Policy, Statistics, and Questions: Reflections on UK Cyber Security Disclosures

Chad Heitzenrater[1,2] and Andrew Simpson[2]

[1] U.S. Air Force Research Laboratory Information Directorate
525 Brooks Road, Rome NY 13441
[2] Department of Computer Science, University of Oxford
Wolfson Building, Parks Road, Oxford OX1 3QD, UK

Abstract. Empirical analysis within the field of information security economics is fraught with difficulty, primarily due to a lack of data. Over the last three 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 point out the dependencies. We aim to provide further insight into the method with a view to helping inform a range of stakeholders: policy-makers; those who make decisions with respect to data disclosures; and those looking to policy to help guide their investment in cyber security.

1 Introduction

In information security economics, models rule the day. A scarcity of data [1] and an uncertainty of fundamental properties of attack and defence [2] are commonly cited reasons for this trend, and have led to calls for increased focus in the area of empirical studies [3]. However, mere 'data' in the information technology space is nearly limitless: web searches for 'cyber threats' or 'computer and network defences' will result in pointers to numerous studies, estimates, and opinions. The majority of these will be, at worst, skewed to sell a product; but, even at best, they are likely to be formed on different bases [4]. This results in real-world analysis that is fraught with interpretation, leading some to characterise these issues as stemming from mis-estimation, uncertainty, absence of information, and ambiguity related to disclosure, bias and missing information [5]. As the

Approved for Public Release; Distribution Unlimited: 88ABW-2015-2269, 20150507

literature in this space is increasingly populated with models and theory, empirical analyses are becoming increasingly important as guideposts to continued development of both.

In this paper, we seek to perform such an empirical analysis through the examination of available public data, for the purpose of characterising 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 — will be exploited. Recent and ongoing publications from the UK Department for Business, Innovation & Skills (BIS)[1] 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 April 2014 [6]. The same UK department is also responsible for the development of a scheme seeking 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* [7], outlines five broad areas of compliance (controls); in turn, these areas are broken into between three and seven practices that constitute the minimal level of exhibited capability to meet the broader security objectives. The nature of these sources will be discussed in Section 2.2.

In performing an analysis on these publications, we seek to overcome some of the perils faced by previous authors. While more limited in scope than efforts such as that of Anderson *et al.* [8], this work will focus upon three specific questions related to UK Government efforts with respect to cyber security:

1. *How do the Cyber Essentials controls relate to the known threat?* This is perhaps the most straightforward part of this analysis (albeit the most subjective). The controls called for by Cyber Essentials are examined with respect to their relevance to the related statistics as reported in the Breaches Survey. While this will not require mathematical rigour, these claims are based in the objective reality of computer security literature. This mapping will form the basis for the two subsequent questions, and shed light on the overlap between the stated policy and the known threat.
2. *Is the effort encompassed within the Cyber Essentials 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 Cyber Essentials. 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

[1] See `www.gov.uk/government/publications`.

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 Cyber Essentials practices will seek to investigate the consistency in approach presented. Specifically, an attempt will be made 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 only had one pound / dollar / euro to spend, where should I put it?" This is particularly relevant due to the prescriptive nature of the Cyber Essentials 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 will supply some simple calculations in order to answer some basic questions, with the desired result being the teasing out of some general insight into the current state of cyber security practice and understanding. The purpose in examining these disclosures through the lens of these questions is to seek 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. In doing so, it is the 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 the purpose of this work is not to critique the Cyber Essentials scheme (or the disclosure efforts of the UK Government). Cyber Essentials represents the culmination of an involved process that included some of the leading thinkers in the area of cyber security, and it is not the authors' goal to question its value. Furthermore, examination of the individual practices and an evaluation of their relative merits against the known threat is in no way intended to advocate for changes in Cyber Essentials or the grounds for anything other than the specified implementation. Rather, in the reality of cyber security today, 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 Cyber Essentials and the Information Security Breaches Survey was due to the openness and availability of data, for which the UK Government (and specifically, BIS) deserve great credit (as noted above, data sources in this field are hard to come by). However, the conclusions drawn are not meant to be specific to the Breaches Survey, Cyber Essentials, 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 data sets — were they to be available for analysis. In this light, the reader should be careful to read the following sections not as directive, but as context for the cyber security decision-making process.

The structure of the remainder of the paper is as follows. Section 2 presents the background necessary to understand Cyber Essentials and the Information Security Breaches Survey. Section 3 describes the method employed for the analysis. Section 4 presents each of the questions in turn, and provides a summary of the analysis behind them. Finally, Section 5 concludes by placing some context around the results and identifying avenues for future investigation.

2 Background

The UK is well known and well respected for its openness, transparency, and disclosure regarding the information produced by its government. For example, the `data.gov.uk` initiative has provided insight and innovation, in part, helping to earn the UK the 'most transparent' government in the world ranking.[1] In addition, 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.[2] While receiving only a portion of this investment, it is in this context that the Information Security Breaches Survey and Cyber Essentials have been created by the Department for Business, Innovation & Skills (BIS). We consider each in turn.

2.1 Information Security Breaches Survey

The Information Security Breaches Survey (ISBS, or 'Breaches Survey') is part of an ongoing series of reports comissioned by BIS. The 2014 report [9] follows reports provided in 2013 [10] and 2012 [11], as part of an ongoing effort by BIS to supply the information in a systematic and consistent way in order to enable analysis and discussion. The survey itself has been conducted by PwC for the past two years, in association with Infosecurity Europe and Reed Exhibitions (in 2012 Infosecurity Europe and Reed Exhibitions are stated to have carried out the survey, with the results analysed and report written by PwC).

Each survey report is presented as both an executive summary, highlighting the main findings and notable statistics, and the main technical report itself, which registers a consistent 22 pages in each of the three offerings examined. This 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.

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

[1] See `www.bbc.co.uk/news/uk-30883472`.
[2] See `www.gov.uk/government/policies/keeping-the-uk-safe-in-cyberspace`.

as a comma-separated value (*.csv) file containing the anonymised responses for all the participants.[3] Further enhancing the usability of the data, a web-site was also created to provide an interface to this data, permitting non-programmers the ability to quickly and easily parse the data set and generate graphs of the primary data aspects.[4] This resource enhances the usability and accessibly of the data, spurring further investigation and use, and has proved invaluable in undertaking the analysis described in this paper.

2.2 Cyber Essentials

The Cyber Essentials (CE) scheme, published in April of 2014, is a laudable attempt by the UK Government to "make the UK a safer place to conduct business online" [7]. As such, the stated goal of the document is to provide "requirements for mitigating the most common Internet based threats to cyber security" — which it goes on to identify as phishing and hacking. Perhaps what is most notable about the scheme is the intended reach: in developing CE, BIS has deemed the content "relevant to organisations of all sizes", noting that, while large organisations would be expected to already have some knowledge or experience with the controls, many small- and medium-sized organisations might not have the necessary support or means. To this extent, such organisations are referred to a set of supporting standards and guidance.

The heart of this policy is composed of five technical controls required for "basic technical cyber protection". These are:

1. Boundary firewalls and internet gateways;
2. Secure configuration;
3. Access control;
4. Malware protection; and
5. Patch management.

These controls are then further sub-divided 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 will be at the heart of the investigation undertaken in Section 3, as the cost and benefit (from a utility perspective) of each measure is explored.

3 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

[3] See `www.gov.uk/government/uploads/system/uploads/attachment_data/file/326419/information-security-breaches-survey-2014-technical-report-data.csv/preview`.

[4] See `https://dm.pwc.com/HMG2014BreachesSurvey/`.

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'. This data represents 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" ([9], with similar statements being made in [10] and [11]). As none of the categories investigated state any such caveat, we will bound our analysis on 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.[1]

3.1 Information Security Breaches Survey

The Breaches Survey has followed a standard methodology for disclosure over the course of the last three years of publication (2012–2014). For 2014, the Information Security Breaches Survey data (covering the preceding year) received approximately 1,125 responses, with roughly half (48%) falling into the target category (with similar figures for 2013, and fewer overall respondents but a similar percentage for 2012).

The respondents represent a variety of sectors (20 as of 2014). While the categorizations have changed slightly between 2012 and 2014, the breakdown has remained relatively consistent over these years. Roughly 20% of those surveyed come from each of the 'technology', 'government, health or education', and 'financial' sectors (although the latter drops off to around 12% for 2014), with another 10-20% categorized as 'other'. The remaining 20–30% consists of a combination of 'telecommunications', 'travel, leisure and entertainment', 'utilities, energy & mining', 'manufacturing', 'retail & distribution' and 'property & construction', each representing 1–6% in any given year. The primary shift in demographics over the three years appears to be related the creation of a 'consultancy & professional' grouping in 2014, representing 15.8% of respondents for that year. Based on the relative distribution, it is likely this represents a mix of those previously categorized as 'financial' and/or 'other'. Unfortunately, the data is 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.

Each survey report is broken into numerous parts, providing insight into attitudes, culture and behaviours, as well as trends on the incidences of security breaches. With focus on the latter, information is provided regarding both the frequency of a malicious security incident (60% for small businesses in 2014, down from 64% the previous year) and the instances of serious incidents (50% for small businesses in 2014, up from 23% the previous year). These breaches are then decomposed by type of incident and reported for each of the three years under consideration. Incidences of these attack types for small businesses

[1] See `http://ec.europa.eu/enterprise/policies/sme/facts-figures-analysis/sme-definition/index_en.htm`.

	2014	2013	2012
Infection by viruses or malicious software	45%	41%	40%
Theft or fraud involving computers	10%	16%	12%
Other incidents caused by staff	22%	41%	45%
Attacks by unauthorised outsiders	33%	43%	41%

Table 1. Type of breach (for the percentage of respondents suffering a breach). 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.

	2014	2013	2012
Infection by viruses or malicious software	31%	14%	33%
Attacks by unauthorised outsiders	23%	18%	9%
Theft or fraud involving computers	4%	3%	1%
Infringement of laws/regulation	4%	4%	1%
Theft of computer equipment	0%	4%	5%
Staff misuse of internet/email	12%	12%	15%
System failure or data corruption	7%	23%	34%
Theft/disclosure of confidential info	19%	10%	2%
Other	N/A	12%	N/A

Table 2. Category of worst attacks suffered. Note that the 2014 and 2012 reports list specific information for small businesses only, where the 2013 report combines small and large businesses. Items above the horizontal line are used for the analysis in Section 4.

is provided in Table 1. Unfortunately (for our purposes, at least), this data is 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. This data was 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 £65,000 – £115,000 for the worst incident in the year of reporting. This is further broken out into the nature of the worst breach, mapped into categories mirroring (but not equivalent to) the overall incident types (Table 2). It is clear from the magnitude of these figures that they are likely skewed toward 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.

3.2 Cyber Essentials

In general, the approach taken to the implementation of Cyber Essentials was that of a 'traditional' small business in a modern office setting that does not

include significant investments in, for example, non-traditional computing platforms (e.g. Supervisory Control and Data Acquisition (SCADA)) or embedded systems. This is the primary audience for the Cyber Essentials policy [7]. As a starting point for analysis, the individual controls were examined for their purpose and approach, and grouped into specific technical or procedural means. The corresponding controls listed in parentheses relate to the control descriptions in [7], and the reader is referred to this document for further details. This grouping, which will serve as the basis for mapping to the Breaches Survey, is presented below.

- Firewall (Controls 1.1 and 1.5, plus cost of firewall)
- Firewall policy (Controls 1.2–1.4)
- System administration (Controls 2.1–2.4)
- Personal firewall (Control 2.5)
- Account administration (Controls 3.1–3.7)
- Antivirus (Controls 4.1–4.4)
- Blacklist (Control 4.5)
- Patching (Controls 5.1–5.4)

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

As a policy, Cyber Essentials 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"

Subsequently, this can, under some simple assumptions, be assigned a value in terms of 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.

Information Security Breaches Survey: Categories	**Cyber Essentials: Technical Controls**
Infection by viruses or malicious software	Antivirus Blacklisting Patching
Attacks by unauthorised outsiders	Firewall Firewall policy Personal firewall Patching
Other incidents caused by staff	System administration Account administration
Theft or fraud involving computers	None

Table 3. Mapping of Information Breaches Survey categories of attack to the Cyber Essentials controls

Each of the malicious incident categories from Table 2 can be mapped against the approaches identified in the Cyber Essentials listing above, resulting in Table 3. Since no direct mapping exists, each control from Cyber Essentials is listed by technology.

Some discussion as to the rationale behind each 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, and is additionally supported by 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 gamut of technologies (as is noted).

The threat of incidents caused by staff is defined in [9] as having a wide range: from unauthorised 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 Cyber Essentials. 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 and not part of the Cyber Essentials focus. Since this is a disconnect between the two documents, these two threat classes will not be considered as a part of the analysis of this paper; rather, we will focus on the remaining two classes.

3.3 Resolving ambiguity

As noted in Section 1, the BIS data — despite being a significant step forward in providing information to enable rational choice analysis — lacks in-depth data on the effectiveness of security measures. This leaves residual ambiguity

in the implementation of security technologies, and is the root of the 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 Cyber Essentials 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. For this analysis 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 vs. worst case.* In some measures, there is a definitive (or at least highly likely) worst-case bound. When definitive evidence as to the actual state was not obtainable or varied widely, such a bound was used. The choice of best vs. 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) and upwards of £50, an effort was made to utilise available data on antivirus software to produce a value that represents the real-world distribution of use.
- For all other calculations, where possible a range is presented that represents the expanse of values possible, in order to gauge the trends that result. As noted previously, it is in these trends — and not in some absolute — that the results of any analysis will be generalisable and of any use to others, given the inherent assumptions and approximations that exist in any tractable model of the real world.

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.

Finally, it is noted that the data from the three Information Security Breaches Survey reports employed (2012 [11], 2013 [10], and 2014 [9]) appear to report inconsistent numbers for previous years in the same measures; 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 paper utilises the data for each year as presented in the year reported; that is, the 2014 report will serve for the source of 2014 data, the 2013 report will serve for the 2013 data, etc.

	2012		2013		2014	
	Low	**High**	**Low**	**High**	**Low**	**High**
ALE_{virus}	£3,465	£6,930	£3,724	£6,916	£12,090	£21,390
ALE_{hacker}	£9,450	£18,900	£4,788	£8,892	£8,970	£15,870
$ALE_{virus+hacker}$	£12,915	£25,830	£8,512	£15,808	£21,060	£37,260

Table 4. Worst case Annual Loss Expectancy for the single worst event, 2012–2014. 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.

This will have the effect of generally using higher rates of occurrence (and will complement the 'worst case' approach that has been taken).

4 Analysis

Our focus now shifts to the challenge of examining our data sets.

4.1 How does Cyber Essentials relate to the threat?

To start the analysis, consider the scenario of a small business potentially facing a singular worst loss in 2014. Using the data of Section 3, 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/unauthorised access by outsiders (23%). It isn't 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 [12]:

$$ALE = (p \cdot \lambda)$$

The ALE_0 (loss with no additional security investment) under these assumptions is presented in Table 4, employing the high and low loss event figures for 2012–2014, as determined by $p_{breach} \cdot loss$. 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 reflects the increase in 'serious' attacks (66% from 32%) and resulting costs.

It is worth emphasising 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 Cyber Essentials, such as those identified by Control 3 (account administration), or fall into categories that are not addressed by Cyber Essentials. If the other incident categories identified in Table 2 related to staff were to be fully correctable through the implementation of the remaining Cyber Essentials controls, this still leaves 15% of incidents unaddressed by this scheme, and a residual ALE of £9,750 – £17,250 per enterprise in 2014. The implications of this will be further explored in Section 5.

4.2 Is the effort encompassed within the Cyber Essentials practices 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 unauthorised outsider as 3 and 5 respectively, with a median of 6 total incidents overall. Normalising 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 2012–2014 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, seeming to motivate investment in cyber defence.

Next, the loss expectancy will be 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 Cyber Essentials 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 license 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:

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

Using this model, the Expected Net Benefit of Information Security (ENBIS) will be 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):

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

In general, one should invest in security at the point that $ENBIS > 0$, 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

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

Given the bounds placed on the value of ALE_0 and under the assumption that any current security investment is uncounted toward 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 Cyber Essentials controls will result in a residual probability of 99%; as anyone involved in cyber security is surely aware, 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 5 lists estimated costs of common cyber security controls

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	N/A

Table 5. Fixed cost estimates for material investments relative to Cyber Essentials.

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. Exact amounts for these are, of course, dependent on the precise implementation; most importantly to this analysis is that they are relative to reality so that the various aspects of resource investment play a role.

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 timeframe of effort by a smaller organisation), the latter is much more complex. Costs included here 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 mostly 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 utilised in [13]) was chosen, as it permits estimates of cost based upon the distribution of costs between software (29%), hardware (21%), manpower (40%) and outsourced (10%) costs. The limitations of this model are well documented [13] 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 5 works out as £1,505.22, or roughly 14 working days of

time at the going rate for IT support personnel in the UK (£26,597 per year[1]). 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 vs. expenditure ratio for maximum security investment [14]. 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 slides 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 organisation at the upper end of the scale is at a loss in each of the three years. Conversely, the hypothetical singe-system organisation 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 £2,000 making this a good investment. While this is admittedly using the lower cost of the firewall in Table 5, 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 inher-

[1] Based on information from `www.payscale.com/research/UK/Job=Information_Technology_(IT)_Support_Specialist/Salary`, accessed in February 2015.

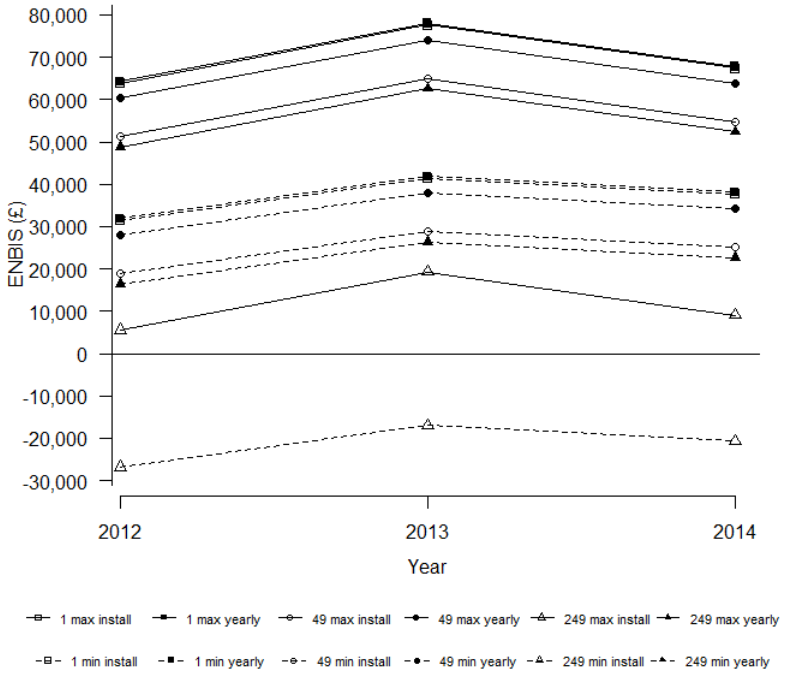

Fig. 1. ENBIS against loss estimates per year. The lines show upper (solid lines) and lower (dashed lines) bounds, using Gartner assumptions of manpower investment.

ently 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 organisation 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%; that is, 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 Cyber Essentials. Here, the simpler question of overall effectiveness will be considered. Returning to the high and low estimates of

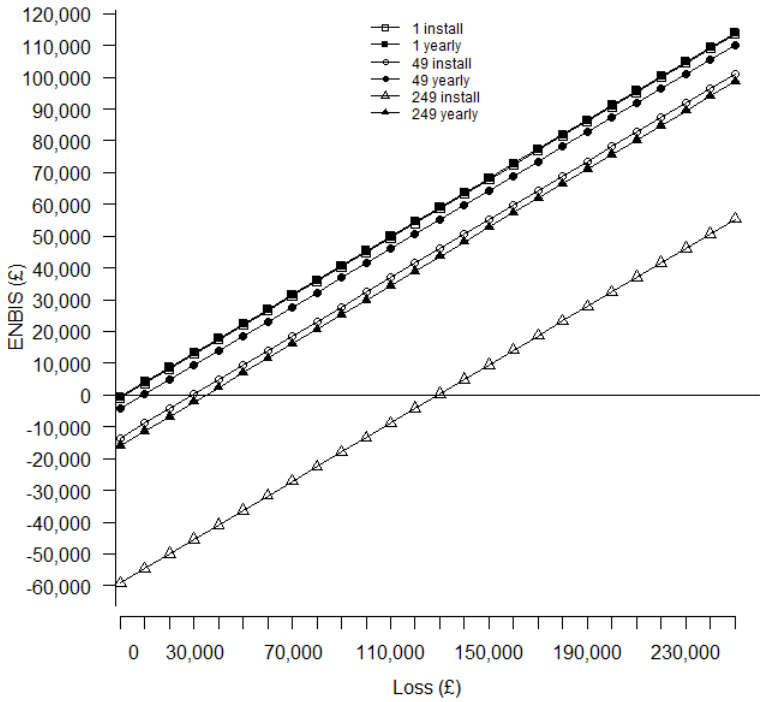

Fig. 2. ENBIS with varied loss. The lines show yearly and install expenditures, using Gartner assumptions of manpower investment.

multiple breach loss, the data for 2014 will again be examined against variability in overall security effectiveness from 50% to 99% effectiveness. This is shown in Figure 3.

It is evident that, as the effectiveness of the controls being invested in 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, under these particular assumptions, 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 a closer at realistic expectations of effectiveness. We discuss this further in Section 4.3.

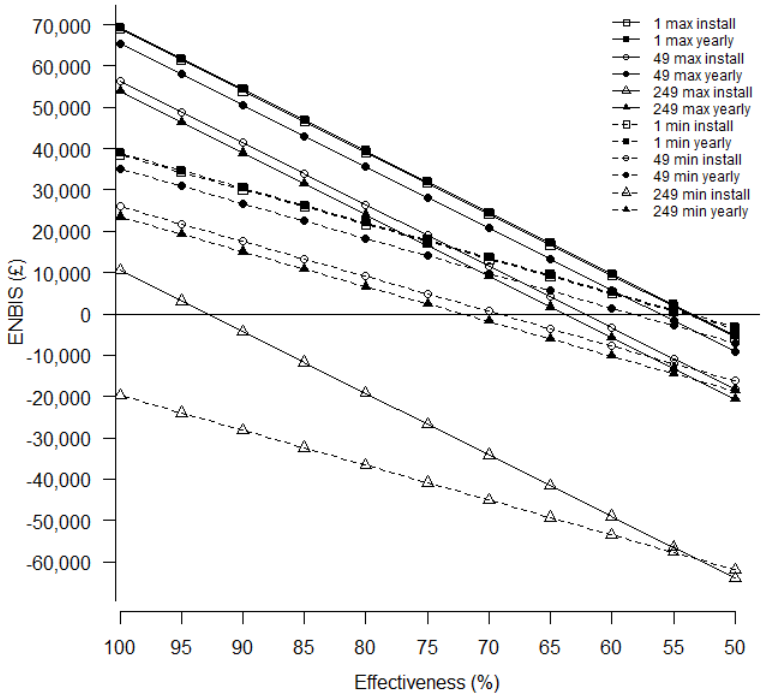

Fig. 3. ENBIS with varied effectiveness of security controls. The lines show upper (solid) and lower (dashed) loss assumptions for various company sizes (1, 49, and 249), using Gartner assumptions of manpower investment.

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 to 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 Cyber Essentials fail, it is worth reiterating that none of these 'views' on the data alone provide a realistic or definitive commentary on the Cyber Essentials 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. It is the 'push-and-pull' between these considerations that is being highlighted here, and the fact that, as discussed in Section 5, the actual decision to implement cyber defences likely relies on much more than an

Control	Effectiveness	Reported ranges / notes
Antivirus	75%	Reported ranges of 5% [15] to 75% [16]
Firewall	60%	Study cited 60% 'out-of-the-box' and
Firewall policy	80%	80% only with skilled administration [17]
Blacklists	73.5%	Lowest coverage for a given malware class by all major AV vendors in [18]

Table 6. Reported effectiveness for various cyber controls.

economic analysis, and by necessity must take into effect regulatory, reputational and ethical considerations.

4.3 How should the threat inform the implementation of Cyber Essentials?

Recalling the controls specified in Cyber Essentials (or rather the corresponding technologies identified) and the mapping provided in Table 3, the question of effectiveness can be further examined. Determining effectiveness of a specific measure can be a difficult exercise; much depends on the specific configuration and deployment scenario, and, to deploy a well-worn cliché, 'the devil is in the details'. 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 of something nearing 100%. However, research, literature and trade publications in the area seem to suggest that this is almost never accomplished, and the bigger the organisation (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 [19]. 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.

The overall cost of a given control will be modelled in the same manner as previously presented. This time, the comparison basis will utilise the Net Present Value (NPV) of the technology in question. NPV seeks to aggregate the benefit to be had over multiple future periods into a singular value, and takes into consideration both one-off and recurring costs [12]:

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

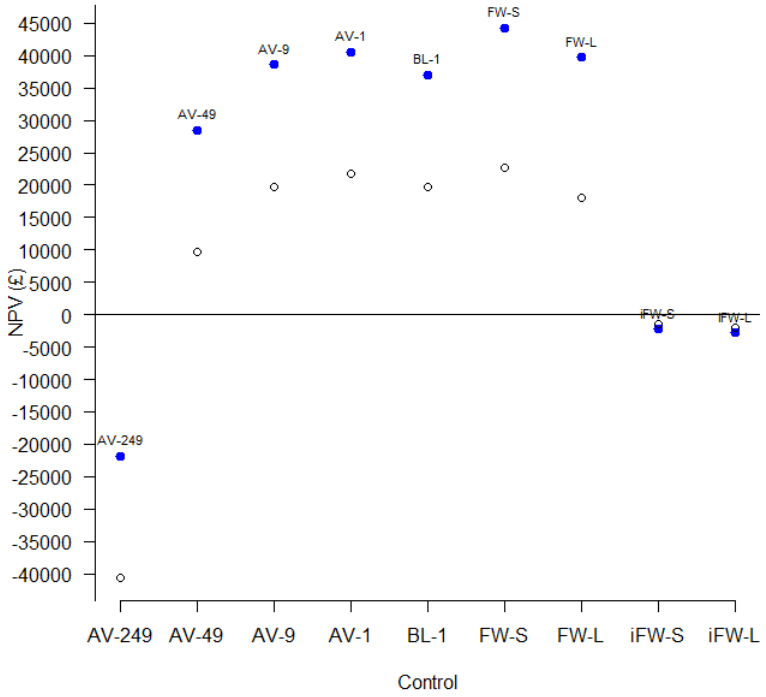

Fig. 4. Net Present Value (NPV) of security controls as calculated using the assumptions contained in this analysis. Solid dots are calculated using the high-bound loss estimate, and circles are calculated using the lower loss estimate.

Employing the same estimates as used in the previous subsection, the NPV for each of the technical controls (except for patching) can be calculated. A rate of return of 5% was used as an ad-hoc estimate, as is common in practice for such calculations [12]. These are plotted in Figure 4 using the data from 2012–2014.

The controls investigated are as follows:

– Host-based antivirus and blacklists for organisations 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 5 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 specialisation in policy or rule set; as 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 organisational 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 setup 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 the point that was made earlier: 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 is going to 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 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 organisation 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 trends toward zero. Further investigation on this topic is left as future work.

One way that 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 organisation, 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 organisation 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. Of course, the compounding of estimates throughout this analysis has impact on this conclusion, as would 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 the concluding section.

5 Conclusions and future work

It is worth *strongly* reiterating that the majority of the analysis presented in this paper was built upon hypothetical 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 Cyber Essentials controls across the variation of organisations. 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.* [2] and by other authors.

- In Question 1, the probability of loss was treated 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 [20]). 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 generalisation, 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 loss 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; that is, 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 utilised 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 who have previously used such methods for measuring cybercrime [4], performing risk analysis in software design [21], conducting quantitative analytics for managing computer security risk [22], and, in close synergy with this work, as inputs to a decision framework for security improvement projects in small companies [23]. 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' [24], 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 contribution of this paper is twofold. The first contribution is the employment of standard security economics metrics to a real data set, for the purpose of exercising the foundations upon which much of the field is based. 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. As such, it provides insights into the state of the field with regards to the ability to utilise such disclosures; it also provides an opportunity to suggest improvements to such publications for their utility as decision aides in analysis or as policy tools. It also points to further questions, such as the applicability of Gordon-Loeb in specific scenarios. The second contribution is 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 Cyber Essentials, but as part of the overall decision-making process that a rational company undergoes when seeking to maximise 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.

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

1. *How does Cyber Essentials relate to the 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 Cyber Essentials scheme — the analysis presented also shows that there remain significant threat gaps which are unaddressed or left to policy. 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 Cyber Essentials practices requisite to the threat?* Implementation of the various aspects of Cyber Essentials lends itself to a wide range of possible costs. Unsurprisingly, the amount of time invested and the number of machines within the organisation are the largest contributing factors in the ENBIS calculations; to paraphrase Herley [25], 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 Cyber Essentials 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 which 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.

In addition to seeking answers to these questions, this report encompasses experiences with Cyber Essentials and the Information Security Breaches Survey 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 Cyber Essentials, 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 Cyber Essentials, might assume that the endeavour is complex, some of the efforts are clearly minimal-effort endeavours that have a great pay-off (for instance, forcing the change of passwords at next login / 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 prioritise accordingly.
- With respect to the Information Security Breaches Survey, from a computational economic standpoint it would be easy to overwhelm the process with a litany of requests for more data; however, it is recognised that this may have adverse effects of lowering participation or adding complexity to the exercise. The primary suggestions that minimise such an outcome are threefold. First, summary data on total losses and total breaches could be made more central; when making investment decisions, this data is 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. Next, 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 is 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. Finally, 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 this data were present and made available.

The analysis presented leaves a number of questions unanswered. To start, the cost, effectiveness and impact of the Cyber Essentials 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 toward 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 Cyber Essentials, 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' [26]. 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. Finally, a number of data sets are starting to be published that may provide additional utility in analysis (e.g., [27]). While being less specific, higher-level treatments may hold key statistics that would be complementary to the Breaches Survey. However, there remains significant difficulty in rectifying data across such sources which prohibited inclusion in this initial work. To the best of the authors' knowledge, none of these data sets are produced by policy-makers, although as the trend towards disclosure increases, the future looks bright for more expansive analysis in this area.

Acknowledgements

The authors would like to thank Emma Osborn for her expertise and insights on the Cyber Essentials scheme, and her beneficial input throughout the development of this work. The authors would also like to thank the anonymous reviewers for their constructive comments.

References

1. Garcia, A., Horowitz, B.M.: The potential for underinvestment in internet security: Implications for regulatory policy. Journal of Regulatory Economics **31**(1) (2007) 37–55
2. Rue, R., Pleeger, S.L., 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). (2007)
3. 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)

4. Anderson, R., Barton, C., Böhme, R., Clayton, R., van Eeten, M., Levi, M., Moore, T., Savage, S.: Measuring the cost of cybercrime. In: Proceedings of the Workshop of Economics and Information Security (WEIS 2012). (2012)
5. Thomas, R.C., Antkiewicz, M., Florer, P., Widup, S., Woodyard, M.: 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)
6. Department for Business, Innovation & Skills: Keeping the UK safe in cyber space. https://www.gov.uk/government/publications/information-security-breaches-survey-2014 (2014)
7. Department for Business, Innovation and Skills: Cyber essentials scheme: overview. https://www.gov.uk/government/publications/cyber-essentials-scheme-overview (2014)
8. Anderson, R. Böhme, R., Clayton, R., Moore, T.: Security economics and European policy. In: Proceedings of the 7th Annual Workshop on the Economics of Information Security (WEIS 2008). (2008)
9. Department for Business, Innovation and Skills: Information security breaches survey 2014. https://www.gov.uk/government/publications/information-security-breaches-survey-2014 (2014)
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 (2013)
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 (2012)
12. Moore, T.: Managing security investment part II. http://lyle.smu.edu/~tylerm/courses/econsec/f12/slides/secinv2-handout.pdf (2012)
13. Brecht, M., Nowey, T.: A closer look at information security costs. In Böhme, R., ed.: The Economics of Information Security and Privacy. Springer (2013) 3–24
14. Gordon, L.A., Loeb, M.P.: The economics of information security investment. ACM Transactions on Information and Systems Security **5**(4) (2002) 438–457
15. Imperva Application Defense Center: Hacker intelligence initiative, monthly trend report #14. http://www.imperva.com/docs/HII_Assessing_the_Effectiveness_of_Antivirus_Solutions.pdf (2012)
16. 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/ (2010)
17. Chai, B.: Firewalls, only 60 per cent effective against malware. http://www.itproportal.com/2011/04/19/firewalls-only-60-cent-effective-against-malware/ (2011)
18. Kührer, 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., Springer (2014) 1–14
19. 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 (2011)

20. Hulthén, R.: Communicating the economic value of security investments: Value at security risk. In Johnson, M.E., ed.: Managing Information Risk and the Economics of Security. Springer (2009) 121–140
21. Verdon, D., McGraw, G.: Risk analysis in software design. IEEE Security & Privacy **2**(4) (2004) 79–84
22. Soo Hoo, K.J.: How Much is Enough: A Risk Management Approach to Computer Security. PhD thesis, Stanford, CA, USA (2000)
23. Xie, N., Mead, N., Chen, P., Dean, M., Lopez, L., Ojoko-Adams, D., Osman, H.: 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)
24. 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
25. 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)
26. Spence, M.: Signaling in retrospect and the informational structure of markets. American Economic Review **92**(3) (2002) 434–459
27. 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` (2014)