
Research paper

Policy, statistics and questions: Reflections on UK cyber security disclosures[†]

Chad D. Heitzenrater^{1,2,*} and Andrew C. Simpson²

¹U.S. Air Force Research Laboratory Information Directorate, 525 Brooks Road, Rome, NY 13441, USA and

²Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD, UK

*Corresponding author. E-mail: chad.heitzenrater@cs.ox.ac.uk.

[†]Approved for Public Release; Distribution Unlimited: 88ABW-2016-3991 20160815

Received 30 September 2015; revised 9 June 2016; accepted 25 July 2016

Abstract

Over the last 3 years, the UK Government, through the Department for Business, Innovation & Skills (BIS), has taken a lead in the area of public disclosure on corporate cyber intrusions via their Information Security Breaches Survey. The recent development of the Cyber Essentials scheme by the same department presents a unique opportunity for reasonably correlated data to be analysed against public policy. We describe some initial steps in undertaking such an analysis by performing standard economics calculations on this data. Through the examination of three key questions that are central to the relationship between these documents, economic implications of the existing policy are highlighted against the reported threats. Somewhat inevitably, the results echo the well worn ‘it depends’ answer to the question of cyber security expenditure need; nevertheless, in doing so, they do help indicate the dependencies.

Key words: information security economics; security policy analysis; security investment; Cyber Essentials; Information Security Breaches Survey.

Introduction

Well-developed insights into the role of economics within information security relies on the development of both a theoretical and an empirical basis of understanding. Theoretical models are built from empirical observation, as observation drives the empirical analysis of real-world situations. The two strands of research are necessary elements to fundamental understanding of the factors at play in security—economic and otherwise.

However, these goals are often undermined by the state of research. An uncertainty of fundamental properties of attack and defence [1] and a scarcity of good data [2] challenge the development of sound models. Existing studies are fraught with interpretation, providing estimates and opinions which are, at worst, skewed to sell a product, and, even at best, likely to be formed on different bases [3]. This has led some to characterize the current state as stemming from mis-estimation, uncertainty, absence of information, and ambiguity related to disclosure, bias and missing information [4]. As the literature is increasingly populated with models and theory, empirical analyses are becoming increasingly important as guideposts to continued development of both.

The resulting state of research undermines the validity of the models they have formed, leading to calls for increased focus in the area of empirical studies [5]. We seek to perform such an empirical analysis through the examination of available public data, for the purpose of characterizing the relationship between public policy and the threats that precipitate its creation. To accomplish this, the unique circumstance of threat disclosure and policy issuance having been generated by the same entity is exploited. Recent and ongoing publications from the UK Department for Business, Innovation & Skills (BIS; www.gov.uk/government/publications) have put a focus on computing and its implications, with 26 publications including the term ‘cyber’ as of January 2015. Included in these publications are ongoing threat analyses and reporting for the UK in the ‘Information Security Breaches Survey’ (Breaches Survey), with the latest report in this series having been issued in June 2015 [6]. The same department is also responsible for a scheme that attempts to establish a common basis for cyber security practices, and a level of ‘cyber hygiene’ to be followed by companies seeking to do business with the UK Government. This scheme, known as ‘Cyber Essentials’

(CE) [7], outlines five broad areas of compliance (controls); in turn, each of these areas is broken into between four and seven practices that constitute the minimal level of exhibited capability to meet the broader security objectives.

While more limited in scope than efforts such as that of Anderson *et al.* [8], our contribution focuses upon three specific questions related to UK Government efforts with respect to cyber security:

1. ‘How do the Cyber Essentials controls relate to the reported threat?’ This is perhaps the most straightforward part of this analysis (albeit the most subjective). The controls called for by CE are examined with respect to their relevance to the related statistics as reported in the Breaches Survey. While this does not require mathematical rigour, these claims are based in the objective reality of computer security literature. This mapping forms the basis for the two subsequent questions, and shed light on the overlap between the stated policy and the reported threat.
2. ‘Is the effort encompassed within the CE controls requisite to the threat?’ While nothing is absolute—let alone the utility and viability of a cyber security scheme—it is reasonable to ask if any insight can be gained as to the relative investment of CE. This is delicate ground, as many assumptions regarding the exact nature of implementation and execution play a large role in cyber security. Just as the construct of the safe and the experience of the lock-picker both play a large role in the success of the bank heist, so too do the skills of the system administrator in configuring the defence interplay with the skills and fortitude of the hacker seeking to infiltrate the system. Efforts will be made to identify the assumptions at play, and, where possible, the discussion will focus on the trend of the model over the established measures of information security economics such as Expected Net Benefit of Information Security (ENBIS), Net Present Value (NPV) and Annual Loss Expectancy (ALE). In this way, the conclusions drawn do not seek to be absolute, but, rather, are indicators of the forces at play, in order to shed light on what is otherwise a very complex, intertwined and ‘dark’ subject.
3. ‘How should the threat inform the implementation of Cyber Essentials?’ Following on from the previous question, an examination of the CE practices will seek to investigate the consistency in approach presented. Specifically, we attempt to examine each individual concept and practice relative to the others, identifying overlap and relative coverage. While the intent is not to advocate anything less than full implementation, these kinds of analyses seek to answer the question, ‘if I had only one pound/dollar/euro to spend, where should I put it?’ This is particularly relevant due to the prescriptive nature of the CE scheme, which requires specific technologies to achieve certification. While not resulting in a definitive answer, the analysis sharpens focus and provides increased understanding, as decisions are made regarding limited resources spent on cyber security.

This analysis supplies simple calculations in order to answer some basic questions, with the desired result being the teasing out of some general insights into the current state of cyber security practice and understanding. The purpose in examining these disclosures through the lens of these questions is three-fold. Our first aim is to shed light on the relationships that exist between the threats faced by companies as they have been known and quantified, and the policies they inform and are informed by. Secondly, we contend that the employment of information security constructs on real-world datasets has the potential to contribute to the ongoing development of the theory and practice in this area. Thirdly, we argue that the exercise

of these constructs on such a data source, regardless of its provenance and validity, is representative of the effort to be taken by a security professional seeking to incorporate current information security economic considerations, and as such reflects the current practical state of the field. In undertaking this examination, we hope that insights into the state of what is known, unknown, truth and belief regarding the state and practice of cyber security will start to become clear.

It is important to note that we do not provide a critique of the CE scheme. Rather, the reality of cyber security is that there are simply too many problems to go around, and so it is reasonable to cast a slanted gaze at the whole, with a view to asking which parts are the most salient—not for the purpose of change, but for the purpose of insight. The choice of CE and the Information Security Breaches Survey (ISBS) was due to the openness and availability of data. However, the conclusions drawn are not meant to be specific to the Breaches Survey, Cyber Essentials (CE), or, indeed, to the UK. Rather, the intention is to present some general observations and leading questions that are likely to also hold for many other policies and datasets. In this light, the reader should be careful to read the following sections not as directive, but rather as context for the cyber security decision-making process.

The structure of the remainder of the article is as follows. ‘Background’ section introduces CE and the ISBS. ‘Method’ section describes the method employed for the analysis. ‘Analysis’ section considers each of the questions in turn, and provides a summary of the analysis behind them. Finally, ‘Conclusions and future work’ section concludes by placing our results in context and identifying avenues for future investigation.

Background

The UK has committed £860 million to the development of a national cyber security strategy, under the ‘Keeping the UK safe in cyber space initiative’ (www.gov.uk/government/policies/keeping-the-uk-safe-in-cyberspace). While receiving only a portion of this investment, it is in this context that the ISBS and the CE scheme have been undertaken and launched by the Department for Business, Innovation & Skills (BIS). This coordinated approach does not imply that the participating populations overlap, as only 49% of ISBS respondents are badged or on their way to being badged to CE [6]. However, the relationship between the overall strategy and the resulting policy and initiatives prompt questions regarding the impact of this investment. We consider these mechanisms in the context of this environment, and seek to understand their relationship.

ISBS

The Information Security Breaches Survey (ISBS or ‘Breaches Survey’) is part of an ongoing series of surveys commissioned by BIS since the early 1990s [6]. This annual survey is part of an ongoing effort by BIS to supply information in a systematic and consistent way in order to enable analysis and discussion. The 2015 report [6] follows a set of standardized reports from 2014 [9], 2013 [10] and 2012 [11], introducing an updated structure and additional questions. The survey itself has been conducted by PwC for the past 2 years, in association with Infosecurity Europe and Reed Exhibitions.

Each report is presented as both an executive summary, highlighting the main findings and notable statistics, and the main technical report itself. From 2012 to 2014, this was a consistent 22 page format, while in 2015 the expanded questionnaire and an additional appendix of charts resulted in a page length of 49. Each document

contains the details behind the headlines provided within the executive summary: information about the respondents, breakdowns of the data by size and type of business, type of cyber security incident, and loss incurred as a result. Information introduced in 2015 includes information pertaining to: governance and risk management; ‘bring your own device’ (BYOD) controls; and incident identification, management and reporting practices.

Perhaps most notable within the 2014 survey report is the additional information provided. For the first time, the entirety of the data was made available as a comma-separated value (*.csv) file containing the anonymized responses for all of the participants (www.gov.uk/government/uploads/system/uploads/attachment_data/file/326419/information-security-breaches-survey-2014-technical-report-data.csv/preview). Further enhancing the usability of the data, a website was also created to provide an interface to this data, granting non-programmers the ability to parse the data set and generate graphs of the primary data aspects (<https://dm.pwc.com/HMG2014BreachesSurvey/>). This resource enhances the usability and accessibility of the data, spurring further investigation and use, and has helped in undertaking the analysis described in this article. This practice was continued with the 2015 report, although with an altered format that replaces the purely raw tabular data with rolled-up statistics arranged by survey question (www.gov.uk/government/uploads/system/uploads/attachment_data/file/432034/Information_security_breaches_data_tables.csv/preview and <https://dm.pwc.com/HMG2015BreachesSurvey/>).

CE

The CE scheme, published in April 2014, attempts to ‘make the UK a safer place to conduct business online’ [7]. It is the result of a multi-year effort by the UK Government to address cybersecurity concerns, following the success of the ‘10 Steps of Cybersecurity’ guidance (<https://www.gov.uk/government/publications/cyber-risk-management-a-board-level-responsibility/>) with standards by which organizations may be assessed. CE was born of a collaboration between industry and government partners, to include the Information Security Forum (ISF), the Information Assurance for Small and Medium Enterprises Consortium (IASME) and the British Standards Institution (BSI) [7], and is focused on the five essential mitigations within the ‘10 Steps’ guidance.

Perhaps what is most notable about the scheme is the intended reach. Although created to address a perceived lack of existing standards to meet the goals of HM Government, it is intended to be compatible with established standards such as the ISO 27000 series (<http://cyberessentials.org/background/>). BIS has deemed the content ‘relevant to organisations of all sizes’, noting that, while large organizations would be expected to already have some knowledge or experience with the controls, many small- and medium-sized organizations might not have the necessary support or means. Such organizations are referred to a set of supporting standards and guidance.

The implementation environment for CE is assumed to be that of a ‘traditional’ small business in a modern office setting that does not include significant investments in, e.g., non-traditional computing platforms (e.g. Supervisory Control and Data Acquisition (SCADA) or embedded systems). Such businesses therefore serve as the primary audience for the CE policy [7]. At the heart of this policy are five technical controls required for ‘basic technical cyber protection’ in such an enterprise. These are as follows:

1. boundary firewalls and internet gateways;
2. secure configuration;

3. access control;
4. malware protection; and
5. patch management.

These controls are then each subdivided into between four and seven specific technical measures. One notable aspect is that the goal and construction of each control varies in terms of technical depth and expertise, such that the individual contributions of each technical measure do not contribute equally to the implementation of the overall technical control. This observation is at the heart of the investigation undertaken in ‘Method’ section, as the cost and benefit (from a utility perspective) of each measure is explored.

Method

Given the challenge of comparing a policy with a presentation of statistical fact, effort was made to carefully consider the steps and assumptions involved. With the variability in the size and complexity of corporate defence postures, focus was placed on the category of enterprise defined by the Breaches Survey as ‘small businesses’. These data represent the findings for companies consisting of fewer than 50 employees, with the caveat made by the reports that the data for medium entities (50–249 employees) is ‘similar to the results for the small ones unless stated otherwise’ [6] (with similar statements being made in [9, 10, 11]). As none of the categories investigated state any such caveat, we will bound our analysis with respect to the EU definition of a small to medium enterprise (SME) consisting of up to 249 employees. This definition encompasses micro (0–9), small (10–49) and medium (50–249) enterprises as defined by the European Commission (http://ec.europa.eu/enterprise/policies/sme/facts-figures-analysis/sme-definition/index_en.htm).

ISBS

Despite the additions made to the 2015 version, the Breaches Survey has followed a relatively standard methodology for disclosure over the last 4 years of publication (2012–15). Table 1 summarizes the target category response, and it is notable for the decline in overall number of responses. Table 2 summarizes the overall response by business sector, with roughly 20 sectors represented. Unfortunately, these data are not separated by category (small or large business), such that conclusions as to the relationship between business sector and cyber security are difficult to draw. The potential impact of the survey methodology is addressed in ‘Threats to Validity’ section.

Each survey report is broken into numerous parts, providing insights into attitudes, culture and behaviours, as well as trends on the incidences of security breaches (with 2015 including additional categories as noted in ‘Information Security Breaches Survey’ section). With focus on the latter, information is provided regarding both the frequency of a malicious security incident (74% for small businesses in 2015, up from 60% in 2014 and 64% in 2013) and the instances of serious incidents (25% for small businesses in 2015, down from 50% in 2014 but on a par with the response of 23% in 2013). These breaches are then decomposed by type of incident and reported for each of the 3 years under consideration. Data pertaining to these

Table 1. Response of SMEs to the Breaches Survey, 2012–15

	2015	2014	2013	2012
Total respondents	664	1,125	1,402	447
% SME (<249)	44	48	45	45

Table 2. Response by sector to the Breaches Survey, 2012–15

Sector	2015 (%)	2014 (%)	2013 (%)	2012 (%)
Technology	19.1	19.5	21	22
Financial services	17.3	11.9	23	17
Government, health or education	16.2	17.0	18	21
Retail and distribution	3.6	3.4	4	3
Utilities, energy and mining	3.0	2.0	3	2
Property and construction	2.9	2.3	1	1
Manufacturing	2.7	5.0	2	6
Travel, leisure and entertainment	2.6	2.3	3	2
Telecommunications	1.5	3.6	5	5
Other	31.1	33.0	20	20

Categories for 2013 ('Financial Services', 'Government, Health or Education' and 'Other') and 2014–15 ('Retail & Distribution' in addition to the 2013 categories) are summarized due to changes in categorization between 2012 and 2015.

Table 3. Type of breach (for the percentage of respondents suffering a breach)

	2015 (%)	2014 (%)	2013 (%)	2012 (%)
Infection by viruses or malicious software	63	45	41	40
Theft or fraud involving computers	6	10	16	12
Other incidents caused by staff	27	22	41	45
Attacks by an unauthorized outsider	35	33	43	41

Note that these numbers do not total 100% for a given year, as they are reports of respondents reporting a given breach type and not distribution of breach types.

attack types for small businesses is provided in Table 3. Unfortunately (for our purposes, at least), these data are not then translated into overall financial losses in these reports.

Financial data is, however, reported in the Breaches Survey for the largest single loss per entity in a given year. These data were provided as a combination of costs, which are then compiled into an overall estimate. Contributing costs include 'business disruption', 'legal implication', 'incident response', 'financial loss' and 'reputation damage'. The survey combines estimates for these figures into a rolled up range estimate of £75 200–£310 800 for worst incident cost to small businesses in 2015, continuing an upward trend (£65 000–£115 000 in 2014, £35 000–£65 000 in 2013 and £15 000–£30 000 in 2012). This is further broken out into the nature of the worst breach, mapped to categories mirroring (but not equivalent to) the overall incident types (Table 4). It is clear from the magnitude of these figures that they are likely skewed towards the upper end of the definition of an SME, with this result an outcome of the granularity of the BIS data reporting approach (especially pre-2014 where the raw data was not published). However, using these figures for the SME definition serves the purpose of analysis that considers a worst-case scenario.

CE

As a starting point for analysis, the individual controls specified by CE were examined for their purpose and approach, and grouped into specific technical or procedural means. The five categories listed in 'CE' section, as identified by BIS, serve as the basis upon which specific technical controls are defined. Each control and sub-control in these categories is stated in a prose form similar to a standard

Table 4. Category of worst attacks suffered and largest single loss

	2015	2014	2013	2012
Infection by viruses or malicious software	10%	31%	14%	33%
Attack/unauthorized access by outsiders	40%	23%	18%	9%
Theft or fraud involving computers	0%	4%	3%	1%
Infringement of laws or regulations	20%	4%	4%	1%
Physical theft of computer equipment	0%	0%	4%	5%
Staff misuse of the internet or email	0%	12%	12%	15%
Systems failure or data corruption	10%	7%	23%	34%
Theft or unauthorized disclosure of confidential information	0%	19%	10%	2%
Compromise of internal systems, and subsequent remote access	20%	N/A	N/A	N/A
Other	N/A	N/A	12%	N/A
Largest single loss (£K)	£75.2	£65	£35	£15
	£310.8	£115	£65	£30

Items above the first horizontal line are used for the analysis in 'Analysis' section, while items below the second horizontal line indicate year-by-year differences in categorization. Numbers below the double line represent the largest single loss range, in thousands of pounds.

Table 5. Mapping of UK BIS category of computer security control to CE specified controls (by name and identifier)

UK BIS Category	CE	
	Technical control(s)	Identifier(s)
Boundary firewalls and internet gateways	Firewall	1.1 and 1.5
	Firewall policy	1.2–1.4
Secure configuration	System administration	2.1–2.4
	Personal firewall	2.5
Access control	Account administration	3.1–3.7
Malware protection	Antivirus	4.1–4.4
	Blacklist	4.5
Patch management	Patching	5.1–5.4

requirement statement, and labelled with an identifier in the [control].[sub-control] format. The resulting mapping is summarized in Table 5, with the reader referred to [7] for further details.

It is clear from Table 5 that the specified controls do not map directly to specific categories of threat; nor do the sub-controls map to specific technical actions—both of which are necessary for an economic analysis. With respect to the latter, each sub-control must be broken into the technical steps for completion, identifying one-time costs versus recurring investment, and providing estimates for items such as time to complete, etc.

As a policy, CE is not prescriptive of specific technical actions, but rather is descriptive of the desired end state. The result is a need to enumerate these unspoken technical actions in order to fully understand the ramification of implementing the stated policy. In some cases, this was a straightforward rewording of the policy into action. For example, the policy

1.5. The administrative interface used to manage boundary firewall configuration should not be accessible from the internet

is easily restated into the action

Turn off external access to the firewall administrative interface

Table 6. Mapping of Information Breaches Survey categories of attack to the CE controls

ISBS Category	CE Technical controls
Infection by viruses or malicious software	Antivirus Blacklisting Patching
Attacks by unauthorized outsiders	Firewall Firewall policy Personal firewall Patching
Other incidents caused by staff	System administration Account administration
Theft or fraud involving computers	None

Subsequently, this can, under some simple assumptions, be assigned a value in terms of either monetary cost or cost of effort (or a combination of both).

Other policies, such as

4.5. Malware protection software should prevent connections to malicious web-sites on the internet (e.g. by using web-site blacklisting)

easily expand into myriad approaches and technical or procedural steps, from which many assumptions can be made as to the most efficient and/or ‘correct’ approach for a given network instantiation or situation.

Each of the malicious incident categories from Table 4 can be mapped against the approaches identified in the CE listing above, resulting in Table 6. Since no direct mapping exists, each control from CE is listed by technology as it relates to the corresponding UK BIS category of control from Table 5.

Some discussion as to the rationale behind this mapping is warranted; fortunately, the Breaches Survey further describes and decomposes each category into more fine-grained actions [9]. Virus detection and mitigation is the main goal of Control 4 (Malware protection), and is additionally supported by Control 5 (Patching) in the ability of the latter to thwart infection once a virus is present. The threat of outsider attack is more difficult to decompose, as, by definition, it ranges in technical manifestation from penetration, to denial of service, to the impersonation of a company (e.g. phishing) or an individual (e.g. identity theft); however, it is arguably well covered by the combination of Control 1 (Boundary firewalls and internet gateways), Control 2 (Secure configuration) and Control 5 (Patching), as is noted. For these categories, focus was placed on technical controls in order to enable economic analysis.

The threat of incidents caused by staff is defined in [9] as having a wide range: from unauthorized access to computer systems to breach of data protection law and loss/leakage of confidential information. This, along with the fourth incident type (Theft or fraud involving computers), is also largely unaddressed within CE. This class of breaches primarily focuses on the physical aspects of cyber security: the theft of machines, of intellectual property, or of time. As such, they have implications for the non-adversarial aspects of cyber security, such as backups, restoration and recovery upon the loss of data, which are not part of the CE focus. Since this is a disconnect between the two documents, these two threat classes are not considered in our analysis; rather, we focus on the remaining two classes.

Threats to validity

The use of CE and ISBS data to perform such an analysis presents challenges that must be considered when interpreting results. These

can be categorized into issues surrounding the survey consistency and methodology, and ambiguity of the underlying data and policy.

Survey consistency

A few changes in the format over the four ISBS reports employed (2012 [11], 2013 [10], 2014 [9] and 2015 [6]) are worth noting, as they have potential impact on the analysis. First, the 2015, 2014 and 2012 reports list specific information for small businesses only, while the 2013 report provides combined data for small and large businesses. The 2015 addition of a new category, ‘Compromise of internal systems with subsequent remote access’, is clearly related to our categories of interest, ‘Infection by viruses or malicious software’ and ‘Attack or unauthorised access by outsiders’ (and likely accounts for some of the variation in their values). Finally, we note that the category ‘Attacks by an unauthorised outsider’ in Table 3 was explicitly labelled as including hacking attempts for the years 2012–2014, while in 2015 this was changed to ‘exclude’ hacking attempts. It is unclear what has been removed from the previous statistics; therefore, we note that—despite being higher in 2015 (35%, from 33% in 2014)—it is apparently an underestimate on the rate when compared to previous years. As a result of these last two points, comparisons between 2015 data and prior data has to be performed with care.

Data from the 2012 to 2015 ISBS reports appear to report inconsistent numbers for previous years in the same measures at times; for instance, the overall probability of a ‘malicious security incident’ in 2013 for a small business was reported as 64% in the 2014 report, but 76% in the 2013 report (page 10, figure 19 in both reports). The analysis reported in this article utilizes the data for each year as presented in the year reported; i.e., the 2014 report will serve for the source of 2014 data, the 2013 report will serve for the 2013 data, etc. This has the effect of generally using higher rates of occurrence (and complements the ‘worst case’ approach that has been taken).

Methodological challenges

With regard to survey methodology, questions can be raised as to the soundness and validity of the statistics reported within the Breaches Survey. Prior to 2008 ISBS data was collected using a formal sampling method; however, in recent years this was changed to a self-selected survey in order to increase participation. As argued in [12], cybercrime surveys—as with any survey of a population where the characteristic of interest is highly concentrated among a small segment of the population—are ripe with errors regarding robustness and statistical validity. The resulting data are in danger of gross inflation or domination by the responses of a few participants, in addition to various response biases. Evidence of such phenomena are on display within the Breaches Survey, with the very low response rate in certain years resulting in a small sample size for some questions (for instance, the 2015 data for small businesses in Table 4 is based on 10 respondents). While such issues can be partially addressed through the employment of robust statistics, these methods are not appropriate in the case of self-selected responses. Such issues are recognized by the ISBS authors, who state: ‘As with any self-select survey of this nature, extrapolation to the wider population should be treated with caution’ [6].

This is not to say that great care is not taken to understand the impact of the methodology on the data. The consistency in the survey questions, which were subjected to a pilot study during their development, is intended to support trend analysis across multiple surveys (Personal communication to the first author, 18 May 2016). Changes to the methodology, such as the adoption of ‘sticky

sessions’—the ability to persist data between online sessions—have been made in an attempt to ‘increase the quality of the raw data by reducing incomplete responses and potential duplication’ [6]. Data are presented in a transparent manner, citing the number of respondents to each question and supported by the publication of some or all of the source data for 2014 and 2015. Finally, each report is reviewed by a panel of independent reviewers drawn from standards, certification and professional associations [6].

Despite the (potential and realized) issues, the employment of data from sources such as the Breaches Survey—especially when gathered on behalf of policy-making bodies such as BIS—serves as both a yardstick by which we measure the validity of developed models and a representative example of practical application. A lack of well-documented source data is in tension with the need to exercise the principles of information security economics. While these issues challenge the validity of any specific analysis, overall trends and underlying themes can be exposed. Where we have sought to provide a reasoned analysis using the best data available, we also recognize the potential pitfalls in interpretation. Nevertheless, we would argue that the value of the principles outlined outweigh that of any specific findings from the analysis.

Resolving ambiguity

While notable for the amount of data made available, the ISBS lacks in-depth data in many areas. This resulting residual ambiguity in the implementation of security technologies is the root of a majority of the assumptions contained in this analysis. The primary classes of ambiguity include the effective security (e.g. the ‘detection rate’) and the time (manpower) invested. The latter especially drives many of the CE controls, as most are inherently IT-related and therefore require the intervention of someone acting as the administrator of the system. These measures may vary widely depending on the skill, complexity, institutional size and automation assumed—as attested to by anecdotal evidence. This was largely treated as the variable aspect, with starting data assimilated from anecdotal evidence, expert opinion and extensive web searches. Certainly, in all cases this did not result in ‘truth’ and so, where applicable, one of two approaches was taken:

- ‘Use of the best versus worst case’. In some measures, there is a definitive (or at least highly likely) worst-case bound. Where definitive evidence as to the actual state was not obtainable or varied widely, such a bound was used. The choice of best versus worst case was made in order to examine border conditions; for the purpose of gaining insight, this did not reflect a state of reality for a specific institution. An example of where this method is employed is in the use of 1 and 249 as the upper and lower bounds for the number of employees (corresponding roughly to the number of machines), representing the boundary of how a small business is defined.
- ‘Use of the expected value’. For items where a range of discrete values is possible and some notion of the distribution is known, the value used is an expected value based upon the known data. An example of this is the calculation of the cost of antivirus software; while the values range from free (e.g. included in the OS or true freeware) to upwards of £50, an effort was made to utilize available data on antivirus software to produce a value that represents the real-world distribution of use.

For all other calculations, a range of possible values (where applicable) is presented, in order to gauge the trends that result.

Analysis

Our focus now shifts to the challenge of examining these data sets. Due to the updated content and format of the 2015 report with respect to those of the previous 3 years, as well as the more comprehensive nature of the 2014 data release, we restrict our focus to the 2012–2014 data. For the purpose of the analysis, any previous security investment that may have been made by an enterprise is not considered as part of the analysis since this information is not readily available from the BIS data.

How do the CE controls relate to the reported threat?

Consider the scenario of a small business potentially facing a singular worst loss in 2014. Using the data of ‘Method’ section, this translates into the conditional probability of a breach, given the probability that the worst security incident is either the result of infection by malicious software (31%) or an attack/unauthorized access by outsiders (23%). It is not clear from the context that the financial loss data provided in the Breaches Survey represents the loss incurred by the worst of ‘any’ malicious breach, or only those considered ‘serious’; therefore, for this analysis, the overall probability of breach (60%) will be used, rather than the 50% figure representing those who incur ‘serious’ breaches. This represents an assumption that any loss will result in a worst-case loss, and will (somewhat inevitably) contribute to a strengthening of the case for security investment.

In order to examine the rationale behind that investment, our calculations employ the Bernoulli Loss Assumption. Simply stated, this reduces the probability of loss to a binary assumption of a set loss with probability, p , or no loss at all. This is consistent with the context of a singular breach, and can be examined using the Annual Loss Expectancy (ALE), defined as [13]

$$\text{ALE} = p \cdot \lambda.$$

(For this analysis we have chosen a simplified ALE calculation that employs a single probability, p_{breach} , as presented in the ISBS. As such, ‘breach probability’ can be interpreted as a combination of the more traditional p_{threat} and $p_{\text{vulnerability}}$ commonly employed in risk analysis.)

The ALE_0 (loss with no additional security investment) under these assumptions is presented in Table 7, employing the high and low loss event figures for 2012–2014, as determined by $p_{\text{breach}} \cdot \text{loss}$. Despite the probability of attack dropping in 2014 from 2013 (to 60% from 76% for overall breaches), the ALE continues to rise due to a significant increase in loss incurred; this may be indicative of the increase in ‘serious’ attacks (66% from 32%), or may simply be tied to escalating costs.

It is worth emphasizing that these numbers only consider the loss incurred by attacks in the category ‘malicious software’ and ‘attack by outsiders’. For 2014, this represents 54% of the worst attacks and 78% of the overall attacks. The remaining 46% of worst security incidents (22% overall) fall into categories that either have only partial coverage in CE, such as those identified by Control 3 (account administration), or fall into categories that are not addressed by CE. If the other incident categories identified in Table 4 related to staff were to be fully correctable through the implementation of the remaining CE controls, this still leaves 15% of incidents unaddressed by this scheme, and a residual ALE of £9750–£17 250 per enterprise in 2014. The implications of this will be further explored in ‘Conclusions and future work’ section.

Is the effort encompassed within the CE controls requisite to the threat?

An even bleaker picture of business loss due to cyber breaches can be painted by incorporating additional information from the

Breaches Survey. Unfortunately, since the report fails to provide total loss numbers beyond the single worst event, overall numbers are at best estimates.

Looking first at the overall number of attacks resulting in loss, the 2014 report cites the median number of breaches suffered by small businesses as a result of malware infection or attacks by an unauthorized outsider as 3 and 5 respectively, with a median of 6 total incidents overall. Normalizing the per-category number of breaches against the median produces an expectation that four of these six attacks will be of one of these two types under consideration (virus; attack). Performing the same analysis on the 2013 and 2012 data produces incidences of 7.9 and 4.5, respectively. This serves as the estimate of the number of attacks per year.

Since, by definition, these additional attacks will be less than the worst reported breach, an extreme worst-case upper bound could be found by multiplying the ALE_0 by the number of incidents; for the 2012–14 data, this would result in losses of £116 235, £124 883 and £149 040 respectively (using the upper bound of the ALE for a given year). While this is rooted in the survey data, it is also an extreme worst case (median number of incidents—each at the highest end of worst reported loss). As an alternative take on this bound, the lower estimate for worst loss could be employed to achieve estimates of £58 117.50, £67 244.80 and £84 240 per annum respectively. These remain dire numbers.

Next, the loss expectancy is compared to the cost and capability to address it. In order to examine these costs, a simplistic cost model is employed for the cost of the CE controls. This model examines costs as a function of:

- the number of machines, n ;
- manpower, m_n , per machine;
- wage, w , per unit of manpower;
- one-time costs per machine, o_n , to include licence fees, etc.;
- associated one-time manpower amount, m_o ; and
- the fixed cost for investments, I , such as infrastructure (e.g. purchasing a firewall).

This is then calculated, where M represents per-machine costs and O represents one-time costs, as:

$$M + O + I = (m_n \cdot w \cdot n) + [(o_n \cdot n) + (m_o \cdot w)] + I.$$

Using this model, the Expected Net Benefit of Information Security (ENBIS) is employed to examine the rationality of defensive investment. This calculation represents the expected loss without security investment (ALE_0) minus the expected loss with the security achieved by investment s (ALE_s) minus the cost to achieve that security s (assuming monotonicity in security investment):

$$\begin{aligned} \text{ENBIS} &= ALE_0 - ALE_s - s \\ &= (p_0 \cdot \lambda) - (p_s \cdot \lambda) - s. \end{aligned}$$

In general, one should invest in security at the point that $\text{ENBIS} > 0$, representing a positive net benefit. Rewriting to solve for the upper bound of security investment, we have $ALE_0 - ALE_s > s$, for which the cost of controls under consideration can be substituted for s . This leaves

$$M + O + I < ALE_0 - ALE_s$$

with an upper bound of

Table 7. Worst-case annual loss expectancy for the single worst event (2012–14)

	2012 (£)		2013 (£)		2014 (£)	
	Low	High	Low	High	Low	High
ALE_{virus}	3465	6930	3724	6916	12 090	21 390
ALE_{hacker}	9450	18 900	4788	8892	8970	15 870
ALE_{both}	12 915	25 830	8512	15 808	21 060	37 260

This is calculated using the overall probability of an adverse event conditional on the probability of a single serious event being a virus or hacker.

$$M + O + I = ALE_0 - ALE_s.$$

Given the bounds placed on the value of ALE_0 and under the assumption that any current security investment is uncounted towards the resolution of the residual probabilities of attack (as presented in the Breaches Survey), estimation of ALE_s follows once the residual probability of loss is known. For now, it will be assumed that the implementation of the CE controls will result in a residual probability of 99%; this is certainly a generous assumption, but it is helpful in investigating the question of resources that is under analysis here. This provides the information required for the right-hand side of the equation.

Turning attention to the left-hand side requires an estimate for the cost of security,

$$\begin{aligned} s &= M + O + I \\ &= (m_n \cdot w \cdot n) + [(o_n \cdot n) + (m_o \cdot w)] + I. \end{aligned}$$

Fortunately, many of the fixed values can be estimated using publicly available data; as Table 8 lists estimated costs of common cyber security controls based on published surveys, reports and literature. This provides estimates for the costs of infrastructure (I) and the fixed costs per machine (o_n), with the simplifying assumption that the number of machines corresponds on a one-to-one basis to the number of employees.

Providing that wage w can also be estimated from available data, and that the number of machines n is bounded by the definition of small businesses to be within the range 1–249—assuming a single machine per employee—most of the values for the model have been identified. The remaining variables m_n and m_o are the most difficult to estimate, as they represent the manpower investment (per-machine and one-time respectively). This includes not only the time to set up and establish the cyber security measure, but also the cost of operation. While the former might be able to be estimated as some percentage of the IT staff budget (or as a bounded time-frame of effort by a smaller organization), the latter is much more complex. Such costs include not only IT-specific functions such as applying updates, but also the user time spent in the execution of security: time lost to applying and rebooting after a patch; waiting for a virus scan to execute; or in conversation with the help desk upon a (true or false) hit by the antivirus or firewall. Adding to the complexity is that these values are also the most likely to exhibit wide variability, with educated IT staff or competent employees engaging in less time—but also exacting a higher cost per unit of time.

In order to place an estimate on these costs such that the analysis could move forward, relevant literature on this topic was consulted. For ease of use and direct applicability, a model originally developed by Gartner (and utilized in [14]) was chosen, as it permits estimates of costs based upon the distribution of costs between software (29%), hardware (21%), manpower (40%) and outsourced (10%)

Table 8. Fixed cost estimates for material investments relative to CE

Control	Cost (£)	Frequency
Control 1 (Firewall)	222.46 (small) 790.40 (large)	One-time
Control 2.5 (Software firewall)	30.57	Per-machine (One-time)
Control 4 (Antivirus; blacklist)	39.37 24.37	Per machine (One-time) Per machine (Yearly)
Control 5 (Patching)	0.00	N/A

costs. The limitations of this model are well documented [14] and acknowledged here; however, for the purposes of providing an analysis of a highly variable quantity for drawing general conclusions, the benefits of this approach outweigh the loss of precision and accuracy in any specific case. As an example, using this method the manpower required for the deployment of the £790.40 router in Table 8 works out as £1505.22, or roughly 14 working days of time at the going rate for IT support personnel in the UK (£26 597 per year; based on information from [www.payscale.com/research/UK/Job=Information_Technology_\(IT\)_Support_Specialist/Salary](http://www.payscale.com/research/UK/Job=Information_Technology_(IT)_Support_Specialist/Salary)). If anything, this is a low-end estimate of cost which may vary from ‘free’ (the spare time of a sole proprietor, which in fact has value likely greater than the estimate), to consultant costs on the order of £50 to £200 per hour. However, this amount was deemed reasonable for the business size under consideration given the nature of the model. As a sanity check, this was shown to correspond to the Gordon–Loeb $1/e$ (37%) security costs versus expenditure ratio for maximum security investment [15]. Using this methodology, the manpower estimates generated (e.g. for firewall maintenance) meet this criterion against the ALE calculations above in each year, with an average of 16.4%. For the ‘installation year’ this holds in each case except the 2013 low loss estimate, which reduces the ratio to 47.9% (although the overall average remains at 25%).

Using the estimates of fixed numbers above, the analysis can now proceed. Examining various size estimates for an SME to include the boundary cases of a single machine, a small company with up to 49 machines and a medium enterprise involving 249 machines, yields the trend lines of Figure 1. The first aspect of note is the existence of scenarios in which the security investment is not a rational choice under the given assumptions: when the lower estimate of loss is applied, the hypothetical organization at the upper end of the scale is at a loss in each of the 3 years. Conversely, the hypothetical single-system organization exhibits a very high ENBIS. As evidenced by this figure, there are clearly related forces at play in these estimates: the manpower investment (as related to the number of machines) and the loss estimates. This deserves some additional attention, starting with the loss estimates.

Recall that the ‘high’ loss estimates were for the median number of breaches per year at the maximum reported worst-case loss given the probability of breach and probability of the type of breach being malware or hacker-related. Likewise, the ‘low’ bounds were set by the same method using the lower end estimate for the single worst loss. As both use the assumption that each loss would be in the range of the ‘worst’ single loss, an obvious line of questioning involves this loss assumption: what happens if the loss is lower for a given year (or indeed higher, as the trend in loss values continues to rise)? This scenario is presented in Figure 2, using the data for 2014.

Here again, the message that the hypothetical single-machine business should invest in such security is clear: the ENBIS quickly

becomes positive in both the installation and annual case, with a loss above £2000 making this a good investment. While this is admittedly using the lower cost of the firewall in Table 8, it is clear from the estimates for the hypothetical larger company that the scale of the investment is highly dependent on the manpower employed to maintain it. While the difference in the fixed costs between installation and annual maintenance total £12 137.33 under these assumptions, the overall difference is more than £30 000 of manpower in addition. Since manpower in this analysis is inherently tied to the fixed outlay (as a result of employing the Gartner model), these costs are inherently driven on a per-machine basis. Therefore, a higher loss is required the larger the organization due to the investment. The resulting effect on the ENBIS supports arguments for automation: the more the manpower can be reduced, the lower the bar for security to be a sound investment.

A final, more subtle aspect will conclude this portion of the analysis. In the previous analysis, an idealistic assumption was employed such that the security provided through the implementation of these controls achieves a level of 99%; i.e., based on the data employed in the Breaches Survey, the probability of compromise is reduced from 60% (the reported incidence of breach for 2014) to just 1%. Clearly, even with the best practices, most IT professionals would be hard pressed to assume their security is so strong. We seek to answer this question relative to the best known data and estimates for given measures relative to the effort called for by CE. Here, the simpler question of overall effectiveness will be considered. Returning to the high and low estimates of multiple breach loss, the data for 2014 will again be examined against variability in overall security effectiveness from 50% to 99%. This is shown in Figure 3.

It is evident that, as the effectiveness of the controls invested decrease, there is a requisite movement in the point at which the endeavour to deploy cyber defences is no longer a rational investment. For the hypothetical larger small business, this happens quite quickly on the higher loss assumptions for the installation costs: at only around 90% effectiveness these costs overcome the net benefit, as happens at around 64% for the yearly costs. At the lower loss probabilities, the benefit for the install costs is never realized under these assumptions, while the yearly expenditure falls short at around 72%. As before, expenditures at the other end of the spectrum prove quite a good investment, especially at this level of loss; although the upper end of small businesses (49 personnel) calls for closer examination at realistic expectations of effectiveness. We discuss this further in ‘How should the threat inform the implementation of Cyber Essentials?’ section.

Each aspect of this analysis contributes key points regarding considerations that must be made by entities seeking to undertake or expand a cyber defence programme. The effectiveness against the threat, the role of manpower and the scalability of use, and the expected business loss are all key aspects of the trade-space that enables cyber defence to be a meaningful and beneficial undertaking. Where this analysis has demonstrated scenarios where the assumptions inherent in policies such as CE fail, it is worth reiterating that none of these ‘views’ on the data alone provide a realistic or definitive commentary on the CE scheme. Where considerations regarding efficient administration, better automation or cheaper software/hardware would reduce costs, the alternative of more time investment spent in labour-intensive tasks of the lost productivity due to the time spent in execution of these controls may induce requisite or higher costs. By virtue of the ‘worst-case’ estimation approach employed, this analysis should be treated as a boundary case rather




Figure 1. ENBIS against loss estimates per year. The lines show upper (solid lines) and lower (dashed lines) bounds, using Gartner assumptions of manpower investment.




Figure 2. ENBIS with varied loss. The lines show yearly and install expenditures, using Gartner assumptions of manpower investment.

than a representative example. As such, it is intended to highlight the ‘push-and-pull’ between various considerations. Additionally, as discussed in ‘Conclusions and future work’ section, the actual decision to implement cyber defences likely relies on much more than an economic analysis, and by necessity must take into effect regulatory, reputational and ethical considerations.

How should the threat inform the implementation of CE?

Recalling the controls specified in CE (or rather the corresponding technologies identified) and the mapping provided in Table 6, the question of effectiveness can be further examined. Determining effectiveness of a specific measure can be a difficult exercise; much




Figure 3. ENBIS with varied effectiveness of security controls. The lines show upper (solid) and lower (dashed) loss assumptions for various company sizes (1, 49 and 249), using Gartner assumptions of manpower investment.

depends on the specific configuration and deployment scenario, and, to deploy a well worn cliché, ‘the devil is in the detail’. Paywall-protected consultancies often perform analyses of specific software or hardware in order to use that data as part of their competitive edge, leaving only the ‘talking-points’ version reported by popular trade magazines as a common source. The best openly published estimates are presented in Table 6.

The effectiveness of the remaining control—patching—is notably hard to estimate. An initial line of thought would seem to suggest that regular, automated patch application would by definition secure one against all known threats, resulting in an effectiveness close to 100%. However, research, literature and trade publications in the area seem to suggest that this is almost never accomplished, and the bigger the organization (thus, the bigger the target), the longer it takes for the company to roll out patches. This is often due to additional testing to ensure non-interference with home-grown applications [20]. Due to the variability inherent in this function, this control will not be considered; although it is worth noting that, under the assumption of high effectiveness, the values cited in the previous subsection (efficiency of 99%) serve as a guide as to what such an analysis might yield.

We now consider the modelling of the overall costs of given controls. This time, the comparison basis will utilize the Net Present Value (NPV) of the technology in question. NPV seeks to aggregate the benefit to be had over multiple future periods (n) into a singular value, and takes into consideration both one-off and recurring costs [13]:

$$NPV = -c_0 + \sum_{t=1}^n \frac{ALE_{0,t} - ALE_{s,t} - c_t}{(1+r)^t}$$

Employing the same estimates as used in the previous subsection, the NPV for each of the technical controls (except for patching) can

Table 9. Reported effectiveness (Eff.) for various security controls

Control	Eff.	Reported ranges/notes
Antivirus	75%	Reported ranges of 5% [16] to 75% [17]
Firewall	60%	Study cited 60% ‘out-of-the-box’ and
Firewall policy	80%	80% only with skilled administration [18]
Blacklists	73.5%	Lowest coverage for a given malware class by all major AV vendors in [19]

be calculated. A rate of return of 5% was used as an ad-hoc estimate, as is common in practice for such calculations [13]. These are plotted in Figure 4 using the data from 2012 to 2014.

The controls investigated are as follows:

- Host-based antivirus and blacklists for organizations at the boundary of small business size, both mapped to the probability of loss due to malware. Each consists of a fixed cost plus annual fee for subscription and maintenance costs. As a result, these controls illustrate the trade-space of effectiveness with measures of 75% and 73.8% respectively. For comparison, the values for antivirus are calculated at the boundary of each class of enterprise that comprise the definition of SME: micro (9), small (49) and medium (249), in addition to a single machine (sole proprietor). The values for blacklisting track these with the same delta as shown in the single machine case, and are omitted from the graph.
- A firewall and an ‘ineffective firewall’ both using the estimates for hardware employed in Table 8 for both small and large businesses. The difference in these categories of control is that the ‘ineffective firewall’ corresponds to an ‘out-of-the-box’ configuration, with no specialization in policy or rule set; as




Figure 4. NPV of security controls as calculated using the assumptions contained in this analysis. Solid dots are calculated using the high-bound loss estimate, and circles are calculated using the lower loss estimate.

such, it does not include the manpower cost, but also operates at an effectiveness of 60% vice the managed firewall effectiveness efficiency of 80%.

From this graph, it is readily apparent that the same conditions (rather unsurprisingly) hold: increased cost comes into play with increased organizational size, yielding higher value for controls at the lower size estimate. Likewise, for all except the inefficient firewall, the lower loss represents a lower NPV (the inefficient firewall is assumed to have no manpower costs in set-up or maintenance). However, a notable aspect of this graph is both in the range of results from high to low loss estimates, as well as the number of estimates that result in negative NPV. This underscores an earlier point: effectiveness matters. Deploying just any defence will not result in a benefit unless efforts are made to ensure it remains effective, and, unfortunately for the case of the technical controls under consideration, the effectiveness will be dependent on the recurring updates and increased manpower costs. This does have limits, as the enterprise security mechanisms show some benefit above manpower-intensive security deployed across an enterprise, under the given assumptions and for the same loss expectancy.

Returning to the previous discussion of the Gordon–Loeb security investment model, this distribution roughly holds with the 1/e guideline—with caveats. In the case of the per-host investment this is most clear, with averages of 91.1% for install expenditure and 56.4% annual maintenance investment for antivirus against expected malware loss for the hypothetical upper-end small business. This reflects the previous assertion regarding the relative size of loss versus the size of the organization and its impact on security investment. However, for the case of the ‘out-of-the-box’ firewall versus the maintained firewall, the latter fares far better despite higher costs that more closely meet (and in one case exceeds) the Gordon–Loeb bound. This reinforces the notion that effectiveness matters

(unsurprisingly), but also suggests that the Gordon–Loeb rule may have a lower bound. Certainly, the lower the expenditure, the better the security investment when all else is considered equal; but when taking effectiveness into account, there appears to be a need for more expressiveness in this guideline—especially as the investment tends towards zero. Further investigation on this topic is left as future work.

One way in which one could interpret this graph is in the following postulation: if I had only one dollar/pound/euro to spend on security, which technology is my ‘best-bet’ for application? While this data is based on a number of assumptions unlikely to hold in totality for any real organization, the assumptions made were held constant throughout and thus provide a basis for investigation of relative merit. Based on the assumptions as stated, some indications emerge: as noted, antivirus is preferable to blacklisting, strictly based on effectiveness. For a small business, the best investment appears to be in the firewall, whereas for a larger organization this is definitively so—to the extent that it may prove more beneficial than other measures even if little care is given to its configuration and administration. This is not to advocate for failing to administer, as the NPV of this control is still negative, and the value increases dramatically with the increased effectiveness that manpower investment brings; rather, it is a good case for investing in controls that secure the network overall, and to push for automation in those that must be host-based. Naturally, the compounding of estimates throughout this analysis impacts upon these conclusions, as does the ability to compose protections (were effectiveness measures for such defensive structures known). A far more interesting question is the relative merits of the individual sub-controls, as their cost and residual impact vary widely across the set. The conditions that are required for a deeper analysis of this aspect are considered in ‘Conclusions and future work’ section.

Conclusions and future work

It is worth reiterating that the majority of the analysis presented in this article was built upon hypothetical (albeit realistic) scenarios. A number of very general assumptions have been made to fill in holes left from the Breaches Survey data and to say general things about the utility of the CE controls across the range of organizations. As such, this analysis does not represent any real-world scenario, and should not be used as the basis for making policy regarding the investment or deployment of cyber defences. Many of these echo concerns expressed by Rue *et al.* [1] and by other authors.

- In Question 1, the probability of loss was calculated as the overall probability of a breach multiplied by the probability of the breach being of the type (malware; hacker). The probabilities for these types of breach were treated as being exclusive in order to present a worst-case and to combat the lack of insight into the underlying probability distributions.
- For Question 2, the loss calculated was based on the Breaches Survey numbers for a worst loss in the year considered. The probability of such an event was based on the probability of loss, in the manner stated above, and each event was considered separate and independent, which may often not be the case (as others have noted, e.g. [21]). The expected number of breaches resulting from a malware or hacker event was based on these probabilities and not on the probability of ‘serious’ event that these numbers (presumably) represent. Again, this was to present a worst-case assumption based on incomplete data.
- The financial losses used were a generalization, based upon the worst-loss event suffered. Presumably, events beyond the worst-loss event would be less (perhaps far less) than that reported event. In absence of any information in the Breaches Survey regarding the total losses, the upper and lower bounds were multiplied by the worst-loss numbers to achieve bounds. To this extent, the absolute upper bounds of such losses were employed; however, it is possible that each loss was below the high estimate but above the lower worst-loss estimate, falling in between these lines. It is far more likely that the severity of events is more graduated; i.e., while a single event may have resulted in such a loss, the other events resulted in losses that were distributed along the continuum between no loss and the worst loss. Indeed, the Breaches Survey indicates that such events may be more common, but of varied cost. In this case we have shown that, as the overall cost of events fall, the resulting net benefit is pushed lower quickly—especially as manpower remains a primary cost.
- For the case of Question 3, the effectiveness utilized for calculation was chosen from the best (read: most beneficial) estimate of effectiveness across a range of reported values. Of course, true effectiveness of any such control is dependent on many aspects not considered here, from the vendor technology to the skill of the administrator.

Given these concerns, while the ALE-based approach described herein is not ideal, it is a well recognized method for performing such analysis. To that end, this work rests on the same principles as those that have previously used such methods for measuring cybercrime [3], performing risk analysis in software design [22], conducting quantitative analytics for managing computer security risk [23], and, in close synergy with this work, as inputs to a decision framework for security improvement projects in small companies [24]. Our contribution lies in the application of these techniques to externally produced data sets for the purpose of providing insight. While some have warned that the very notion of quantifying security in

this way may be a ‘weak hypothesis’ [25], it is—as we have noted—continued empirical analysis that will provide the basis for further comparison and discussion. Our future work will build upon such approaches for the purpose of making more informed security design decisions throughout the system life cycle.

The goal of the preceding sections was to examine the constructs of CE to discover what lessons might be learned. These can be summarized as follows.

1. ‘How do the CE controls relate to the reported threat?’ It was shown using ALE calculations that the potential loss resulting from the breaches identified within the Breaches Survey is high, based upon a singular loss event and reported probabilities. Additionally, this trend continues to rise, driven by increases in the probability of serious events and costs associated with resolution and suggesting that investment today may have an even greater future return. While this bolsters the case for investments in cyber security technology—as advocated by the CE scheme—the analysis presented also shows that there remain significant threat gaps that are unaddressed or left open ended. This would seem to indicate a need for additional policy and/or a need for additional technology investment to spur development towards some of these means that are difficult or unavailable to smaller companies.
2. ‘Is the effort encompassed within the CE controls requisite to the threat?’ Implementation of the various aspects of CE lends itself to a wide range of possible costs. Unsurprisingly, the amount of time invested and the number of machines within the organization are the largest contributing factors in the ENBIS calculations; to paraphrase Herley [26], user time is not free—even if that time is employed in the name of security. The analysis performed placed some bounds on the amount of time per machine that results in a positive outcome, and has demonstrated that, even under the most ideal of circumstances, companies must take care when implementing defensive programmes, as their benefit will rely heavily on the amount of time invested and the benefit they provide. To fail on this point can quickly lead to scenarios where the venture fails to provide a solid return on investment for the company—even when the implementation is flawless.
3. ‘How should the threat inform the implementation of Cyber Essentials?’ The controls called for within CE were analysed against real-world effectiveness measures as reported in the literature and trade publications, leading to an ability to form a gross relative comparison between them (all assumptions being held equal). Some results were unsurprising, in that—all else being equal—higher effectiveness translates to better value, even if that comes with a maintenance cost (as long as that cost is moderate and not labour-intensive). The effect of manpower on the overall value proposition of cyber defence was once again confirmed, to the effect that in some cases for our hypothetical larger small business it may prove less beneficial to deploy manpower-intensive security across an enterprise than to deploy less effective, but less consuming, technologies. This should not be seen as a call to throw hardware out and hope for the best, but, rather, to underscore the importance of automation and cost control for technologies that must touch each user node. Rather than an expensive one-time outlay, these scenarios are likely to result in hidden costs that are accrued in small increments, but if left uncontrolled have the potential to overwhelm the benefit of the security investment. It was shown that this finding appears to hold with the Gordon–Loeb model for

maximum security investment, suggesting a need for a more expressive measure that considers effectiveness in establishing a lower bound for security investment.

These questions continue to increase in both importance and complexity. The 2015 Breaches Survey [6] finds that spending on security continues to increase for small businesses (with 44% of small businesses increasing security spending in the last year, up by 27% over 2014), even while the size, scope and number of incidents continue to increase. On the question of CE, about half (49%) have either started, completed, or have plans to implement these controls (although only 6% are ‘badged’ and 23% are still planning). However, 26% of companies do not evaluate the effectiveness of their security expenditure. As small business owners wrestle with allocating investments between certification, technology, policy and training, it is increasingly important that the basis for such investment is rational and well supported.

In addition to seeking answers to these questions, this article encompasses experiences with CE and the ISBS that provides some perspective on these efforts, leading to the ability to make some recommendations to improve their use in future ventures. These can be stated as follows.

- In the case of CE, the primary finding of this analysis is the inconsistency of the depiction and the high level at which these controls are presented. For instance, while the employment of a corporate firewall consumes the discussion of Control 1, the employment of host-based software firewalls is placed in Control 2.5. While certainly not equal, the presentation of these controls implies a certain requisite investment that may not hold. Given that the presentation is likely more logical than cost-driven, it could be improved by placing some indication on the expected investment required for given controls in an effort to increase compliance. Where a non-technical business owner, upon seeing Control 3 of CE, might assume that the endeavour is complex, some of the efforts are clearly minimal-effort investments that have a great pay-off (for instance, forcing the change of passwords at next login or after a set period is often a simple tick-box exercise). This would differentiate such actions from the more involved actions such as developing a maintaining policies or expensive hardware/software investments, and permit the business owner to prioritize accordingly.
- With respect to the ISBS, from a computational economic standpoint it would be easy to overwhelm the process with a litany of requests for more data; however, it is recognized that this may have adverse effects of lowering participation or adding complexity to the exercise. The primary suggestions that minimize such an outcome are three-fold. First, summary data on total losses and total breaches could be made more central; when making investment decisions, these data are far more useful than worst cost. While the latter may prove a good motivation tool for a company, the former would be far more useful for analysis. Secondly, a finer-grained breakdown of the data into sub-categories would go far to limit the variability that arises. The differences in defence posture and resources by companies that consist of 10, 50 and 250 employees are likely to be significant. Greater differentiation along the accepted definitions of ‘micro’, ‘small’ and ‘medium’ would improve specificity, and result in more interesting and useful analysis. Thirdly, continued and improved publication of the background data—as was accomplished with the 2014 report—will be valuable for research in this area. However, more needs to be done to make this data

usable, and decoding the form—or at least publishing the data key—would go far in this respect. This last measure could overcome the other two, provided that these data were present and made available.

The contribution of this article is three-fold. The first contribution is in the application of standard security economics metrics to conduct empirical analysis of openly available data. To the best of the authors’ knowledge, this is the first treatment of these UK Government policies and data disclosures through the lens of information security economics. The second contribution lies in the methodology employed, which could be replicated by a company seeking to make such an investment. As noted previously, it is not advocated that such an analysis be the sole basis for implementation of a policy such as CE, but as part of the overall decision-making process that a rational company undergoes when seeking to maximize the utility of their investments. Many of the challenges faced in this analysis can be overcome by a company that has solid data in terms of their hardware and software costs, manpower employed towards security, and possibly effectiveness of their current investment. This leaves loss as the primary variable—which the company can set according to their taste. Finally, and perhaps most importantly, it highlights the danger of properly interpreting such an analysis and points to badly needed improvements in data and approach. Insights into theory and practice are a natural result of such exercises, and provide an opportunity to suggest improvements to data publications and analytical frameworks for their utility as decision aides in analysis or as policy tools.

This analysis leaves a number of questions unanswered. To start, the cost, effectiveness and impact of the CE Controls 2 and 3 remain open questions. The lack of fixed hardware or software costs prohibited the ability to provide a reasonable manpower estimation using the methods employed in this work, and quantifiable measures of effectiveness towards the threats identified in the Breaches Survey remains unclear. This is a problem to be tackled by those who engage in analyses that take into account human behaviour and actions. Along a similar route, this analysis does not take into effect other benefits of achieving (or not achieving) a security accreditation such as CE, which on the negative end include loss business and reputation damage and on the positive end include increased business opportunity, goodwill from consumers, or as a signal as part of a ‘sheepskin effect’ [27]. The latter could have complex implications, to include a change in the perceived likelihood of success of an attacker that results in a reduction of attack probability as the attacker moves on to easier targets. Analysis of such deterrence effects are left for future work. A key remaining issue lies in the availability of breach data and the validity of such data when available. While repeating this analysis on other available sources (e.g. [28]) will support general trends, there remains significant difficulty in rectifying data across such sources (which prohibited inclusion of additional datasets in this initial work). Such difficulties underscore the need for datasets that support statistically valid analysis that can be applied generally. To the best of the authors’ knowledge, none of these datasets are produced by policy-makers, although as the trend towards disclosure increases, more expansive analysis in this area should be possible in the near future.

Acknowledgements

The authors would like to thank Emma Osborn for her expertise and insights on the CE scheme, and her beneficial input throughout the development of

this work. The authors would also like to thank the anonymous reviewers for their helpful and constructive comments.

References

- [1]. Rue R, Pleeger SL, Ortiz D. A framework for classifying and comparing models of cyber security investment to support policy and decision-making. In: *Proceedings of the 6th Annual Workshop on the Economics of Information Security (WEIS 2007)*. <http://www.econinfosec.org/archive/weis2007/papers/76.pdf>
- [2]. Garcia A, Horowitz BM. The potential for underinvestment in internet security: implications for regulatory policy. *J Regul Econ* 2007;31:37–55.
- [3]. Anderson R, Barton C, Böhme R, et al. Measuring the cost of cyber-crime. In: *Proceedings of the Workshop of Economics and Information Security (WEIS 2012)*. 2012. http://www.econinfosec.org/archive/weis2012/papers/Anderson_WEIS2012_old.pdf
- [4]. Thomas RC, Antkiewicz M, Florer P, et al. How bad is it? — a branching activity model to estimate the impact of information security breaches. In: *Proceedings of the 12th Workshop on the Economics of Information Security (WEIS 2013)*. 2013. <http://dx.doi.org/10.2139/ssrn.2233075> or <http://www.econinfosec.org/archive/weis2013/papers/ThomasWEIS2013.pdf>
- [5]. Moore T, Anderson R. Economics and internet security: a survey of recent analytical, empirical and behavioral research. Technical Report TR-03-11, Computer Science Group, Harvard University (2011)
- [6]. Department for Business, Innovation & Skills: Information security breaches survey 2015. <https://www.gov.uk/government/publications/information-security-breaches-survey-2015> (16 August 2016, date last accessed)
- [7]. Department for Business, Innovation and Skills: Cyber essentials scheme: overview. <https://www.gov.uk/government/publications/cyber-essentials-scheme-overview> (16 August 2016, date last accessed)
- [8]. Anderson R, Böhme R, Clayton R, et al. Security economics and European policy. In: *Proceedings of the 7th Annual Workshop on the Economics of Information Security (WEIS 2008)*. 2008. <http://www.econinfosec.org/archive/weis2008/papers/MooreSecurity.pdf>
- [9]. Department for Business, Innovation and Skills: Information security breaches survey 2014. <https://www.gov.uk/government/publications/information-security-breaches-survey-2014> (16 August 2016, date last accessed)
- [10]. Department for Business, Innovation and Skills: Information security breaches survey 2013. <https://www.gov.uk/government/publications/information-security-breaches-survey-2013-technical-report> (16 August 2016, date last accessed)
- [11]. Department for Business, Innovation and Skills: Information security breaches survey 2012. <http://www.pwc.co.uk/audit-assurance/publications/uk-information-security-breaches-survey-results-2012.jhtml> (16 August 2016, date last accessed)
- [12]. Florêncio D, Herley C. Sex, lies and cyber-crime surveys. In: *Economics of Information Security and Privacy III*. New York: Springer, 2013, 35–53
- [13]. Moore T. Managing security investment part II. <http://lyle.smu.edu/tylerm/courses/econsec/f12/slides/secinv2-handout.pdf> (6 February 2015, date last accessed)
- [14]. Brecht M, Nowey T. A closer look at information security costs. In: Böhme R (ed.), *The Economics of Information Security and Privacy*. Berlin Heidelberg, Germany: Springer, 2013, 3–24.
- [15]. Gordon LA, Loeb MP. The economics of information security investment. *ACM Trans Informat Syst Security* 2002;5:438–57.
- [16]. Imperva Application Defense Center: Hacker intelligence initiative, monthly trend report #14. http://www.imperva.com/docs/HII_Assessing_the_Effectiveness_of_Antivirus_Solutions.pdf (16 August 2016, date last accessed)
- [17]. Greenberg A. Study finds Microsoft's free antivirus as effective as Symantec's Norton. <http://www.forbes.com/sites/andygreenberg/2010/10/19/study-finds-microsofts-free-antivirus-as-effective-as-symantecs-norton/>
- [18]. Chai B. Firewalls, only 60 per cent effective against malware. <http://www.itproportal.com/2011/04/19/firewalls-only-60-cent-effective-against-malware/> (16 August 2016, date last accessed)
- [19]. Kühner M, Rossow C, Holz T. Paint it black: evaluating the effectiveness of malware blacklists. In: Stavrou, A, Bos, H, Portokalidis, G (eds.), *Proceedings of the 17th International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2014)*. Volume 8688 of Lecture Notes in Computer Science. Cham, Switzerland: Springer, 2014, 1–14.
- [20]. Ms. Smith: Patching Windows is a major time sink for IT departments. <http://www.networkworld.com/article/2229227/microsoft-subnet/patching-windows-is-a-major-time-sink-for-it-departments.html> (16 August 2016, date last accessed)
- [21]. Hulthén R. Communicating the economic value of security investments: value at security risk. In: Johnson, ME, (ed.), *Managing Information Risk and the Economics of Security*. New York, NY: Springer, 2009, 121–40.
- [22]. Verdon D, McGraw G. Risk analysis in software design. *IEEE Security & Privacy* 2004;2:79–84.
- [23]. Soo Hoo KJ. How much is enough: a risk management approach to computer security. Ph.D. Thesis, Stanford University, Stanford, CA, USA, 2000.
- [24]. Xie N, Mead N, Chen P, et al. SQUARE project: cost/benefit analysis framework for information security improvement projects in small companies. Technical Report CMU/SEI-2004-TN-045, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA (2004)
- [25]. Verendel V. Quantified security is a weak hypothesis: a critical survey of results and assumptions. In: *Proceedings of the 2009 Workshop on New Security Paradigms Workshop*, ACM, 2009, 37–50.
- [26]. Herley C. Why do Nigerian scammers say they are from Nigeria? In: *Proceedings of the 11th Annual Workshop on the Economics of Information Security (WEIS 2012)*. 2012. http://www.econinfosec.org/archive/weis2012/papers/Herley_WEIS2012.pdf
- [27]. Spence M. Signaling in retrospect and the informational structure of markets. *Am Econ Rev* 2002;92:434–59.
- [28]. Ponemon Institute LLC: 2014 cost of data breach study: Global analysis. http://www-935.ibm.com/services/multimedia/SEL03027USEN_Poneman_2014_Cost_of_Data_Breach_Study.pdf (16 August 2016, date last accessed)