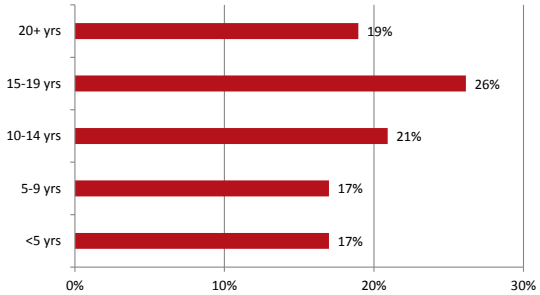

Figure 5: Years Spent in the Security Industry

Assurance Schemes: The assurance schemes that respondents reported to be involved in are shown in Figure 6. These assurance schemes fall into three categories.

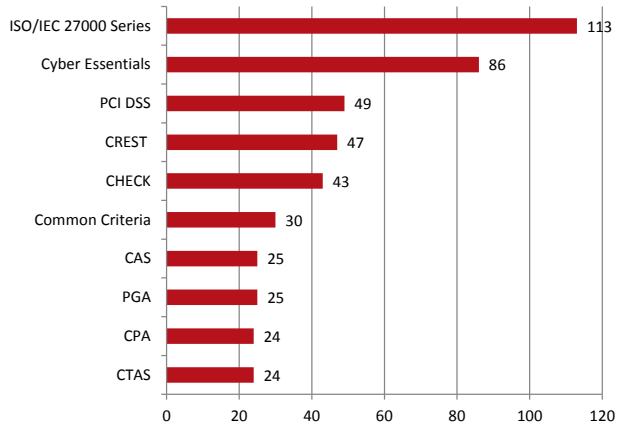

Figure 6: Involvement in Assurance Schemes

First, those where assessment provides a certification for managing security risks. The three most frequent schemes fell into this category. The de facto Information Security Management System (ISMS) standard, ISO/IEC 27001, featured most predominantly with 113 respondents. Cyber Essentials, a UK-specific, entry-level certification was

also widely represented (86). The Payment Card Industry Data Security Standard (PCI DSS) for those who process credit cards placed third (49). CESG, the information security arm of GCHQ, operates two schemes here. CESG Assured Service (CAS) (25) assesses a multitude of services with industry-specific implementations such as CAS-T for telecommunications. The CESG Pan Government Accreditation (PGA) (25) scheme was designed to manage combined risks within the UK public sector.

Second, those for assessing the security of products. This includes ISO/IEC 15408 (Common Criteria) (30), which plays a further role within the CESG scheme, Commercial Product Assurance (CPA) (24). Furthermore, the CESG Tailored Assurance Service (CTAS) (24) is intended to provide answers to specific assurance concerns during the accreditation process.

Third, those proving *organisational* competency to deliver security assessments. CHECK (43) is a mandatory certification for organisations delivering IT Health Checks (a form of penetration test) to UK public sector bodies. CREST (47) is the predominant (optional) alternative for organisations delivering penetration tests and related services in the private sector.

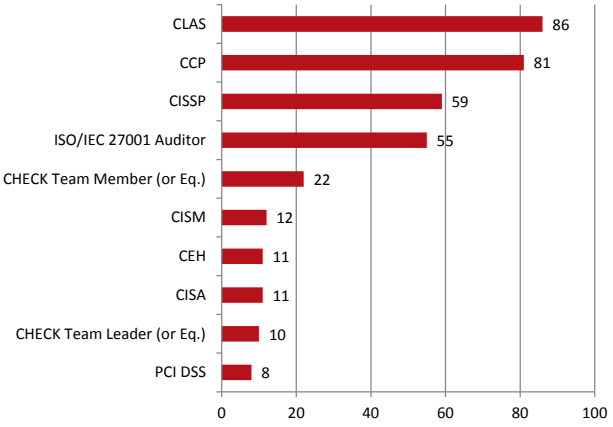

Figure 7: Individual Qualifications Held

Individual Qualifications: The frequency of industry qualifications held by respondents is shown in Figure 7. The total qualification frequency exceeds the respondent count, as respondents were encouraged to list all qualifications held.

Both of the highest frequency schemes are broad in scope and related in design. The CESG Certified Professional (CCP) (81) scheme defines seven roles that candidates can be assessed against. The most popular, however, was the CESG Listed Advisor Scheme (CLAS) (86) which provides additional assurance about individuals through its assessment requiring CCP and security clearance. The respondent discrepancy is understood to be the result of some respondents only listing the "higher" CLAS qualification.

The $(ISC)^2$ Certified Information Systems Security Professional (CISSP) (59) qualification assesses broad domains of security knowledge, and is targeted at managerial roles. A further managerial qualification, although appearing less frequently is the ISACA Certified Information Security Manager (CISM) (12). CISM has been positioned within the market between CISSP and the audit-specific ISACA qualification, Certified Information Systems Auditor (CISA) (11). Auditing qualifications for ISO/IEC 27001 come in a variety of forms (e.g., Lead or Internal Auditor). The high frequency of appearance (55) is not surprising given the high frequency of assurance scheme involvement (see Figure 6). Both ISO/IEC 27001 and PCI DSS (8) show less frequency within qualifications to assurance scheme involvement due to differing individual and organisational competency requirements within the schemes. For example, for PCI DSS, for auditing, both the organisation and individual must achieve Qualified Security Assessor (QSA) status; for vulnerability scans only the organisation must be an Approved Scanning Vendor (ASV); for penetration testing there are no requirements, only guidelines [31].

The remaining qualifications target the penetration testing industry. EC-Council's Certified Ethical Hacker (CEH) (11) is an entry-level qualification. The CHECK Team Member (22) and CHECK Team Leader (10) are both CESG qualifications. Similar to CLAS, the CHECK qualifications require a pre-requisite qualification plus security clearance. In this case, these pre-requisites are based on a mapping to industry "equivalent" qualifications by three providers: CREST, Tigerscheme, Cyber Scheme. A tabular mapping can be found in [32].

7. Results

The emphasis of the following analyses fall on representing the perceptions of the cyber security industry from an holistic perspective. However, given the specialisms inherent within the industry, it is reasonable to suspect that these perceptions may differ between role groups. To evaluate such a hypothesis, independent sample t-tests were also conducted to compare the perceptions of those from Security Practitioner and Information Security Manager roles. This role-based analysis was restricted to the two highest frequency roles, due to the remaining containing inadequate frequencies for a representative statistical analysis. As highlighted in Section 6.1 some stakeholders reported multiple day-to-day roles as part of this study. Such is the case here, with 15 stakeholders holding a day-to-day role as both a Security Practitioner (out of 124) and Information Security Manager (33). This caveat must be considered in the interpretation of results. The t-tests were conducted for each of the 20 assurance techniques, in both the evaluation of respondent confidence in Section 7.1, and the five characteristics under evaluation in Section 7.2. In order to conduct such an evaluation, the qualitative labels to describe perceptions of assurance technique characteristics had to be converted to a quantitative form. A linear mapping between the incrementing qualitative label (e.g.,

of skill level or time required) and an incrementing numerical value (i.e., 1,2,... n) was used for this purpose. For the sake of brevity, it will only be reported where there is a statistically significant difference ($p<0.05$) in the perceptions between roles.

7.1. Stakeholder Confidence

Table 1: Confidence of Respondents About Input

Assurance Technique	Confidence Level			Total
	Low	Med	High	Resp.
Review of [...] PPP	3%	39%	**58%**	93
[...] Self Assessment Form	17%	**45%**	38%	84
Architectural Review	1%	36%	**63%**	83
Configuration Review	4%	**51%**	46%	79
Source Code Review	17%	**43%**	39%	69
Observation	11%	**55%**	34%	56
Interview	8%	38%	**54%**	71
Red Team Exercises	6%	42%	**52%**	66
Penetration Tests	5%	30%	**66%**	88
Vulnerability Scan	6%	36%	**58%**	85
Social Engineering	17%	38%	**45%**	65
Threat Assessment	3%	47%	**51%**	73
Static Analysis	27%	**61%**	11%	44
Dynamic Analysis	26%	**57%**	17%	42
Fuzzing	29%	**49%**	22%	41
Formal Verification	15%	**55%**	30%	40
Cryptographic Validation	20%	**46%**	34%	41
Emanation Security Analysis	34%	**51%**	14%	35
Witnessed Test	10%	**63%**	28%	40
Public Review	37%	**47%**	16%	38

Confidence Level: Respondents were asked to provide a measure of confidence about their answers for each assurance techniques, on a scale of "Low", "Medium" and "High". Table 1 displays the results. Bolded figures represent the category with the highest frequency for each assurance technique. Notable within these findings is a clear trend of decreasing confidence as total respondent frequency decreases. This trend continues throughout the survey data.

Assurance techniques with the highest proportion of "High" confidence answers were found to be *Penetration Testing* (66%) and *Architectural Reviews* (63%). The highest proportion of "Low" answers were *Public Review* and *Emanation Security Analysis*, with a combined "Low" and "Medium" total of 84% and 85% respectively. Potential reasons for the low confidence trend may be the result of some assurance techniques only having specific use cases (e.g., *Emanation Security Analysis* and high security environments) or their relative novelty (e.g., *Public Review* and bug bounties).

No assurance techniques were found to have a statistically significant difference ($p<0.05$) in perceptions of expertise requirements between the Security Practitioner and Information Security Manager roles.

7.2. Assurance Technique Characteristics

Number of People Required: For the reference scenario used within this study, Table 2 shows that at least 75% of respondents perceived that 19 of the assurance techniques could be completed by one or two people. For successful completion by one person, *Vulnerability Scans* and the *Review of Client-Completed Self-Assessment Forms* are particularly notable, having received 84% and 81% respectively. Furthermore, *Architectural Reviews*, *Threat Assessment*, and *Static Analysis* each received figures greater than 70%. Of the perceptions for completion by two people, *Penetration Testing* received 59%; this may be a result of the specialism that characterises the profession (e.g., web applications and infrastructure). One marked exception is that of *Red Team Exercises* for which 58% of respondents perceived that it required three or more people, and 35% perceiving a requirement of four or more. This may once again be the consequence of specialism, given the strong relationship between *Red Team Exercises* and *Penetration Testing* and *Social Engineering* (see Figure 3).

Table 2: Number of People Required

Assurance Technique	Number of People				Total
	1	2	3	4+	Resp.
Review of [...] PPP	**54%**	38%	6%	1%	94
[...] Self Assessment Form	**81%**	13%	4%	2%	84
Architectural Review	**70%**	20%	6%	4%	83
Configuration Review	**63%**	28%	5%	4%	82
Source Code Review	**49%**	29%	12%	10%	69
Observation	**64%**	30%	4%	2%	56
Interview	44%	**48%**	8%	-	71
Red Team Exercises	10%	31%	25%	**33%**	67
Penetration Tests	28%	**59%**	11%	1%	88
Vulnerability Scan	**84%**	14%	2%	-	86
Social Engineering	42%	**45%**	5%	9%	65
Threat Assessment	**73%**	21%	4%	3%	73
Static Analysis	**75%**	18%	5%	2%	44
Dynamic Analysis	**69%**	19%	7%	5%	42
Fuzzing	**78%**	17%	-	5%	41
Formal Verification	**40%**	38%	10%	13%	40
Cryptographic Validation	**56%**	29%	7%	7%	41
Emanation Security Analysis	**60%**	34%	6%	-	35
Witnessed Test	**55%**	33%	13%	-	40
Public Review	**49%**	27%	5%	19%	37

A significant difference was identified for *Witnessed Test* in the perceptions of Security Practitioners (M=1.48, SD=0.71, N=33) and Information Security Managers (M=2, SD=0.53, N=8); $t(14) = -2.28$, $p = 0.039$, two tailed. This difference may be a consequence of the lack Security Practitioner familiarity with *Witnessed Test* usage, which is largely constrained to the compliance assessment process (e.g., for standards and regulations). In this case, perceptions of Information Security Managers may have greater accuracy due to their potential oversight and involvement in such a process, which leads to a greater awareness of the parties involved.

Expertise Required: A cyber security skills gap is a frequently cited issue facing the industry. However, assurance techniques must continue to be appropriately skilled; therefore, to understand job role requirements, respondents were asked what they perceive to be the required consultant role to successfully complete each assurance technique. Results are shown in Table 3.

Table 3: Expertise Required — P: Practitioner (also known as Junior); P(W): Practitioner with Supervision; S: Senior; Pr: Principal.

Assurance Technique	Expertise Required				Total
	P	P(W)	S	Pr	Resp.
Review of [...] PPP	33%	**35%**	32%	-	94
[...] Self Assessment Form	**48%**	29%	23%	1%	84
Architectural Review	8%	12%	**71%**	8%	83
Configuration Review	23%	**44%**	33%	-	82
Source Code Review	14%	25%	**46%**	14%	69
Observation	30%	**41%**	25%	4%	56
Interview	13%	19%	**58%**	10%	72
Red Team Exercises	9%	14%	**55%**	23%	66
Penetration Tests	14%	36%	**45%**	5%	88
Vulnerability Scan	**50%**	36%	14%	-	86
Social Engineering	20%	32%	**45%**	3%	65
Threat Assessment	5%	22%	**59%**	14%	73
Static Analysis	27%	32%	**41%**	-	44
Dynamic Analysis	17%	36%	**48%**	-	42
Fuzzing	29%	**34%**	**34%**	2%	41
Formal Verification	13%	25%	**48%**	15%	40
Cryptographic Validation	7%	12%	**59%**	22%	41
Emanation Security Analysis	9%	**46%**	34%	11%	35
Witnessed Test	18%	30%	**50%**	3%	40
Public Review	**38%**	27%	24%	11%	37

The findings suggest that many assurance techniques require Senior consultants. Notable within this category are *Architectural Review, Interview, Red Team Exercises, Threat Assessment* and *Cryptographic Validation*. Of these, *Architectural Review* ranked the highest with a 71% Senior requirement. Few assurance techniques were perceived to require Principal consultants, although both *Red Team Exercises* (23%) and *Cryptographic Validation* (22%) are prominent within this role. Each also contained distinctly high perceptions of a minimum Senior requirement, giving a combined total for Senior and Principal as 78% and 81% respectively.

If provided with supervision, some assurance techniques were perceived as adequate to be performed by the Practitioner role. *Configuration Review* and *Emanation Security Analysis* can both be seen to have marked increases in this scenario, compared to when to supervision is not provided. This has implications for its cost-effectiveness as it constitutes time from two roles. For Practitioners without supervision, the dominant assurance techniques were *Review of Client-Completed Self-Assessment Forms*, *Vulnerability Scan*, and *Public Review*.

No assurance techniques were found to have a statistically significant difference ($p<0.05$) in confidence levels between the Security Practitioner and Information Security Manager roles.

Time Required: Perceptions of the duration required to complete each assurance technique for the given reference scenario can be found in Table 4. The majority of assurance techniques can be completed within 10 days according to the perceptions of respondents, with the highest frequencies largely falling within the 2-10 day range. Of those in the range of two days or less, once again *Review of Client-Completed Self-Assessment Forms* and *Vulnerability Scans* are prominent in comparison to other assurance techniques. Both received 79% of their respondents for two days or less, and 52% and 39% respectively for one day or less.

It is noteworthy that a sizeable fraction of respondents perceived some assurance techniques to require greater than 10 days to complete: *Source Code Review* (44%), *Cryptographic Validation* (32%) and *Red Team Exercise* (29%). Perceptions of *Public Review* also totalled 35% in this category; this may be because certain schemes which use this assurance technique are often not goal-specific, instead running for long periods of time and soliciting wider contributions (e.g., bug bounties).

No assurance techniques were found to have a statisti-

Table 4: Time Required to Complete

Assurance Technique	Time Required to Complete						Total
	<1 Day	1 Day	2 Days	2-10 Days	10-20 Days	20+ Days	Resp.
Review of [...] PPP	2%	8%	16%	**55%**	15%	4%	93
[...] Self Assessment Form	21%	**31%**	27%	17%	4%	-	84
Architectural Review	4%	14%	23%	**43%**	13%	2%	83
Configuration Review	11%	14%	21%	**40%**	14%	1%	81
Source Code Review	1%	4%	9%	**42%**	25%	19%	69
Observation	4%	14%	38%	**41%**	2%	2%	56
Interview	14%	18%	**32%**	31%	6%	-	72
Red Team Exercises	-	9%	21%	**41%**	21%	8%	66
Penetration Tests	1%	1%	14%	**66%**	17%	1%	88
Vulnerability Scan	15%	24%	**40%**	19%	2%	-	86
Social Engineering	6%	14%	29%	**45%**	3%	3%	65
Threat Assessment	-	22%	27%	**36%**	12%	3%	73
Static Analysis	2%	11%	**43%**	30%	7%	7%	44
Dynamic Analysis	-	17%	**36%**	**36%**	7%	5%	42
Fuzzing	-	24%	**37%**	27%	10%	2%	41
Formal Verification	3%	10%	13%	**53%**	8%	15%	40
Cryptographic Validation	-	15%	15%	**39%**	12%	20%	41
Emanation Security Analysis	-	20%	**37%**	**37%**	3%	3%	35
Witnessed Test	3%	13%	**48%**	35%	3%	-	40
Public Review	-	14%	14%	**38%**	8%	27%	37

cally significant difference ($p<0.05$) in perceptions of time requirements between the Security Practitioner and Information Security Manager roles.

Effectiveness: Table 5 displays respondent perceptions of how effective the assurance techniques were *"in achieving their objectives"*. Out of the 20 assurance techniques, 11 received their highest frequency of respondents within the "Good" rating of effectiveness. A further seven assurance techniques were predominantly rated as having a "Fair" effectiveness. It is notable that many of these assurance techniques were earlier rated with "Low" or "Medium" answer confidence within Table 1. Only two assurance techniques received more than 50% of their responses with a combined rating of "Fair" and "Poor": *Review of Client-Completed Self-Assessment Forms* and *Public Reviews*. In particular, it is striking that *Review of Client-Completed Self-Assessment Forms* received a 31% "Poor" rating.

Table 5: Effectiveness of Assurance Techniques — P: Poor; F: Fair; G: Good; VG: Very Good; E: Excellent

Assurance Technique	Effectiveness					Total
	P	F	G	VG	E	Resp.
Review of [...] PPP	2%	27%	**46%**	19%	5%	93
[...] Self Assessment Form	31%	**35%**	29%	4%	2%	84
Architectural Review	1%	6%	**46%**	39%	8%	83
Configuration Review	-	21%	**41%**	35%	3%	80
Source Code Review	7%	9%	**49%**	28%	7%	69
Observation	2%	**39%**	38%	18%	4%	56
Interview	7%	22%	**36%**	32%	3%	72
Red Team Exercises	5%	5%	27%	**38%**	26%	66
Penetration Tests	-	5%	34%	**47%**	15%	88
Vulnerability Scan	6%	28%	**37%**	24%	5%	86
Social Engineering	5%	22%	**45%**	15%	14%	65
Threat Assessment	-	19%	**47%**	29%	5%	73
Static Analysis	2%	**41%**	39%	18%	-	44
Dynamic Analysis	2%	**40%**	36%	19%	2%	42
Fuzzing	5%	**39%**	32%	24%	-	41
Formal Verification	3%	28%	**40%**	30%	-	40
Cryptographic Validation	-	20%	**49%**	24%	7%	41
Emanation Security Analysis	6%	**40%**	34%	20%	-	35
Witnessed Test	10%	28%	**45%**	15%	3%	40
Public Review	16%	**37%**	26%	13%	8%	38

At the higher end of the scale, only *Penetration Tests* and *Red Team Exercises* received their highest frequencies above a "Good" rating, while also being the only assurance techniques to receive greater than 50% of respondents for a combined effectiveness of "Very Good" and "Excellent". Despite this, however, other assurance techniques continued to report respectable effectiveness ratings. In particular, *Architectural Review*, *Configuration Review* and *Interview* saw 30% or more respondents with a "Very Good" effectiveness. Furthermore, *Architectural Review* received a combined 47% "Very Good" and "Excellent" effectiveness.

The role-based analysis identified two instances of significant differences in stakeholder perceptions for two assurance techniques. First, a significant difference was identified for *Social Engineering* in the perceptions of Security Practitioners (M=3.16, SD=1.06, N=56) and Information Security Managers (M=2.43, SD=0.85, N=14); $t(24) = 2.73$, $p = 0.011$, two tailed. These results may reflect a greater familiarity by Security Practitioners about the potency of *Social Engineering* in simulated security assessments where human behaviour is often seen as a weak link in defensive operations. Such reasoning may explain the higher mean effectiveness score for Security Practitioners. Second, a significant difference was identified for *Formal Verification* in the perceptions of Security Practitioners (M=3, SD=0.88, N=32) and Information Security Managers (M=3.57, SD=0.53, N=7); $t(14) = -2.24$, $p = 0.042$, two tailed. Unlike *Social Engineering*, *Formal Verification* is a largely theoretical exercise, which may have reflected negatively in the perceptions of Security Practitioners, and would explain the higher mean rating of Information Security Managers.

Table 6: Cost of Assurance Techniques — C: Cheap; M: Moderate; E: Expensive; VE: Very Expensive; EE: Extremely Expensive

Assurance Technique	Cost					Total
	C	M	E	VE	EE	Resp.
Review of [...] PPP	19%	**68%**	13%	-	-	93
[...] Self Assessment Form	**65%**	31%	4%	-	-	84
Architectural Review	11%	**58%**	28%	4%	-	83
Configuration Review	11%	**66%**	21%	1%	-	80
Source Code Review	4%	26%	**36%**	19%	14%	69
Observation	21%	**64%**	14%	-	-	56
Interview	18%	**57%**	22%	1%	1%	72
Red Team Exercises	3%	21%	**50%**	21%	5%	66
Penetration Tests	1%	43%	**47%**	8%	1%	88
Vulnerability Scan	**52%**	30%	15%	2%	-	86
Social Engineering	15%	**55%**	28%	2%	-	65
Threat Assessment	14%	**55%**	23%	8%	-	73
Static Analysis	14%	**64%**	20%	2%	-	44
Dynamic Analysis	10%	**60%**	29%	2%	-	42
Fuzzing	10%	**66%**	17%	5%	2%	41
Formal Verification	-	**38%**	23%	23%	18%	40
Cryptographic Validation	5%	27%	**34%**	24%	10%	41
Emanation Security Analysis	9%	**43%**	29%	17%	3%	35
Witnessed Test	13%	**45%**	33%	8%	3%	40
Public Review	**42%**	34%	16%	5%	3%	38

Cost: Perceptions of the costs for completing each assurance technique are shown in Table 6.

Three assurance techniques were considered predominantly "Cheap": *Vulnerability Scan*, *Review of Client-Completed Self-Assessment Forms* and *Public Review*. The *Review of Client-Completed Self-Assessment Forms* was perceived to be the cheapest of all assurance techniques, receiving a 65% "Cheap" rating along with a further 31% as "Moderate".

The majority of assurance techniques (13 of 20) received the highest frequency of perceptions within the "Moderate" cost category. Out of these 13, over 50% of respondents rated "Moderate" costs for: *Review of Documented Policies, Procedures and Processes, Architectural Review, Configuration Review, Observation, Interview, Social Engineering, Threat Assessment, Dynamic Analysis, Static Analysis* and *Fuzzing*.

More than 60% of respondents perceived the following techniques to be "Expensive" or greater: *Source Code Review*, *Red Team Exercises*, *Formal Verification* and *Cryp-*

tographic Validation. Penetration Testing while receiving a large proportion of respondents as "Expensive" (47%), received few at the higher ratings (only 9%). A limited number of assurance techniques received an "Extremely Expensive" rating, although both *Source Code Review* and *Formal Verification* are prominent with 14% and 18% respectively.

No assurance techniques were found to have a statistically significant difference ($p<0.05$) in perceptions of cost between the Security Practitioner and Information Security Manager roles.

7.3. Cost-Effectiveness of Assurance Techniques

Based upon the findings for the perceptions of assurance technique characteristics in Section 7.2, a measure of cost-effectiveness was further derived. Cost-effectiveness as defined within this study is a metric that considers the relative cost and effectiveness of each assurance technique, along with the confidence of respondent answers. Due to the infrequent cases of statistical significance in perceptions in the role-based analysis, this section considers cost-effectiveness purely from the holistic perspective of stakeholders. As defined within Section 4, an assurance target can constitute a security control or competence requirement. A separate approach is required for each. Here we present that for measuring the cost-effectiveness of assurance techniques that assess security controls. Computing the cost-effectiveness metric requires a mapping between the qualitative online survey criteria and a quantitative value. The assignments used within this mapping are presented below:

$$\begin{aligned}&\text{Effectiveness\_Weighting} = \\ &\quad \{(\text{Excellent} = 1), (\text{Very Good} = 0.8), \\ &\quad (\text{Good} = 0.6), (\text{Fair} = 0.4), (\text{Poor} = 0.2)\} \end{aligned} \quad (1)$$

$$\begin{aligned}&\text{Cost\_Weighting} = \\ &\{(\text{Extremely Expensive} = 1), (\text{Very Expensive} = 0.8), \\ &(\text{Expensive} = 0.6), (\text{Moderate} = 0.4), (\text{Cheap} = 0.2)\} \end{aligned} \quad (2)$$

Using the mappings, the cost-effectiveness of each assurance technique can be calculated using the following formula:

$$\begin{aligned}&\text{Cost-Effectiveness} = \\ &\quad \text{Overall\_Effectiveness} \times (1 - \text{Overall\_Cost}) \end{aligned} \quad (3)$$

Within this formula, overall cost is subtracted from 1, as it is considered to be inversely proportional to the overall effectiveness. Overall cost and overall effectiveness are each the aggregate values of three sub-measures of cost and effectiveness, where each measure represents all respondent answers at one of the three confidence levels: "High", "Medium" and "Low". For example, the effectiveness rating of all respondents who reported "High" answer confidence. Cost and effectiveness are therefore three element sets, which are aggregated to provide an "overall" value. Within these formulae, we refer to a "Valid Proportion" (VP). This consists of the frequency that a particular variable was chosen by survey respondents, relative to the cumulative frequency of all answers, which excludes missing cases (i.e., where respondents did not answer the question). For example, how many respondents answered "Excellent" effectiveness relative to the total number of answers across "Excellent", "Very Good", "Good", "Fair" and "Poor". VP is then represented within the range of [0,1]. The formula can be expressed as:

$$VP_{value} = \frac{\text{Value Occurrences}}{\text{Total Number of Values}}, VP \in [0, 1]. \quad (4)$$

The values at each confidence level are calculated as:

$$\begin{aligned}\text{Cost}_{\{high,medium,low\}} = (&1 \times VP_{extremely\ expensive} \\ &+ 0.8 \times VP_{very\ expensive} + 0.6 \times VP_{expensive} \\ &+ 0.4 \times VP_{moderate} + 0.2 \times VP_{cheap}) \end{aligned} \quad (5)$$

$$\begin{aligned}\text{Effectiveness}_{\{high,medium,low\}} = (&1 \times VP_{excellent} \\ &+ 0.8 \times VP_{very\ good} + 0.6 \times VP_{good} \\ &+ 0.4 \times VP_{fair} + 0.2 \times VP_{poor}) \end{aligned} \quad (6)$$

Weightings of 0.6, 0.3 and 0.1 are then applied to the confidence levels of "High", "Medium" and "Low" respectively to calculate the overall cost and overall effectiveness.

$$\begin{aligned}\text{Overall\_Cost} = (&0.6 \times Cost_{high} + 0.3 \times Cost_{medium} \\ &+ 0.1 \times Cost_{low}) \end{aligned} \quad (7)$$

$$\begin{aligned}\text{Overall\_Effectiveness} = (&0.6 \times Effectiveness_{high} + 0.3 \\ &\times Effectiveness_{medium} + 0.1 \times Effectiveness_{low}) \end{aligned} \quad (8)$$

To determine the extent that confidence affects cost-effectiveness, a separate analysis was conducted that does not consider it. Therefore, a separate formula is required, although this approach continues to use the same basic components. The output of computing cost (5) and effectiveness (6) are no longer three element sets that represent each confidence level; instead, a singular level consisting of responses from all confidence levels is used. As confidence is not considered, the weighting added through (7) and (8) is not required. Instead, the single element outputs of (5) and (6) are directly used as Overall_Cost and Overall_Effectiveness respectively within (3).

The cost-effectiveness scores for each assurance technique using both the weighted and unweighted formulae is shown in Figure 8.

Within the weighted analysis, the two least cost-effective assurance techniques were found to be *Formal Verification* and *Cryptographic Analysis*. Such positioning may be the consequence of two factors. First, the extensive amount of manual analysis used within such activities, which here requires both a high level of expertise (see Table 3) and time

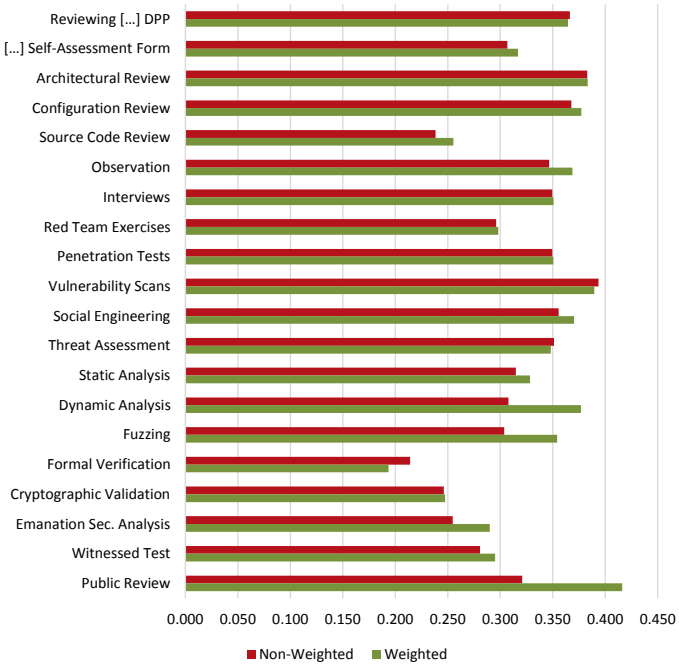

Figure 8: Cost-Effectiveness of Individual Assurance Techniques

for assessment (see Table 4). Second, that the dominant use case lies within niche environments or scope-restricted scenarios. These factors can be further expanded to explain the low ranking of two further assurance techniques for cost-effectiveness. In the case of *Emanation Security Analysis*, the context of its use is significant, due to it being primarily used within high security environments to assess risks that are not considered in the majority of information security scenarios. In contrast, *Source Code Review* has wider application; however, time and competence requirements have both been perceived to be notably high. Such characteristics are not the case for two of *Source Code Review's* optional contributing assurance techniques: *Static Analysis* and *Dynamic Analysis* (see Figure 3). These assurance techniques involve large degrees of automation, and are widely used in combination with *Source Code Review* (notably *Static Analysis* as a first phase). Their contributory nature and automation may explain their distinctly higher cost-effectiveness scores.

For the weighted analysis, the assurance technique that was perceived to be most cost-effective was *Public Review*. Such a finding may support the meteoric rise of crowdsourced bug bounty programs within recent years as an affordable security assessment model. *Vulnerability Scan* closely followed in second place. However, unlike *Public Review* with moderate to high expertise requirements, most respondents perceived *Vulnerability Scan* to require only a Practitioner role (either alone or with supervision) in 86% of cases (see Table 3) while requiring only one person (see Table 2). Further highly cost-effective assurance techniques were *Architectural Review, Configuration Review* and *Dynamic Analysis*.

Despite subtle changes in cost-effectiveness scores, there is consistency between the distributions of the weighted and unweighted analyses. Such a finding suggests confidence plays a limited role in determining *rankings* of cost-effectiveness (i.e., those that are cost-effective and those that are not remain largely consistent). Three exceptions to this exist: *Public Review, Dynamic Analysis* and its optional contributing assurance technique, *Fuzzing*. Of these *Public Review* is the most striking having been the most cost-effective in the weighted analysis, but was influenced in the unweighted analysis by the large percentage of low confidence answers (the highest across all assurance techniques; see Table 1). Due to the consistency and the authors' belief that the confidence of respondents in their perceptions should be considered, further analyses within this paper will use the weighted analysis results.

7.4. Analysing Combinations of Assurance Techniques

Assurance techniques may be used in isolation, or combined in an effort to establish greater synergy in the assessment of security. Such synergy is dependent upon the extent to which assurance techniques are complementary. This section reports on the perceptions of respondents about complementary assurance techniques, before assessing their combined effectiveness and cost-effectiveness.

7.4.1. Establishing Complementary Assurance Techniques

For each assurance technique within the online survey, respondents were asked what they perceived to be its first, second and third most complementary assurance technique. Results were then analysed through the use of stacked bar charts created for each assurance technique. An example can be seen in Figure 9.

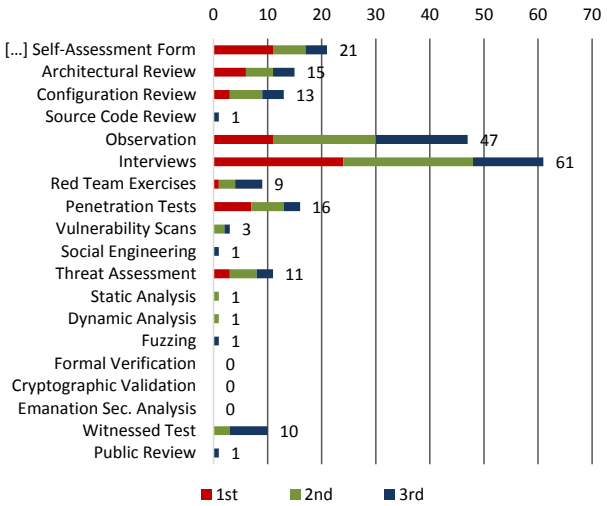

Figure 9: Complementary Assurance Techniques for the Review of Documented Policies, Procedures and Processes

The outcome of this analysis was twofold. First, it enabled the identification of what were perceived to be popular, complementary assurance techniques. Second, it facilitated the grouping of assurance techniques that provide perceived synergistic outcomes.

Table 7: Popular Complementary Assurance Techniques

Assurance Technique	1st	2nd	3rd	Total
Review of [...] PPP	10	0	2	12
[...] Self Assessment Form	0	2	0	2
Architectural Review	2	1	1	4
Configuration Review	1	4	3	8
Source Code Review	0	2	3	5
Observation	0	4	5	9
Interview	2	6	3	11
Red Team Exercises	0	0	1	1
Penetration Tests	3	4	2	9
Vulnerability Scan	2	1	1	4
Social Engineering	0	0	1	1
Threat Assessment	0	0	1	1
Static Analysis	1	1	0	2
Dynamic Analysis	1	3	1	5
Fuzzing	0	1	1	2
Formal Verification	0	0	0	0
Cryptographic Validation	1	0	0	1
Emanation Security Analysis	0	0	0	0
Witnessed Test	0	0	0	0
Public Review	0	0	0	0

Most Commonly Chosen Techniques: Popularity here is split between first, second, and third choice of perceived complementary assurance technique. For each assurance technique analysed, the most popular assurance technique in the first, second, and third categories received an increment in Table 7. For example, as shown in Figure 9, for the *Review of Documented Policies, Procedures and Processes*, *Interview* was perceived to be the highest first choice complementary assurance technique, and was incremented as such within the appropriate table column. The process continued for second and third choice, and then the remaining assurance techniques.

Table 7 shows a clear dominance of *Review of Documented Policies, Procedures and Processes* as the most popular complementary assurance technique, not only in terms of the cumulative frequency across the three categories, but also as the primary complementary assurance technique. The second and third most frequently chosen were *Interview* and *Observe* respectively. These three techniques can be seen to be related, and are considered the fundamental assurance techniques used within auditing. Interestingly, the main proportion of their values lays within the second and third choice complimentary categories, rather than the primary. In contrast, *Penetration Test* ranked joint third in overall popularity, but saw its values even distributed across the three categories. *Configuration Review* also featured frequently, but again, was predominantly considered a second or third choice complimentary assurance technique. It is notable that these five assurance techniques received high cost-effectiveness scores within Section 7.3.

The assurance techniques that earlier received low confidence ratings, also received low popular complimentary values, with the exception of *Dynamic Analysis*. Arguably this may be the result of respondents being unlikely to recommend assurance techniques as complementary when they lack confidence in their answers.

Combinations of Complementary Assurance Techniques: Beyond understanding popularity, the perceptions of stakeholders regarding complementary assurance techniques were also used to establish complimentary combinations. The intention being to determine combinations that can lead to greater levels of effectiveness and cost-effectiveness. For each assurance technique, the cumulative frequency (across all three categories) to which the other 19 assurance techniques were listed as complementary was reviewed. The top three complementary assurance techniques were chosen to create a combination of four. Duplicate combinations were removed from the analysis, of which the authors found five. The remaining 15 combinations can be seen in Table 8. Each combination is assigned a numerical value for reference purposes. An example of this process can be seen by once again drawing on Figure 9. For the *Review of Documented Policies, Procedures and Processes* the three most popular complementary assurance techniques were *Interview, Observation* and *Review of Client-Completed Self-Assessment Forms*. In Table 8 this is represented as Combination 1.

7.4.2. Effectiveness and Cost-Effectiveness for Combinations of Complementary Assurance Techniques

Using the respondent-derived complementary-based combinations of assurance techniques, this section analyses the resulting effectiveness and cost-effectiveness. Aggregation through multiplication was the chosen approach for this purpose, as it allows ease of analysis for uniformly sized combinations.

The effectiveness of the 15 combinations is presented within Figure 10a. The combination ranked most effective is "Comb. 8" (*Red Team Exercise; Penetration Test; Dynamic Analysis; Fuzzing*), and is closely followed by "Comb. 3" (*Red Team Exercise; Penetration Test; Vulnerability Scan; Social Engineering*. Both "Comb. 4" (*Architectural Review; Configuration Review; Penetration Test; Vulnerability Scan*) and "Comb. 10" (*Review of Documented Policies, Procedures and Processes; Architectural Review; Configuration Review; Penetration Test*) also received amongst the highest effectiveness ratings.

Section 7.4.1 earlier highlighted the popularity of the three main assurance techniques used within audits (*Review of Documented Policies, Procedures and Processes; Interview; Observe*). Despite this, of the two combinations that included all three such assurance techniques, both were determined as showing weak effectiveness. Furthermore, in both cases, the fourth assurance technique holds close ties to the real-world audit process. For "Comb. 1" this was the *Review of Client-Completed Self-Assessment Form* which is often used in the preliminary phases of an audit, or as a lightweight substitute for an audit; for "Comb. 2" this was the *Witnessed Test* which is widely used by industry regulators, and whose higher effectiveness may be explained from this independence.

The cost-effectiveness of the 15 combinations is shown in Figure 10b. A change in the distribution compared to

Table 8: Numbered Combinations Defined Based on the Complementary Analysis

Assurance Technique	Combination													
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**
Reviewing [...] PPP	✓	✓						✓	✓	✓	✓	✓	✓		
[...] Self-Assessment Form	✓														
Architectural Review				✓		✓				✓	✓	✓		✓	
Configuration Review				✓	✓					✓					
Source Code Review					✓		✓					✓			✓
Observation	✓	✓							✓						✓
Interview	✓	✓									✓				
Red Team Exercises			✓					✓							
Penetration Tests			✓	✓				✓		✓			✓	✓	
Vulnerability Scans			✓	✓											✓
Social Engineering			✓						✓						
Threat Assessment									✓		✓				
Static Analysis					✓	✓	✓								
Dynamic Analysis					✓	✓	✓	✓							
Fuzzing						✓	✓	✓							
Formal Verification												✓	✓		
Cryptographic Validation													✓		
Emanation Sec. Analysis														✓	
Witnessed Test		✓													
Public Review															✓

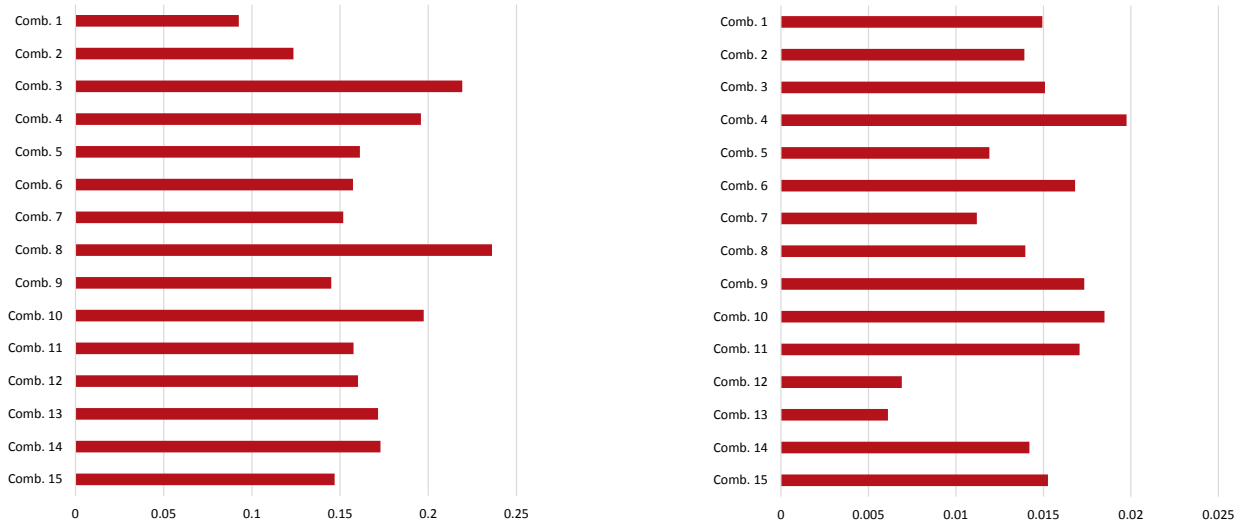

(a) Effectiveness (b) Cost-Effectiveness

Figure 10: Combined Assurance Techniques

the effectiveness analysis is immediately apparent.

The most cost-effective combinations were "Comb. 4" (*Architectural Review; Configuration Review; Penetration Test; Vulnerability Scan*) and "Comb. 10" (*Review of Documented Policies, Procedures and Processes; Architectural Review; Configuration Review; Penetration Test*). Both of which received amongst the highest effectiveness scores. "Comb. 8" (*Red Team Exercise; Penetration Test; Dynamic Analysis; Fuzzing*) and "Comb. 3" (*Red Team Exercise; Penetration Test; Vulnerability Scan; Social Engineering*) which ranked highest for effectiveness, were ranked notably lower for cost-effectiveness (in particular for "Comb. 8"). *Red Team Exercise* is present in both combinations, whose low cost-effectiveness adversely affects combination performance. Despite this, it is notable that both of the most effective and cost-effective combinations include the alternative simulated security assessment assurance technique, *Penetration Test*. Furthermore, although the two audit-focused combinations ("Comb. 1" and "Comb. 2") received the lowest effectiveness scores, their rankings improved considerably for cost-effectiveness, which suggests there is some merit in their use as the dominant method for assessing conformance to assurance schemes.

The least cost-effective combinations were "Comb. 12" (*Review of Documented Policies, Procedures and Processes; Architectural Review; Source Code Review; Formal Verification*) and "Comb. 13" (*Review of Documented Policies, Procedures and Processes; Penetration Test; Formal Verification; Cryptographic Validation*). Each ranked respectably for effectiveness; however, they are the only com-

binations that contain *Formal Verification* and *Cryptographic Validation* which received the lowest individual scores for cost-effectiveness.

8. Conclusion

The assurance technique is a fundamental component of the assurance ecosystem; it is the mechanism by which we assess security to derive a measure of assurance. Despite this importance, the characteristics of assurance techniques have not been comprehensively explored within academic research. This paper addresses this gap through the definition of 20 assurance techniques and their interdependencies, and a stakeholder survey to gather perceptions of their characteristics. This survey received responses from 153 industry practitioners, of which 81% had over five years of experience.

With respect to assurance technique base characteristics the survey collected perceptions across five major areas, for a reference scenario of a "medium" sized target. First, the number of people required to perform assurance techniques, with respondents perceiving that 19 could be completed by 1 or 2 people. Second, the expertise requirements, which found most assurance techniques required Senior Consultants, although Practitioners would suffice if given supervision. Third, the time required, which found most assurance techniques could be completed within 2-10 days, although a fraction of respondents perceived some assurance techniques required longer (e.g., *Red Team Exercise* and *Cryptographic Validation*). Fourth, their effectiveness, with stakeholders perceiving *Red Team Exercise* and *Penetration Test* as the most effective, and *Review of Client-Completed Self-Assessment Form* as the least. Fifth, the cost, where *Review of Client-Completed Self-Assessment Form* and *Vulnerability Scan* were perceived the cheapest, and *Source Code Review*, *Red Team Exercises*, *Formal Verification* and *Cryptographic Validation* as the most expensive. A role-based analysis was further conducted to identify differences in perceptions between Security Practitioners and Information Security Managers. Three cases of statistical significance were found: *Witnessed Test* (people required), *Formal Verification* (effectiveness), and *Social Engineering* (effectiveness).

These base characteristics were then used to derive a measure of cost-effectiveness. This analysis saw *Public Review* and *Vulnerability Scan* as the most cost-effective, and *Formal Verification* and *Cryptographic Validation* as the least.

Survey respondents were also asked what they perceived to be the first, second, and third most complementary assurance technique, for each of the 20 defined assurance techniques. This data was used to establish combinations of complementary assurance techniques, which were then analysed for their effectiveness and cost-effectiveness. The findings determined that the most effective combination was "Comb. 8" (*Red Team Exercise; Penetration Test; Dynamic Analysis; Fuzzing*), with "Comb. 4" (*Architectural Review; Configuration Review; Penetration Test; Vulnerability Scan*) as the most cost-effective.

Three limitations can be seen within the findings presented within this paper. First, the limited number instances of statistical significance between the perceptions of Security Practitioners and Information Security Managers may potentially be a consequence of 15 stakeholders reporting both roles (with role sample sizes as 124 and 33 respectively); in the wider population greater divergence in perceptions may exist between those who clearly identify as occupying a singular role. This hypothesis could not be tested here due to the lack of a adequately sized dataset if multi-role stakeholders were omitted. Second, perceptions by their very nature are subjective constructs and their interpretation must be pursued with caution. However, the emphasis of this paper on the holistic perspective, which encompasses different groups of stakeholders who may potentially differ in their perceptions, mitigates this to some degree. Third, the qualitative labels used within the ratings scales of the survey were chosen for their relatable, and easily understood terminology; however, there are alternative approaches that in retrospect may have been more appropriate to minimise the effect of subjectivity (e.g., Likert Scales).

The findings presented within this paper are intended to provide guidance for the cost-effective design and implementation of security programmes and future assurance schemes, along with establishing the foundations for related research activity. Future work in this area will follow three tracks. First, the assessment of respondent perceptions for assurance techniques that can be used to assess individual competencies for using other assurance techniques (e.g., those used within qualification assessments). Second, an analysis of how assurance techniques are used to assess conformance to assurance schemes in practice. The authors have already prepared a dataset for this analysis. The framework for its creation along with its results across 17 assurance schemes can be found in a separate technical report [29]. Third, the application of these assurance techniques to particular environments (e.g., safety-critical such as Industrial Control Systems). Future research should seek to complement the results obtained in this paper through determining to what extent expert perceptions may align with other more objective approaches, such as those based on security metrics; however, this should be pursued with the caveat that these security metrics are known not to render fully-accurate results.

Acknowledgements

This cyber security research was funded by the UK government.

References

[1] Department of Business Innovation and Skills, 2014 Information Security Breaches Survey, Tech. rep. (2014).

[2] PricewaterhouseCoopers, Managing cyber risks in an interconnected world: Key findings from The Global State of Information Security Survey 2015, Tech. rep. (2014).

[3] R. L. Jones, A. Rastogi, Secure Coding: Building Security into the Software Development Life Cycle, Information Systems Security 13 (5) (2004) 29–39.

[4] G. Mcgraw, Software Security, IEEE Security & Privacy Magazine 2 (2) (2004) 80–83.

[5] B. Arkin, S. Stender, G. McGraw, Software penetration testing, IEEE Security and Privacy Magazine 3 (1) (2005) 84–87.

[6] G. McGraw, Software Security, Datenschutz und Datensicherheit - DuD 36 (9) (2012) 662–665.

[7] N. Davis, Secure Software Development Life Cycle Processes, Tech. rep., Software Engineering Institute (2013).
URL `http://resources.sei.cmu.edu/asset_files/whitepaper/2013_019_001_297287.pdf`

[8] D. Jackson, D. Cooper, Where do Software Security Assurance Tools Add Value, in: Workshop on Software Security Assurance Tools, Techniques, and Metrics. SSATTM05, 2005, pp. 14–21.

[9] P. E. Black, Counting Bugs is Harder Than You Think, in: 2011 IEEE 11th International Working Conference on Source Code Analysis and Manipulation, IEEE, 2011, pp. 1–9.

[10] P. Black, Static Analyzers: Seat Belts for Your Code, IEEE Security & Privacy Magazine 10 (3) (2012) 48–52.

[11] NIST, Special Publication 500-297: Report on the Static Analysis Tool Exposition (SATE) IV, Tech. rep. (2013).
URL `http://samate.nist.gov/docs/NIST_Special_Publication_500-297.pdf`

[12] B. Chess, G. McGraw, Static analysis for security, IEEE Security and Privacy Magazine 2 (6) (2004) 76–79.

[13] A. Bessey, D. Engler, K. Block, B. Chelf, A. Chou, B. Fulton, S. Hallem, C. Henri-Gros, A. Kamsky, S. McPeak, A few billion lines of code later, Communications of the ACM 53 (2) (2010) 66–75.

[14] E. J. Schwartz, T. Avgerinos, D. Brumley, All You Ever Wanted to Know about Dynamic Taint Analysis and Forward Symbolic Execution (but Might Have Been Afraid to Ask), in: 2010 IEEE Symposium on Security and Privacy, IEEE, 2010, pp. 317–331.

[15] D. Washington Navy Yard, B. A. Hamilton, Software security assessment tools review, Mar 2 (2009) 145.

[16] D. Geer, J. Harthorne, Penetration testing: a duet, in: 18th Annual Computer Security Applications Conference, 2002. Proceedings., IEEE Comput. Soc, 2002, pp. 185–195.

[17] P. Midian, Perspectives on Penetration Testing, Computer Fraud & Security 2002 (6) (2002) 15–17.

[18] M. Bishop, About Penetration Testing, IEEE Security & Privacy Magazine 5 (6) (2007) 84–87.

[19] K. Xynos, I. Sutherland, H. Read, E. Everitt, A. Blyth, Penetration Testing and Vulnerability Assessments: A Professional Approach, in: International Cyber Resilience Conference, 2010, pp. 126–132.

[20] A. Tang, A guide to penetration testing, Network Security 2014 (8) (2014) 8–11.

[21] H. Thompson, Application penetration testing, IEEE Security and Privacy Magazine 3 (1) (2005) 66–69.

[22] R. Anderson, T. Moore, The economics of information security, Science (New York, N.Y.) 314 (5799) (2006) 610–3.

[23] R. Anderson, T. Moore, S. Nagaraja, A. Ozment, Incentives and information security, Algorithmic Game Theory (2007) 633–649.

[24] R. Pal, L. Golubchik, K. Psounis, P. Hui, Will cyber-insurance improve network security? A market analysis, in: IEEE INFOCOM 2014 - IEEE Conference on Computer Communications, IEEE, 2014, pp. 235–243.

[25] N. Kshetri, The simple economics of cybercrimes, IEEE Security & Privacy Magazine 4 (1) (2006) 33–39.

[26] R. Anderson, C. Barton, R. Böhme, R. Clayton, M. J. G. van Eeten, M. Levi, T. Moore, S. Savage, Measuring the Cost of Cybercrime, in: R. Böhme (Ed.), The Economics of Information Security and Privacy, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 265–300.

[27] G. Tassey, The economic impacts of inadequate infrastructure for software testing, Tech. Rep. 011 (2002).

[28] A. Drommi, D. Shoemaker, J. Ingalsbe, J. Bailey, N. Mead, Models for assessing the cost and value of software assurance, Tech. rep., Software Engineering Institute, Carnegie Mellon University (2007).

[29] J. Such, A. Gouglidis, W. Knowles, G. Misra, A. Rashid, The Economics of Assurance Activities, Tech. Rep. SCC-2015-03, Security Lancaster, Lancaster University (2015).

[30] A. Jones, D. Ashenden, Risk Management for Computer Security: Protecting Your Network & Information Assets, Butterworth-Heinemann, 2005.

[31] PCI Security Standards Council, Information Supplement: Penetration Testing Guidance, Tech. rep. (2015).
URL `https://www.pcisecuritystandards.org/documents/Penetration_Testing_Guidance_March_2015.pdf`

[32] W. Knowles, A. Baron, T. McGarr, Analysis and recommendations for standardization in penetration testing and vulnerability assessment: Penetration testing market survey, Tech. rep., British Standards Institution (BSI) (2015).