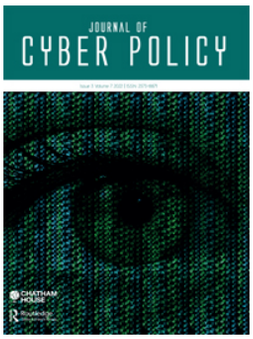

Journal of Cyber Policy

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/rcyb20

Patchy incentives: using law to encourage effective vulnerability response

Andrew Cormack & Éireann Leverett

To cite this article: Andrew Cormack & Éireann Leverett (12 Dec 2023): Patchy incentives: using law to encourage effective vulnerability response, Journal of Cyber Policy, DOI: 10.1080/23738871.2023.2284233

To link to this article: https://doi.org/10.1080/23738871.2023.2284233

Published online: 12 Dec 2023.

JOURNAL OF CYBER POLICY
https://doi.org/10.1080/23738871.2023.2284233

Routledge
Taylor & Francis Group

Patchy incentives: using law to encourage effective vulnerability response

Andrew Cormack[a] and Éireann Leverett [b]

[a]Cardiff, Wales; [b]Cambridge, England

ABSTRACT

Data breach reports suggest that managing patches is hard: too many major incidents are caused by well-known software vulnerabilities with available fixes. Legal sanctions – from mandates to liability – apparently have limited effect. This paper discusses how an effective vulnerability response process can help software users allocate their remediation effort to minimise overall risk and disruption. We analyse laws and regulations on liability, product quality and patching mandates to see why they fail to promote good practice. Recent cases under privacy laws highlight features that make risk-based patching a better basis for system managers, executives and regulators to agree a common approach to effective vulnerability response.

ARTICLE HISTORY
Received 25 April 2023
Revised 20 August 2023
Accepted 1 September 2023

KEYWORDS
Attack surface management; law; regulation; vulnerabilities

Introduction: vulnerabilities and fixes

Modern software is too complex to be perfect: implementation flaws and design choices with unexpected consequences are inevitable. Some of these flaws prevent a system or feature working as intended. Others – usually harder to find – affect its security. ENISA defines 'a weakness an adversary could take advantage of to compromise the confidentiality, availability, or integrity of a resource' to be a 'vulnerability' (ENISA n.d).

Organisations writing software – here referred to as 'vendors', whether or not the software is actually sold – should have processes to fix vulnerabilities and issue updates. Around fifty vulnerabilities are announced per day (Cyentia Institute n.d.a), so major vendors will need to prioritise work and effort among several problems. Prioritisation (within the vendor) will usually depend on the impact if the vulnerability is exploited and the deployed footprint of the affected software. The means of discovery may be a factor: external reporters often expect a fix within sixty days under Coordinated Vulnerability Disclosure (NCSC n.d.); internally discovered vulnerabilities may use different processes and longer timescales. A 'zero-day' vulnerability, exploited by attackers before a fix is available (Fruhlinger 2019), may trigger an accelerated process. All processes should develop a fix or mitigation: test it in commonly-used configurations and make it available to users. Many vendors schedule regular patch releases to help customers' planning: the second Tuesday of each month is Microsoft's 'Patch Tuesday' (Wikipedia n.d.). Some patches may appear at any time: from vendors without a regular cycle, or where a fix is too urgent to delay (Miranda et al. 2021).

CONTACT Éireann Leverett eleverett@concinnity-risks.com

This paper considers the next link in the chain: organisations that receive vendors' vulnerability notices. They need a policy and process for vulnerability response, 'deal[ing] with vulnerabilities affecting the software the organisation uses but for which it is not the code maintainer' (Cyber Safety Review Board 2023). We first consider why 'patch everything immediately' is neither efficient nor feasible. Instead, it is better to decide when and what to patch, and what tools can inform prioritisation decisions. Although the need for prioritisation was identified in 2007 (Miura-Ko and Bambos 2007), economic inefficiencies pointed out in 2009 (Romanosky and Acquisti 2009), and 'more effective (and statistically significant) strategies' proposed in 2014 (Allodi and Massacci 2014), this risk-based approach to vulnerability response is not widely adopted, often resulting in 'disproportionate mitigations' (Allodi, Massacci, and Williams 2021). Our principal focus is whether regulation could provide better incentives and adopt the new paradigm. Examining mandates, laws that regulate outcomes, behaviour and harms, we conclude that behaviour-based legislation provides an excellent framework for encouraging risk-based vulnerability response. Analysing UK data protection cases involving missing patches highlights four helpful characteristics in behaviour-based laws. We hope that linking vulnerability response to these laws can provide common ground for technologists, executives, compliance departments and regulators to discuss.

Vulnerability response: the challenge for patch consumers

Organisations that use third-party software and receive vulnerability notifications must prioritise. It might seem that they should install every patch immediately, but patching can cause disruption and may introduce new risks. Taking time to plan and test reduces undesirable results. Effective vulnerability response cannot eliminate these impacts, but it can, and should, balance the risk of an unpatched system being attacked against the certain disruption to organisation, users and customers. The risk of an insufficiently tested patch creating problems is real. Vulnerability response must identify the appropriate time and preconditions for patching each individual system.

Perhaps surprisingly, risk-based vulnerability response will usually conclude that, at any given time, some patches should not be installed. Patching will either cause more disruption than the current risk justifies or divert effort from more urgent vulnerabilities. The median 'remediation effort' – irrespective of organisation size – enables planning, testing and installing fixes for 15 per cent of known vulnerabilities in the IT estate. This may be enough to maintain, or even improve, the security of infrastructure, data and customers. Most vulnerabilities are never investigated by attackers or developed into exploit code. The process for prioritising vulnerabilities is critical: for the same level of effort, the choice of prioritisation algorithm can make a twenty-fold difference to the organisation's exposure (Cyentia Institute n.d.a, n.d.b).

This section first considers factors influencing the timing of patching, then introduces tools to inform effective risk-based vulnerability response programmes, then reviews how a programme can prioritise patches.

When to patch

Vulnerabilities create risks to the confidentiality, integrity or availability of the data accessed, stored or processed by computer systems. Vulnerability announcements help

software users remediate these risks: they also draw the attention of malicious actors. Though vendors try to avoid this, they inevitably give clues about how the vulnerability might be exploited (Shahzad, Shafiq, and Liu 2012). Announcements increase the risk of exploitation, but in most cases, only gradually. Writing code to exploit a vulnerability is a complex technical task, and unless an attacker is targeting a specific organisation, delivers little benefit while earlier vulnerabilities of similar severity remain widespread (Allodi, Massacci, and Williams 2021). Of thirteen vulnerabilities in 2021 that became widespread threats after public announcement (i.e. excluding zero-days), only two were exploited within a week, the rest took from seven to ninety-three days (Condon et al. 2021).

Organisations can use this period of slowly increasing exploitation risk wisely to reduce deployment risks. Although responsible vendors test patches and mitigations against common software configurations, individual installations may encounter unexpected interactions with rarer configurations. These may crash the patched system (Condon et al. 2021), but some create new – or re-create old – vulnerabilities in either that system or others it interacts with (Ormandy 2021). This is a particular concern when patching 'technology cornerstones', where an incompatibility can impact the organisation's entire IT operation (Condon et al. 2021). Organisations should first test patches or mitigations on a copy of their live system, especially where unusual configurations have not been covered by vendor testing (Leverett, Coburn, and Woo 2019). Pre-testing also familiarises staff with the patching process: identifying any required information or – particularly for mitigations – side-effects, and local specificities.

Even with preparation, installing a patch or mitigation often involves disruptions. Many patches require rebooting, causing interruption to services. Miranda et al. (2021) note that for components of industrial systems, even brief outages require significant pre-planning. Mitigations – including disabling or removing vulnerable functions, blocking or filtering network or application traffic – may make affected functions unavailable until full remediation becomes possible (Cyber Safety Review Board 2023). Users may need to take individual action – from re-authenticating to upgrading their client software (CISA 2016). These service and user impacts may be unavoidable but can be reduced by scheduling maintenance windows for patching. This means communicating to users times to avoid critical changes: these may match regular vulnerability announcement cycles (Leverett, Coburn, and Woo 2019), but out-of-cycle, highly urgent patches must also be handled similarly. Organisations often have critical periods or events – sports events or speeches or exam results or tax deadlines – when disruption is particularly costly. One online bookmaker served 150,000 customers per minute during the 2021 Grand National, with the industry as a whole taking over £100 million from 13 million people (Muvija 2021). For such companies, 'disruptive maintenance is at best a negotiation, and at worst an instant deal breaker' (Condon et al. 2021). The commercial benefit of re-scheduling patching to avoid these times may well justify the increased risk of a shorter testing period or a longer period of exposure. Deciding when to deploy a patch requires balancing these diverse risks.

Tools for assessing patch risk

Every organisation should have processes to track vulnerability announcements and fixes in the software it uses, to prioritise and take appropriate actions (Miranda et al. 2021). This process must consider context, including critical systems and times of use, existing local

protection measures, and the significance of different types of security impacts. Several catalogues, metrics and statistical tools can contribute to vulnerability responses that are locally specific, but also systematic and evidence-based.

MITRE's Common Vulnerabilities and Exposures (CVE) catalogue (CVE n.d.) establishes a unique reference number for each disclosed vulnerability since 1999. The linked US National Vulnerability Database (NVD) (US Government n.d.) provides additional information such as affected software versions. Both are structured and machine-readable: their reference numbers (e.g. CVE-2014-0160) let both human and automated patching systems identify vulnerabilities and patches or alternative mitigations.

The Common Vulnerability Scoring System (CVSS) (Common Vulnerability Scoring System n.d.), established in 1999, captures factors that make a vulnerability impactful. These include intrinsic ('Base') features – what access is required for exploitation, how complex exploitation would be, what authentication and user interactions are needed, and the impacts on confidentiality, integrity and availability. It also records features that vary in time ('Temporal') – availability of exploits and remedies, and the confidence level of the report. CVSS scores are widely used in vendor announcements and by NVD announcements. A set of 'Environmental' features let individual organisations adjust the CVSS score for local factors affecting the severity of exploitation – the type of damage likely in their ICT estate, the prevalence of vulnerable systems and the relative importance of confidentiality, integrity and availability.

In November 2021, the US Government's Cybersecurity and Infrastructure Security Agency (CISA) announced a new catalogue of currently exploited vulnerabilities for which mitigations are available (CISA 2021b). Initially, this contained 290 vulnerabilities (2017–2021): nearly *two orders of magnitude fewer* than the CVE list from the same period; by June 2022 it had grown to 770 (CISA 2022).

Whereas CISA highlights current exploits, the Exploit Prediction Scoring System (EPSS) (EPSS n.d.) predicts which vulnerabilities will be exploited in the near future (in EPSS v2 within the next 30 days). This statistical model uses correlations from historic data to support sound risk-based prioritizations even before exploits occur in the wild (Jacobs et al. 2021). More broadly, inspired by EPSS and other work, Leverett, Rhode, and Wedgbury (2022) identified statistical patterns in the rate of vulnerability announcements that could help organisations plan their allocation of resources to vulnerability response up to a year ahead. That work has since expanded into an open source proof of concept and a yearly conference (Leverett, Manion, and Rhode n.d.). The success of these and other forecasting efforts (Sarabi et al. 2016) indicates the existence of patterns that risk-based approaches can use to improve the security of systems before they are compromised.

What to patch

These catalogues make it clear that a risk-based approach to vulnerability response is both essential and appropriate. The rate of discovery is increasing: 18,000 new CVEs in 2020; 20,000 in 2021; 25,000 in 2022; more than 50 daily (CVE 2023). It 'isn't even remotely feasible' for most organisations to assess, find and remediate everything (Cyentia Institute n.d.b). Allodi and Massacci (2014) confirm that the 'timely patching of all vulnerabilities is infeasible.' Figure 1 – CVEs published each month – shows the scale and intractability of the problem: even ignoring CVSS Base severities below 'High' (7.0) and 'Medium' (4.0)

leaves a monthly workload of between 300 and 1000 CVEs to consider. Worse, simplifying by impact may omit vulnerabilities where organisation-specific threats or dependencies increase the local risk or conversely, waste attention on others already covered by existing mitigations (Jacobs et al. 2021).

Of course, most organisations will not use all the software that those vulnerabilities are in; but larger organisations will use more of them. Even after filtering vulnerabilities for just the software or firmware they use (filtering by the asset register), they still struggle. 'Organisations cannot fix all the vulnerabilities across all their assets all of the time.' Perhaps surprisingly, 15 per cent is enough, if efficiently used, to maintain or reduce the level of exploitation risk. Large-scale internet measurements suggest only 16 per cent of known vulnerabilities are even investigated (Cyentia Institute n.d.a) and less than 5 per cent have proof-of-concept or exploit tools, and even these may never be used (Jacobs et al. 2021). Allodi and Massacci (2014) confirm that 'the vast majority of attacks recorded in the wild are driven by only a small fraction of known vulnerabilities.' It seems the rate of vulnerability discovery is as daunting to attackers as to defenders. Thus we expect this percentage to fall over time, as the number of vulnerabilities far outstrips their usage.

This means defenders can and should choose which vulnerabilities to patch next. Poor choices, according to Jacobs et al. (2021), 'waste countless hours and resources remediating vulnerabilities that could be delayed, or conversely delay remediation of critical vulnerabilities.' Unnecessary downtime causes financial loss. Instead, each organisation should prioritise vulnerabilities that represent both a high likelihood of exploitation and a high impact to the organisation. Likelihood and impact change as new vulnerabilities are discovered, exploit tools developed (Cyentia Institute n.d.a) and defences updated, so prioritisation must constantly adapt to fresh information (Miranda et al. 2021).

Even simple strategies can produce dramatic improvements. CVSS Base score – a 'suboptimal' strategy that ignores temporal and local factors (Jacobs et al. 2021) – can halve the vulnerability exposure compared to applying patches as they are released. Prioritising

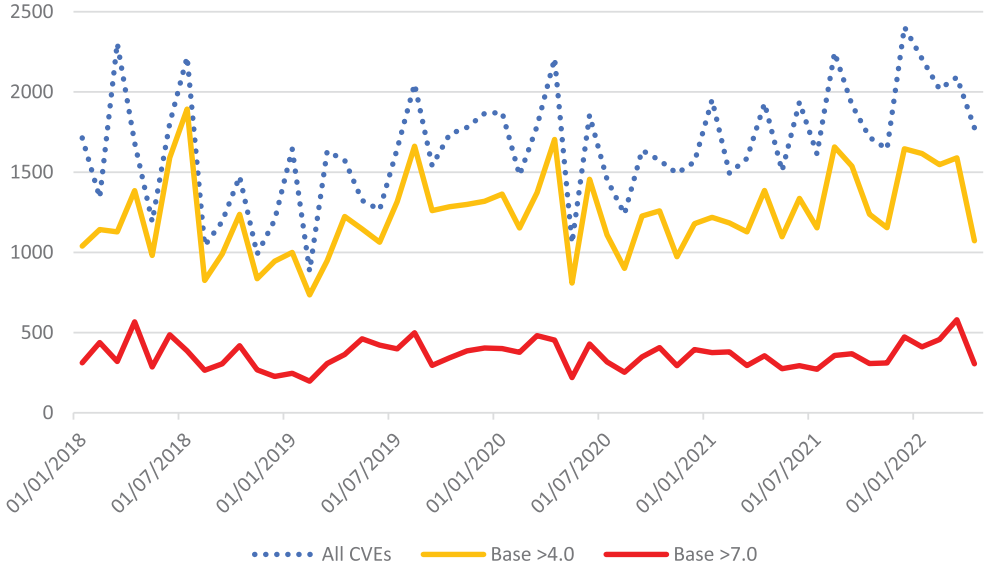

Figure 1. Number of CVE notices per month, by CVSS Base score.

vulnerabilities with known exploits can be 'more effective than quadrupling capacity to patch what CVSS deems to be critical', while using EPSS to predict future exploitation risk could further halve exposure. Incorporating local knowledge of exploitability and attack surface can do even better (Cyentia Institute n.d.a). Effective prioritisation could deliver a target level of exposure with four times less effort (Jacobs et al. 2021), or help those with limited resources improve their security position twenty-fold.

Vulnerability response must not be 'a mindless, endless loop of finding and fixing. 'Organisations have a great deal of control over their attack surface through the strategies and capabilities they employ' (Cyentia Institute n.d.a). Yet many immature vulnerability management and patch processes are basically binary: is this particular vulnerability important enough to patch? Then they simply move on to the next vulnerability, failing to record the decision and revisit the prioritisation decision in the future. By contrast, we think a much better way to approach it is to filter by the asset register and assign a timeline to patch vulnerabilities according to their risk factors. An abstract diagram can be found below, though the exact timelines and volumes will be specific to each organisation. The important takeaway is that changes in prioritisation over time move a vulnerability from one priority queue to another (Figure 2).

Crucially, the last priority queue is the patch time of last resort … and its timing is different for different organisations. Consider an industrial system in which devices have high uptime requirements and upgrades may be performed on a yearly cycle versus a virtual machine or software-defined network in the cloud where old VMs are retired within minutes and new ones are created with updated operating systems or libraries. These two organisations will have wildly different approaches to prioritisation because of this baseline lifecycle. In short; if your median lifecycles are short, your vulnerability management programme may be lighter and faster too. It is the relationship of the distribution of lifecycles to the timeline of exploitation risk that matters.

Promoting risk-based vulnerability response

Risk-based vulnerability response can improve the security of organisations' systems, data and customers, using resources and skills that any organisation should already possess. No approach can eliminate all possibility of breaches: some zero-day attacks occur without any warning and others too soon after a warning for organisations to reasonably respond, but equally, unpatched systems should not be an 'easy meal ticket' for attackers (Roncevich 2018). In short, we can do much more risk reduction with the same workload and talent.

Nevertheless, risk-based responses are not commonly adopted and unpatched high-priority vulnerabilities still enable damaging breaches long after fixes were available. We therefore consider whether laws could encourage risk-based approaches that allocate remediation effort optimally. It would certainly help regulators and vulnerability managers develop and promote best practice. All this would benefit the public good. Laws can create incentives both as a 'sword' – punishing those who do not act as desired – and a 'shield' – rewarding those who do. These sticks and carrots are relevant to vulnerability response where even perfect behaviour cannot prevent occasional bad outcomes.

In 2009, Romanosky and Acquisti considered whether US data breach laws created economic incentives 'to protect personal information, and decrease the harm … as a

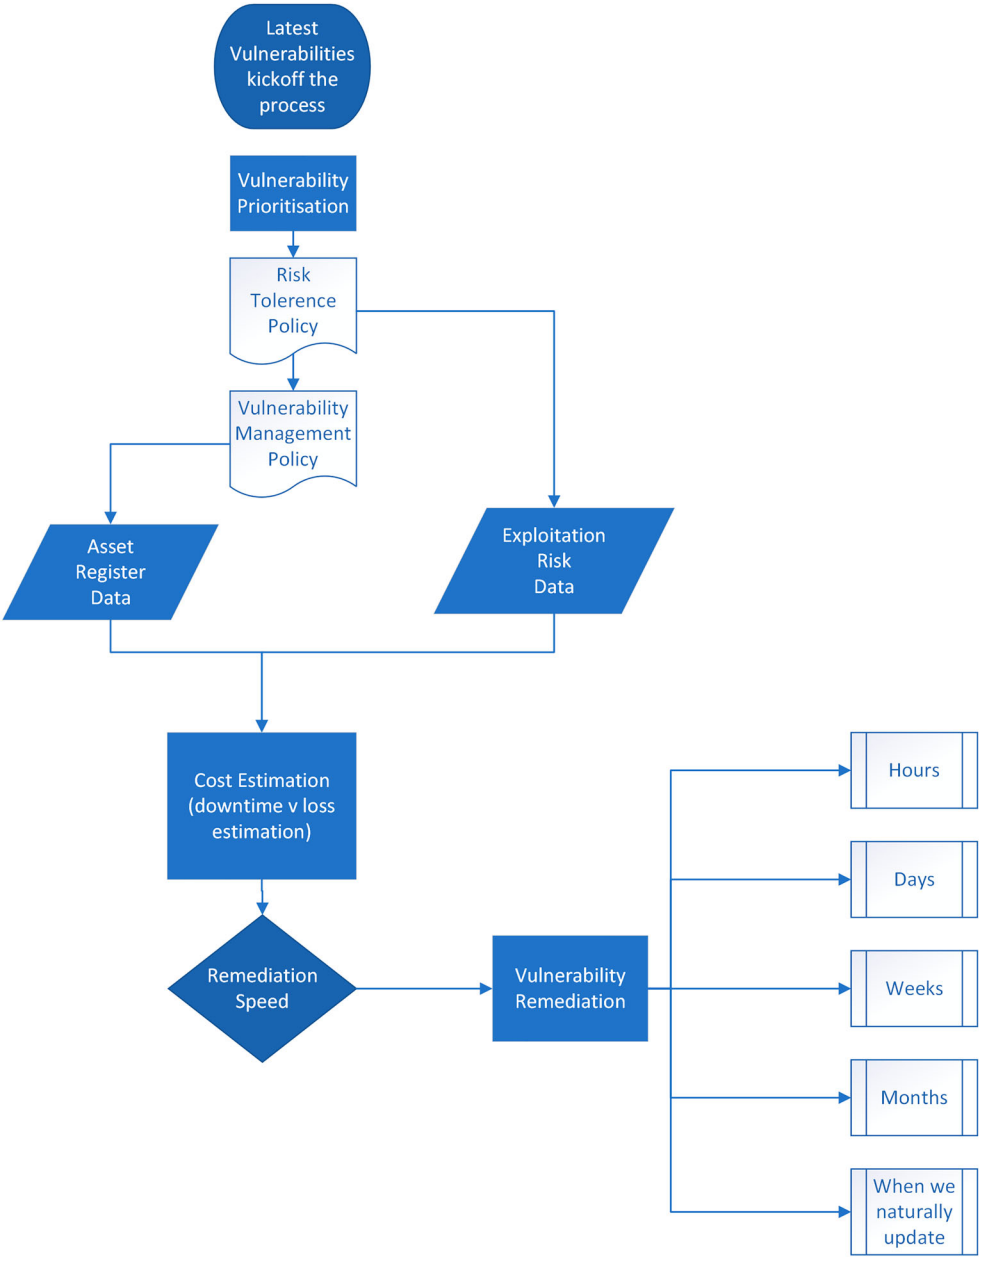

Figure 2. An idealised patching process example.

result of breaches of this information', finding only 'lukewarm results'. They concluded *ex ante* regulation 'is efficient only for a single set of firms causing the average amount of harm', *ex post* liability 'only when suits are always initiated and firms always pay for their harm' and breach notification '[only] when firms bear all of the consumer harm and will reduce total social loss when consumers take action to reduce their harm'. In practice, the diversity of regulated firms and the low level of enforcement, the low

likelihood of successful claims for harm, and the externalities in privacy breach costs dilute these incentives.

In the rest of this paper we narrow Romanosky and Acquisti's focus from protecting personal information to the behaviours that constitute risk-based vulnerability management. Do, or could, laws create incentives that encourage these? Over the past decade their 'three policy approaches' have been supplemented by laws creating specific mandates to patch vulnerabilities. Their categories are still helpful, though we have reframed them based on whether the laws create requirements on behaviour or primarily on *ex ante* product safety (UK Government 2019) that create requirements on the outcomes of behaviour. Those are principally *ex post* 'liability', that assign responsibility for harm. Since different jurisdictions have favoured different approaches, our illustrations of the capabilities and limitations of each category necessarily come from different parts of the world: mandates from US and UK Government internal policies; behavioural requirements from a wide range of data protection laws (European Union 2016); specific UK cases on vulnerability management; outcome requirements from EU product regulation; and harm responsibility from English and US legal systems.

Laws that mandate

Some legal requirements simply specify required vulnerability response practices against CVSS and specific examples of that temporal and local interpretation inflexibility follow.

The US Government's Binding Operational Directive (BOD) 19–02 insists that 'Critical' vulnerabilities in internet-facing systems (ratings from CVSS Base or the Nessus vulnerability scanner) be remediated within 15 days and 'High' severity ones within 30 (CISA 2019). The UK's Cyber Essentials – mandatory for anyone bidding for government contracts involving personal and sensitive information (NCSC 2014) – requires all critical and high-risk patches to be implemented within 14 days. This is considered 'a reasonable period' whereas 'a shorter period may not be practical' (NCSC 2022). BOD 21–01 requires that all currently exploited vulnerabilities (listed in the CISA catalogue) be patched within two weeks (CISA 2021b). Outside government, the Payment Card Industries (PCI DSS) contract requires (Romanosky and Acquisti 2019) that systems handling payment card information must not expose any vulnerability 'scored 4.0 or higher by the CVSS' to external scans (PCI Council n.d.).

Such literal requirements have limitations. First, they only affect organisations subject to the mandate. Others may be encouraged to conform, but CISA can only 'strongly recommend that private businesses and state, local, tribal and territorial (SLTT) governments prioritise mitigation of vulnerabilities listed in CISA's public catalog' (CISA 2021b). Second, as Jacobs et al. (2021) point out, mandates may reinforce incorrect use of the standards. CVSS recommends adjusting 'intrinsic, constant' Base scores using Temporal and (local) Environmental factors (Common Vulnerability Scoring System n.d.); in CVSS version 2.0 these could 'modify the score as much as plus or minus 2.5' (Mell, Scarfone, and Romanosky n.d.), more than the difference between 'critical' and 'medium'. BOD 19-02 ignores these modifiers and uses only the initial vendor severity assessment. BOD 21-01 applies a temporal metric ('currently exploited') to both internal and external systems (CISA 2021a), but ignores local safeguards, such as firewalls and intrusion prevention systems, that reduce the likelihood and impact of exploitation (Allodi and Massacci

2014). It is also unclear whether CISA has a *minimum* severity threshold, for which we might do nothing until more evidence of exploitation emerges. This leaves organisations attempting to comply with no limit on how much labour they should apply and no certainty on whether their approaches will be looked upon favourably, though it is possible to find some acknowledgment that not all CVEs must be addressed deep in the FAQs (CISA 2021b). Maximum work for maximum uncertainty is not an effective policy approach, and publishing clear guidance on CVEs that can remain unaddressed is just as useful. However, it also is not done because it would give clear guidance to malicious actors that these CVEs could work forever. This is a policy paradox that needs resolving; while actors shouldn't be given such information, it is also important for organisations to have clarity on what is below the threshold of remediation for compliance purposes.

PCI DSS includes Environmental factors – respecting firewalls and exempting vulnerabilities that only impact availability – but no obvious Temporal aspect, such as the availability of exploit tools (PCI Council n.d.). Yet we know that the amount of exploitation changes over time, and how the response should change over time continues to go unaddressed. As others have said, we must be careful when we mandate; what was once a floor of best practice can become a ceiling of compliance.

Mandates are helpful but over-simplified risk factors can also produce inaccurate prioritisation (Jacobs et al. 2021). A two-week deadline may not be enough to properly test, schedule and implement all the current mandated patches. The Log4j vulnerability affected 'thousands of software components'; one US Government department spent 33,000 h' effort 'over many weeks and months' on mitigation. During this period, 'mission-critical' work on other vulnerabilities was postponed. The opportunity cost on other vulnerability remediation of any prioritisation approach is rarely considered or measured. Mandates using only Log4j's CVSS Base score of 10.0 (Cyber Safety Review Board 2023) cannot assess if this prioritisation was appropriate or whether, at some stage in its mitigation, local factors including system criticality and attractiveness to intruders meant that unaddressed vulnerabilities – despite a lower CVSS Base score – presented a higher risk of successful attack. It has been suggested that a vulnerability with CVSS Base 7.5 discovered around the same time was 'lost to the hysteria of [Log4j], and … was pushed to the backburner before fading into obscurity.' A year later, this vulnerability is reported to be still present in 8 per cent of network-accessible instances (Baines 2023). This reference represents one of the rare exceptions that the opportunity cost has been documented and two CVEs are compared, and we hope to see much more of it in future reviews of patching policy. It is still insufficient in that it only addresses two vulnerabilities and a single period of time. Yet it raises an interesting question about the measurement of effectiveness of mandates: what is the yardstick (*Quis custodiet ipsos custodes*)?

There is also a risk that overly-specific mandates will constrain the adoption of newer versions of a standard or improved ways of using it (e.g. although CVSS version 3.1 was released in June 2019, BOD 19–02 did not adopt the new version until June 2022). Romanosky and Acquisti (2009) warn that mandates could even 'create a false sense of security' if organisations believe that the actions required by the mandate are all that is needed.

To comment briefly on the levers of enforcement, while PCI can be used to levy fines for non-compliance or in the event of a breach (Fruhlinger 2019), the BOD is more about compliance with US federal law, and cyber essentials take another turn again in that it is guidance-only, but non-compliance may prevent you from getting government contracts.

The fines from the PCI take two primary forms, general non-compliance discovered through audits, and specific non-compliance discovered through breaches. The fines are fiscally structured differently in both cases. There are other PCI consequences available too, which aren't overtly fiscal, but translate into the same effect; for example, banks can increase their processing fees or you can be blocked from processing card data entirely.

We believe these mandates have been helpful as an evolutionary approach but fall short in a number of policy dimensions. Firstly, they chose CVSS scores as a cut-off while ignoring variability in frequency and exploitation. CVSS proportion amongst the total volume of vulnerabilities is very stable over time, but that is a problem as vulnerability numbers keep increasing at a faster rate. In short, there are more and more vulnerabilities even of the critical variety. Secondly, they are silent on *which vulnerabilities* to ignore-until-further-evidence. Thirdly, they don't respect the management of the risk at the local environment or other mitigating controls. Fourthly, they ignore the change of risk profile over time. Fifthly and finally, there is precious little review of their mis-prioritisation of vulnerabilities over time. If they seek to be effective as mandates, they must also seek out evidence of their own mistakes and successes.

However, we must also note the positive quality that these mandates are easy to understand. That is particularly useful when trying to get the attention of boards who have limited time to address a large number of risks (NCSC 2023). Risk is hard to measure, quantify and articulate (Risk Metrics Working Group 2022), thus these mandates emerged precisely because they were easily explainable. That is a quality that needs to be preserved as we move towards approaches that respect evidence of risk, as opposed to estimation of impact.

Laws on behaviour: e.g. data protection

Many organisations are subject to laws that demand particular kinds of behaviour. Fifty years ago, public sector organisations in the state of Hesse, Germany were required to ensure that data 'shall be obtained, transmitted and stored in such a way that they cannot be consulted, altered, extracted or destroyed by unauthorised persons' (Hessen Datenschutzgesetz 1970). Subsequent European laws expanded on this obligation:

Council of Europe Convention 108 (1981): 'Appropriate security measures shall be taken for the protection of personal data stored in automated data files against accidental or unauthorised destruction or accidental loss as well as against unauthorised access, alteration or dissemination' (Council of Europe 1981).

European Data Protection Directive (1995): 'The controller must implement appropriate technical and organisational measures to protect personal data against accidental or unlawful destruction or accidental loss, alteration, unauthorised disclosure or access, in particular where the processing involves the transmission of data over a network, and against all other unlawful forms of processing' (European Union 1995).

General Data Protection Regulation (2018): 'Personal data shall be … processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures' (GDPR 2018).

Similar formulations appear around the globe:

California, Privacy Rights Act, revising the current Consumer Privacy Act, from 1st Jan 2023: 'A business that collects a consumer's personal information shall implement reasonable security procedures and practices appropriate to the nature of the personal information to protect the personal information from unauthorised or illegal access, destruction, use, modification, or disclosure' (California Code Civil Code 2018).

Brazil, LGPD (2018): 'Processing agents shall adopt security, technical and administrative measures capable of protecting personal data from unauthorised access and from accidental or unlawful destruction, loss, alteration, communication or any form of improper or illegal treatment' (Brazilian Government 2018).

India, Personal Data Protection Bill (2022): 'Every data fiduciary and data processor shall protect personal data in its possession or under its control by taking reasonable security safeguards to prevent personal data breach' (Indian Government 2022).

China, Personal Information Protection Law (2021): 'A personal information processor shall, according to the purpose and method of processing personal information, type of personal information, impact on individual's right and interest, and possible security risk, etc., take the following measures to ensure the compliance of personal information processing activities with provisions of laws and administrative regulations, and prevent unauthorised visit, or leakage, falsification, and loss of personal information' (China Briefing 2021).

Nigeria, Data Protection Regulation (2019): 'Anyone involved in data processing or the control of data shall develop security measures to protect data; such measures include but not limited to protecting systems from hackers, setting up firewalls, storing data securely with access to specific authorised individuals, employing data encryption technologies, developing organisational policy for handling Personal Data (and other sensitive or confidential data), protection of emailing systems and continuous capacity building for staff' (Nigerian Government 2020).

These examples share four key features:

- a requirement on **behaviour** – 'shall be taken', 'shall be processed', 'shall implement', 'shall adopt', 'shall take' – irrespective of outcomes or harms that may result;
- the behaviour is measured against an **external standard**, not 'industry practice', which may be inadequate (Solove and Hartzog (2022) report that 58 out of 59 random audits in 2014 found breaches of HIPAA rules when handling health data; in October 2022 more than 160 breach notification letters were reported to the US State of Massachusetts (Massachusetts Government 2022));
- use of 'appropriate' and similar formulations, implying an **assessment of risk**; and,
- failure to meet the requirement can result in **sanctions**, again even if no harm or damage has (yet) been caused.

The first and fourth of these encourage organisations to adopt relevant behaviours, by threatening sanctions if they do not. The second and third place actions that are known to reduce risk among those desired behaviours. Indeed, the European Data Protection Board (2021) has identified 'proper patch management' as 'one of the most important' protective measures, linking it directly to this incentive.

Cases on vulnerability response

Vulnerability response is analysed below in four Monetary Penalty Notices (MPNs, commonly referred to as 'fines') issued by the UK Information Commissioner (hereafter 'ICO') between 2018 and 2023. After introducing the cases, we consider how the four key features above were used to assess practice, then identify specific behaviours that the Regulator expected. One case was reviewed by an Appeal Tribunal, which confirmed the importance of patch management.

Tuckers Solicitors (2022) self-reported a ransomware attack: the attacker encrypted nearly a million files and published sixty document bundles from legal cases on an underground market site. Subsequent investigation found a serious vulnerability left unpatched for at least five months after remedies were announced by the software vendor (ICO 2022). Tuckers, by investigating fastest, allowed the most helpful analysis: in the case of DSG (2022) a system holding over twenty-five million people's personal details and five million payment card details was vulnerable for four years (ICO 2020a); the case of Carphone Warehouse (2018) found a six-year-old Wordpress installation (ICO 2018a); the case of Cathay Pacific (2020) identified one ten-year-old vulnerability, one server lacking sixteen patches from an eight-month period and a domain controller unpatched for a year despite Microsoft releasing twelve updates (ICO 2020b). The ICO's and Tribunal's analyses of these incidents highlight links between the four common features of behaviour laws and risk-based vulnerability response.

Each organisation was sanctioned for behaviour – failing to manage vulnerabilities – *without needing proof that this caused the actual data breach*. For example, 'Neither Tuckers nor [their investigator] was able to determine conclusively how the attacker was able to access Tuckers' network. However, it did find evidence of an unpatched vulnerability … that could have been used to either access the network, or further exploit areas of Tuckers once inside the network (ICO 2022)'. In Carphone Warehouse, the breach used valid user credentials, but the investigation's identification of poor patching practice increased the penalty imposed.[1]

The comments on the cases of Carphone Warehouse, Cathay Pacific and DSG clearly show that vulnerability and patch management are required behaviours: 'Carphone Warehouse's approach to software patching was seriously inadequate';[2] 'Cathay Pacific could not provide any evidence of up-to-date patch management for either System A or System C servers … If Cathay Pacific had operated more effective patch management, attackers would have had less opportunity to exploit known vulnerabilities';[3] 'DSG allowed the critical risks relating to patch management and administrator password management to persist from May 2017, even though this was highlighted as a critical risk in penetration tests carried out several months apart.'[4] In DSG's appeal, the Tribunal confirmed the ICO's view 'that DSG's failure to take appropriate measures in relation to this risk was a contravention of [the Data Protection Security Principle] for which it is appropriate to hold DSG to account,' adding that 'As a matter of common sense, we find that security patches are generated in response to known vulnerabilities.'[5]

Policies must be followed: this was a failing of both Carphone Warehouse (ICO 2018a) and DSG (ICO 2020a), where an essential post-patch step was not performed, letting the attackers obtain administrator powers three years later.[6] 'Industry practice' comparisons may condemn – Tuckers against standards from the Law Society (ICO 2022) and Solicitors'

Regulatory Authority – but cannot excuse. The ICO 'rejects' accusations of demanding more than contemporary industry practice: 'The deficiencies set out … represent appropriate measures that data controllers such as [DSG/Carphone Warehouse] should have had in place [at the time of the incident]' (ICO 2020a; 2018a). Indeed, 'Carphone Warehouse is a large, well-resourced and experienced data controller … A company of this size and standing was well placed to assess any weaknesses in its data security arrangements and to take appropriate action': such organisations should lead 'the state of the Art' in security measures, not follow others (European Union 1995).

The need for a risk-based approach is explicit in the GDPR's requirement for 'appropriate technical and organisational measures to ensure a level of security appropriate to the risk [to the rights and freedoms of natural persons]' (GDPR 2018); other legislation demands 'appropriate measures' and equivalent formulations. In DSG's appeal, 'There is no evidence before us of any risk assessment or decision(s) … relating to the critical risks of security patch management and password practices.'[7] Risk factors to be considered, according to the ICO Notices, include:

- the quantity and nature of data being processed: DSG held data about 25 million individuals (ICO 2020a) and over 5 million payment cards,[8] Tuckers should have taken account of 'the highly sensitive nature of the personal data' (ICO 2022);
- the impact on individuals of any breach: DSG's records were likely to be taken for 'nefarious and criminal purposes': more harmful than accidental loss (ICO 2020a);
- the likelihood of the vulnerability being exploited: in Cathay Pacific, 'The vulnerability had been described as allowing 'remote attackers to bypass authentication and gain administrative access via direct request'. The complexity of the vulnerability was described as low – meaning 'very little knowledge or skill is required to exploit it' (ICO 2020b);
- third-party assessments: in DSG, 'Microsoft announced this as an 'important' update that should be applied at the 'earliest opportunity' in order to be fully protected (ICO 2020a),' and Tuckers' failure was 'particularly negligent given that the NCSC had published an Alert drawing attention to [the vulnerability]' (ICO 2022).

Risk assessment should consider the 'costs of implementation' (European Union 1995). In the case of Tuckers, the ICO recognised that even a free software update has costs, 'such as the cost of personnel to test the patch prior to deployment' (ICO 2022); costs to data, service and users from planned and unplanned downtime can be included. Each case highlights a particularly urgent vulnerability, recognising the need to prioritise (ICO 2022). Effectively re-emphasizing that risk informs two choices: which vulnerability to address first, and how and when to mitigate it. Some vulnerabilities have to wait, either until higher-risk vulnerabilities have been addressed, or until a particularly high mitigation cost can be reduced, and it is legally recognised as right to do so.

Inadequate behaviour leads to significant sanctions: Cathay Pacific was fined the maximum £500,000 permitted by the 1998 Data Protection Act (ICO 2020b); Carphone Warehouse £400,000 under that Act (ICO 2018a); and Tuckers £98,000 under the 2018 Data Protection Act (ICO 2022). DSG's original £500,000 fine was reduced on appeal, but £250,000 imposed for 'inconsistent patch management' alone.[9] Investigations and fines focus on 'the kind of contravention rather than the actual consequences' (ICO

2020a). In Tuckers, although the harm was caused by the attacker, 'the infringements identified by the Commissioner were relevant to the personal data breach because they gave the attacker a weakness (vulnerability) to exploit and/or because they increased the risks to personal data once the attacker entered Tuckers' network' (ICO 2022). The Carphone Warehouse attacker used valid credentials but this 'does not … absolve' their 'seriously inadequate' software management that 'exposed the content of the system to very serious risk' (ICO 2018a). Elsewhere, fines have been imposed for breaches discovered in time to prevent harm, for example where patient records were left in a vacant building (ICO 2018b).

These Monetary Penalty Notices show that risk-based vulnerability response should reduce the risk of sanctions. We next examine how detailed analysis supports the adoption of the tools and techniques discussed above.

As an interesting aside, the money gathered from such cases essentially funds the ICO, though there are some extra steps and independent accounting via the UK Government's Consolidated Fund (ICO n.d.). To those unfamiliar with the UK's budget and financing, this is essentially income usable directly by government (UK Parliament n.d.). The details of how many MPN notices were delivered by sector can be found in the Table 1 below (ICO n.d.). The maximum such fine imposed can be £17.5 million or 4 per cent of the global turnover.

In the 2021–2022 year, the ICO issued fines of around £633,000. In 2022–2023, the DPA fines increased to £15.2 million. Note that we're comparing GDPR fines to DPA-related fines because they are both about data protection, but they are different for a variety of reasons too lengthy to go into here.

By comparison, GDPR cumulative fines are orders of magnitude higher (GDPR Enforcement Tracker n.d.). We can't compare data yet for Brazil's LGPD, but we can say that they published their first fine in July this year (Mari 2023).

Support for risk-based vulnerability response

The MPNs highlight triggers that should prompt action: 'both the vulnerability and the fix [were] public knowledge' in Cathay Pacific (ICO 2020b); 'malicious actors were exploiting the … vulnerability' in Tuckers (ICO 2022). Publication and exploitation are common factors in risk-based vulnerability response, as is the use of relevant third-party information sources: vendors such as Microsoft in DSG (ICO 2020a); industry sources such as

Table 1. The number of MPNs issued by the UK's ICO.

Sector	2021–2022	2022–2023
Marketing	2	9
Finance	2	9
Retail and manufacture	4	8
General business		2
Land or property		2
Technology and telecoms	1	1
Utilities	1	
Health	7	1
Legal		1
Central government	1	
Charity and voluntary	1	
Membership association	1	

CVE in Cathay Pacific (ICO 2020b); and national centres such as the UK National Cyber-Security Centre (NCSC) in Tuckers (ICO 2022).

Organisations should respond to external severity assessments, but not by simply patching immediately. Tuckers approves the 'ISO27002 suggestion that organisations should define a timeline to react to notifications of potentially relevant technical vulnerabilities, and once a vulnerability has been identified, associated risks should be identified and actions taken, such as patching the system to remove the vulnerability' and such as recognising that mitigations other than patching may be appropriate. This Notice expects organisations will 'test the patch prior to deployment' and even for a 'critical' (CVSS Base 9.8) vulnerability, remedial action should be 'prompt', citing the NCSC Cyber Essentials deadline of 14 days after patch release. 'Not applying a high-risk security patch until four months after it was released, despite it being listed as 'critical'' was considered 'negligent practice' (ICO 2022).

None of these cases discusses prioritisation among vulnerabilities, but linking timescale to severity implies that this is both necessary and appropriate. Properly assessing, testing and remediating a single vulnerability with complex dependencies (such as Log4j, discussed above (Cyber Safety Review Board 2023)) would require most of the 14-day window. This may require postponing less urgent work, deprioritizing vulnerabilities that are likely to affect fewer systems or have lighter impact. The Regulator's recognition that remedial measures have costs (ICO 2022) and that vulnerability risks vary also supports prioritisation among remedies. If no external source has indicated particular urgency and statistical measures suggest immediate attacks are unlikely, the organisation can consider whether the impact of patching could be reduced by delaying. That allows more compatibility testing, choosing a less disruptive time, and reducing disruption risk from over-rapid reaction.

These Notices confirm that, under a behavioural law, vulnerability response is one of the legally-required behaviours. This does not depend on industry practice, risk assessment or reputable external evidence. That should be part of behaviour; and sanctions may be imposed for inadequate behaviour, even if no harm resulted. We conclude these laws provide strong support for risk-based vulnerability response.

Laws on outcomes: e.g. product law

We now consider laws that regulate outcomes rather than processes. For example, manufacturers of many physical products are held responsible if the product is defective. This encourages process improvements that reduce the likelihood and expense of defects, but can only do so for processes before the 'outcome' is assessed. Many laws judge fitness/defectiveness at the point where a product is placed on the market, which has the virtues of simplicity and clear responsibility. However, many processes affecting the fitness of software products and services – notably vulnerability response – occur *after delivery to the customer*.

Point of sale laws ignore these activities: they could de-prioritise them or encourage manufacturers to postpone risky activities until after the legal 'outcome' is assessed. Post-sale fitness assessment raises significant policy questions and conflicts: how much should we reduce the manufacturer's responsibility when users do not look after their products? Conversely, how much freedom should users/owners have to use, modify or repurpose their purchase outside the manufacturer's control? This section first analyses

traditional 'pre-sale' outcomes laws, then considers whether other 'post-sale' laws suggest ways to encourage vulnerability response.

Regulation before sale

Any law that regulates outcomes must decide when the outcome is assessed. Physical product liability laws usually select the moment when a product is put on sale. Until then, the manufacturer has almost complete control over quality and fitness: afterwards, once the purchaser has the product, almost none. Having bought a device, I can (ab)use it however I choose, whatever the manufacturer intended. This transfer of factual control is an obvious, and usually efficient, moment to also transfer legal responsibility. Making manufacturers responsible for the outcome of what they do before sale creates an incentive to do that well: the results benefit all purchasers, so any cost can fairly be added to the purchase price. Responsibility for the results of post-sale abuse by a few purchasers could encourage manufacturers to remove useful features that might be misused – to over-engineer the product to work under extreme mistreatment. This reduces functionality and increases price for all purchasers, most of whom will never benefit.

The idea that manufacturers completely control outcomes often produces an absolute test of fitness. The European Product Liability Directive (1985) is an example of legislation that makes manufacturers strictly liable – irrespective of fault or mitigating circumstances – for any damage caused by any defect that could have been detected given the 'state of scientific and technical knowledge at the time when he put the product into circulation'. This binary rule may be efficient when allocating the costs of faulty devices – a disappointed purchaser need not prove fault nor where in a complex supply chain it lies – but can be unhelpful in managing software vulnerabilities where a risk-based approach is necessary and efficient to minimise harm.

Fundamentally, software product-makers have capabilities after sale: by offering updates (although the 2009 amendment to the ePrivacy Directive (European Union 2009) entitles device owners to withhold their consent to install them) or by providing managed software as a 'hosted' or 'cloud' service (ENISA 2009). Unlike physical products, defects in software can be repaired, discovered and even introduced long after the moment of sale (Coatanroch et al. 2022). Laws should encourage designs and business models that improve software quality throughout its life, whether delivered direct to a user or providing a third-party service.

The 2021 expansion of the European Radio Equipment Directive (2014/53/EU) (RED) (European Union 2014) still applied the 'market placement' model to radio-equipped toys, smart devices, cameras, mobile phones, laptops, dongles, alarm systems, home automation systems and wearables (European Union 2021). Post-sale software updates are vital for these products (European Parliament 2022), but this law gives manufacturers no incentive to provide them. The UK's Product Security and Telecommunications Infrastructure Act attempts by requiring – at point of market placement – a vulnerability policy that provides 'transparency about the length of time for which the product will receive important security updates' (UK Government 2021). A European Parliament report notes accountability challenges 'if developers fail to meet promised development or update deadlines': a pre-sales law cannot provide ongoing monitoring or enforcement! (Coatanroch et al. 2022, 36).

Simply changing the moment when outcome is assessed is unlikely to provide appropriate incentives because it makes the manufacturer design for the most reckless user,

reducing quality and increasing costs for the reasonable majority. To explore these issues, we consider laws that address faults discovered after sale.

Regulation after sale

Unlike the Product Liability and Radio Equipment Directives, the European Product Safety Directive (2001/95/EC) creates obligations before and after sale, but only for products that pose a safety risk and faults affecting the 'safety and health' of humans. Their manufacturers have a

> duty to adopt measures commensurate with the characteristics of the products, enabling them to be informed of the risks that these products may present, to supply consumers with information enabling them to assess and prevent risks, to warn consumers of the risks posed by dangerous products already supplied to them, to withdraw those products from the market and, as a last resort, to recall them when necessary. (European Union 2001)

This changes the point of assessment and what is assessed – to the moment when a safety risk is discovered in a 'product already supplied'. Sending information, warnings and recalls are tests of the manufacturer's behaviour, not of whether owners' responses to those warnings produce good outcomes. This assessment reflects 'commensurate' risk, not a strict binary test.

Product safety law might thus also inform behaviour-based laws on vulnerability response. Patches, like physical safety measures, should be developed and deployed quickly to remedy an unsafe situation (usually involving a risk of physical harm). Clearly, pre-sale and post-sale activities are both needed. Some post-sale remedial actions are similar: issuing a vulnerability notice is analogous to 'warning', disabling an unsafe software function to 'recalling'. The extensive RAPEX Risk Assessment Guidelines for Consumer Products might also be a useful model for software (European Union 2019).

Safety law covers limited products, in circumstances involving extreme risk. Patching should be routine for all digital services and should address many more risks by contrast with product recalls. Digital service providers have capabilities that physical manufacturers do not: a software update of a connected device costs less than a physical product recall. Effective laws should encourage providers to act on these possibilities (Coatanroch et al. 2022).

Post-sale regulation has policy conflicts: including sustainability, the environment, consumer rights and safety/security. Digital components make products easier to repair, recycle, refurbish, upgrade and re-purpose, offering a 'whole new range of sustainable services, product-as-service models and digital solutions' (European Commission 2022). Sometimes legislation can discourage these beneficial activities. A European Parliament report warns that treating software updates as new products, with burdensome transparency requirements, might 'encourage firms to make fewer or less frequent updates' or to address safety risks by informing customers rather than updating the product (Coatanroch et al. 2022). The original Radio Equipment Directive (2014/53/EU) warns that laws restricting software updates by users or third parties could be 'abused' for anti-competitive purposes (European Union 2014). Unclear responsibilities when 'the producer is no longer able to control software or other technical features subsequently installed in or

learned by the product' discourage makers, maintainers and users alike (European Union 1985).

Recognising that the RED expansion only 'address[es] the problems regarding products lacking [all] security features' (European Union 2014), the European Commission's Cyber Resilience Act was developed to 'address cybersecurity risks in all connected products and associated services and throughout their entire lifecycle'. The 2022 Call for Evidence identifies problems both before and after sale: for economic reasons 'vendors … often do not put in place adequate cybersecurity safeguards' and 'vendors' response to vulnerabilities throughout their products' lifecycle is too often inadequate.' The legal framework is insufficient, both because it 'does not cover all types of digital products … [and] non-embedded software products are not addressed', and because it 'does not prescribe specific cybersecurity requirements, e.g. covering the whole life cycle of a product'. Vulnerabilities are specifically identified – 'Whole life cycle' requirements are crucial in the case of digital products and ancillary services, as software needs to be updated on a regular basis' – noting that 'vulnerabilities in software products are increasingly serving as a channel for cybersecurity attacks, causing significant societal and economic costs' (European Parliament 2022).

The draft Cybersecurity Regulation covers 'all connectable hardware and software products' that are supplied 'in the course of a commercial activity, whether in return for payment or free of charge'. Software-as-a-Service platforms are excluded as likely to be covered by the cloud computing provisions of the Network and Information Security Directive. However, non-critical applications are included, recognising the risk that attackers will use them as an entry point from which to traverse to systems with greater impact. Software manufacturers are required to satisfy a wide range of essential requirements, including point-of-sale outcomes such as 'Products with digital elements shall be delivered without any known exploitable vulnerabilities', but the extensive post-sale obligations are all behavioural in nature: including receiving vulnerability reports, acting to address problems, and publishing the resulting mitigations. These activities must continue for at least five years, or the expected lifetime of the product if shorter. Manufacturers must also inform ENISA and customers of any 'actively exploited vulnerability' or 'incident having impact on the security of the product'. Surprisingly, few of the policy conflicts mentioned in preparatory work are addressed: 'Refurbishing, maintaining and repairing' will 'not necessarily' constitute a new product, but re-purposing a product in accordance with sustainability objectives seems likely to constitute a 'substantial modification' and thereby incur full manufacturer obligations under Article 15. The exemption for 'free and open source software' is narrowly drawn, excluding those who offer commercial support or monetise other services (European Union 2022a).

This new Regulation represents a considerable expansion of the scope of the Product Safety and Radio Equipment laws and provides a detailed and demanding specification of vendors' post-sale vulnerability management behaviour. Article 43, highlighting vulnerability handling as an activity that may present a significant cybersecurity risk, requires national authorities to assess and take 'all appropriate provisional measures to prohibit or restrict that product' if the manufacturer does not correct any non-compliance found. Fines up to €15 million or 2.5 per cent of global turnover are available. However, the subjects of this paper – the organisations that receive patches and mitigations – are not addressed. Patches, mitigations and information may become more readily available to such software users, especially if the suggestion of 'automatic updates' is

adopted. However, there is no new incentive for them to act on these! Even in a clear case of a widely available product used in critical sectors, whose vulnerabilities are being exploited by malicious actors, regulators can only intervene with the product's manufacturer, importer or distributor, not its users (European Union 2022a).

Finally, it is notable that this legislative proposal does not attempt to impose outcomes-based duties after market placement. Any future regulation of post-sale vulnerability response seems likely to be behaviour-based, following the models of the Cybersecurity Regulation and existing data protection enforcements.

Laws on harm: e.g. liability

The final group of laws respond to harm resulting from an organisation's action or inaction: often called 'liability', 'negligence' (for inaction) or, in English and US legal systems, 'tort'. They aim to compensate victims of harm, spread its costs (for example by encouraging insurance) and deter harmful behaviour (Hedley 2002).

As a way to promote specific behaviour such as vulnerability response, however, their power is limited. The link between behaviour – particularly absence of behaviour – and a successful claim for harm involves several steps. In tort law: did the entity said to be responsible owe the specific victim a duty of care?; did its (in)action breach that duty?; did that breach (directly) cause the harm?; and was the harm of a kind that can be claimed? (Romanosky and Acquisti 2009). Sometimes also: was the harm a reasonably foreseeable result of the behaviour? (Overseas Tankship 1961).

Uncertainty at any stage can break the liability chain or, viewed in reverse, remove the incentive to behave differently. US data breach cases have refused to recognise harms, costs and even victims' standing to bring claims (Solove and Hartzog 2022). The case of Tulip Trading dismissed many arguments for a duty of care on software developers (Royal Courts of Justice 2023); and the case of Cathay Pacific found a vulnerability but could not prove it was used by the attacker (ICO 2020b); the actual source of leaked private information may be impossible to determine (Romanosky and Acquisti 2009). The European Commission note 'The complexity of digital technologies (e.g. within IoT systems) could make it very challenging for injured parties to identify the producer responsible' (European Commission 2022).

Even successful claims create little incentive to improve software management, as compensation awards are typically limited to physical damage and personal injury (Bussani 2011). Systems likely to cause these harms are often subject to 'behaviour' laws – including consumer product safety (UK Government 2019) (discussed above) and medical device regulation (European Union 2017) – that provide stronger and clearer incentives than general liability. Victims of poor software management are more likely to suffer economic or emotional harm, including possible future consequences of a data breach, and these are very rarely covered by a negligence claim (Romanosky and Acquisti 2009). The case of Smith and Others v TalkTalk dismissed as 'not within … scope' a claim 'that the Defendant negligently published webpages which, via a vulnerability which was known or should have been known, enabled criminal hackers to access the information'. The claimants were directed instead to data protection law (Graeme Smith n.d.) which did, in the case of Vidal-Hall, allow a claim for 'anxiety and distress' (Royal Courts of Justice 2015). However that claim involved active misuse of personal

data not, as in the case of Smith, inaction that let a third-party cause harm. Overall, the many ways to dispute liability, and the tiny proportion of breaches that result in claims, dilute motives for organisations to change behaviour (Romanosky and Acquisti 2009).

Despite economically inefficient liability laws creating little direct incentive on digital service providers, they may indirectly promote insurance markets, whether for traditional high-cost, low-likelihood liabilities such as errors and omissions (Floresca 2014), or by mandating insurance (Woods and Simpson 2017), or by creating post-breach costs that insurance can mitigate. Fines themselves cannot normally be insured according to some scholars but costs imposed by privacy and data protection laws – including forensic investigation, legal advice, victim notifications, credit monitoring and public relations (Floresca 2014) – can. Once insured, moral hazards emerge; organisations might feel 'lethargic in taking ownership of compliance policies and procedures' (Talesh 2018).

Insurance can encourage preventive measures that make claims less likely. These may be a prerequisite for obtaining insurance and a way to reduce premiums (Talesh 2018). Insurers know from the claims they receive (Woods and Simpson 2017) which measures best reduce risk and can access 'subject matter experts who know the latest vulnerabilities and the cyber risk landscape and are able to provide specialised knowledge to clients to ensure that their cyber infrastructure is secure' (Talesh 2018). Harm-based laws that encourage cyber insurance could promote risk-based vulnerability remediation.

Conclusion

Having examined four types of incentivizing laws – mandates, behaviour-based, outcomes-based, and harm-based – we conclude that, beyond a mandate requiring organisations to adopt specific practices, behaviour-based laws are most likely to encourage the adoption of effective vulnerability remediation.

General harm-based (liability) laws demand close links in causation, foreseeability and responsibility between (in)action and resulting harms. Though foreseeability is increasing with respect to both exploitability (EPSS) and forecasting (Vulnerability Forecasting), the causal chains remain muddy and inefficient.

Although the Mirai botnet relies on unpatched webcams and video recorders, its victims – including Twitter, Netflix and CNN (Wolf 2016) – are unlikely to have any claim against their manufacturers given the UK decision in the case of Smith. The low risk of having to pay damages gives organisations little incentive to improve practice. Even this only affects products and services that cause physical harm; here behaviour and outcome-based laws, such as product and medical safety, should create more direct vulnerability response obligations (European Union 2022a).

Indirectly, privacy and other laws that impose costs (rather than fines) may support an insurance market in which insurers create requirements or incentives to adopt practices, including vulnerability response, that reduce claims. The cyber insurance industry has also remained silent on what CVEs they expect to be patched, essentially treating such knowledge as proprietary to their underwriting.

Outcome-based laws must define when the outcome will be assessed. To ensure a single entity can reasonably be held responsible for that outcome, this is typically the point of sale or market placement. Such laws cannot influence activities – including vulnerability response – that occur after this point. Postponing the assessment allows

attributing responsibility among multiple parties – at least the manufacturer and the owner – raising significant public policy conflicts. It is perhaps for this reason that we were unable to identify any post-sale outcome-based laws: the post-sale obligations in the proposed Cybersecurity Regulation are behavioural in nature.

Mandates are more promising, though they are limited in scope and sometimes embed inefficient practices. They may aim for 'extreme deterrence' rather than 'optimal risk management' (Solove and Hartzog 2022). For example, the US Government's BOD 19–02 requires patching every vulnerability whose CVSS Base score exceeds 7.0 (CISA 2019): most of these will never be exploited and local conditions may make a lower Base scoring vulnerability more urgent for a particular organisation. BOD 21–01 uses a more efficient prioritisation metric (Cyentia Institute n.d.a) of 'known exploited vulnerabilities' (CISA 2021b). Mandates that consider local risk factors can offer better results for the same effort.

Behaviour-based laws appear most effective. These require a broad range of organisations to adopt proactive risk-based security processes, with penalties even if risky behaviour has not caused harm. Regulators' rulings have identified vulnerability response as a required process and provided considerable detail on expected features. The risk-based requirement supports tools that help service managers prioritise among the competing risks of (un)scheduled downtime, taking time to test mitigations, and leaving vulnerabilities temporarily unpatched. Phrases such as 'having regard to … the cost of their implementation' (European Union 1995) recognise limits – both to organisations' vulnerability response resource and tolerance of disruption.

Such laws also avoid the inefficiencies identified by Romanosky and Acquisti (2009), that dilute legal incentives. Regulators need not identify every harm arising from an organisation's behaviour: they can assess behaviour directly and demand an 'appropriate' standard reflecting the organisation's particular risk. They can investigate and sanction risky behaviour before harm occurs. This protects individuals by encouraging organisations to identify and fix problems before a breach (Solove and Hartzog 2022). Penalty Notices have particularly criticised organisations that were, or should have been, aware of risks but did not act to mitigate (ICO 2018a).

The proposed European Cybersecurity Regulation takes behaviour-based approaches to patch management by vendors, but more relevant to the subject of this paper is the Regulation on Digital Operational Resilience for the Financial Sector (DORA), which applies it to vulnerability response by some software users. DORA will, according to the agreed final text (European Union 2022b), require financial services organisations (Art.2(1)) to 'have appropriate and comprehensive documented policies for patches and updates' (Art.9(4)(f)) as part of a 'risk management framework' (Art.6(1)) that includes 'policies, procedures and controls for ICT change management … based on a risk assessment approach' (Art.9(4)(e)). Not implementing such policies should carry 'effective, proportionate and dissuasive' penalties, even if harm has not yet occurred (Art.50(3)). DORA therefore has the four characteristics likely to incentivize desired practice: requirements on organisations to adopt a specific, risk-based behaviour, with direct sanctions for failure to do so. Its explicit vulnerability response requirement strengthens the GDPR case law expectations above.

Linking patching practices to legal requirements supports good practice. Developers and implementers of vulnerability remediations can promote them in terms of legal analysis and requirements. They are relevant to executives, managers, compliance officers and regulators alike. A common language of risk helps communication between silos:

vulnerability response is not just a 'technical' issue (Solove and Hartzog 2022). Effective tools and practices should become good practice that propagates through outcome-based laws. There is 'significant room for improvement' (Jacobs et al. 2021), perhaps twenty-fold for some organisations and data subjects (Cyentia Institute n.d.a). Remediating vulnerabilities already challenges most organisations (Cyentia Institute n.d.b) and their volume is growing inexorably. Risk-based vulnerability response is essential if we are to protect our systems from harm.

Notes

1. Carphone Warehouse (n 102) [15].
2. Carphone Warehouse (n 102) [22(2)].
3. Cathay Pacific (n 103) [24(8)].
4. DSG Retail appeal (n 101) [114].
5. DSG Retail Limited v Information Commissioner.
6. DSG Retail Limited v Information Commissioner.
7. DSG Retail Limited v Information Commissioner.
8. DSG Retail Limited v Information Commissioner.
9. DSG Retail Limited v Information Commissioner.

Notes on contributors

Andrew Cormack was the chief regulatory adviser at JISC, and a long time liaison of FIRST (Forum of Incident Response and Security Teams). As well as his prodigious security, data protection, and legal knowledge, his acknowledged mastery of the interplay between legal and engineering concerns made him one of the most respected experts in the global cyber security and policy communities. Andrew graduated in Mathematics from Cambridge University in 1984. As a life-long distance learner, he has also obtained degrees in law and humanities from the Open University and a Masters in Computer and Communications Law from Queen Mary, University of London (2015). He worked for Plessey Telecommunications, the Natural Environment Research Council's Research Vessel Services, and Cardiff University, before being appointed head of Janet-CERT in 1999. He was a member of the Permanent Stakeholders' Group of ENISA for ten years, and chair of the Funding Council of the Internet Watch Foundation from 2009 to 2013.

Eireann Leverett is a technologist and entreprenuer, with a focus primarily on incdient response and cyber crime. He initially studied psychology and philosophy at Antioch College. He graduated with a BEng from Edinburgh University in 2005 in AI and Software engineering, and a Masters in Advanced Computer Science from Cambridge in 2011. He is particularly motivated by quantifying risks in digital domains. He has spent time in recent years doing data science around vulnerabilities, and sees a great leap forward in the predictive and legal understanding of vulnerability and exploit risk. He and Andrew shared many walks and a more than few jokes of life in Scotland, Wales, Cambridge, and technology policy.

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Éireann Leverett http://orcid.org/0000-0001-6586-7359

References

Allodi, Luca, and Fabio Massacci. 2014. "Comparing Vulnerability Severity and Exploits Using Case-Control Studies." *ACM Transactions on Information and System Security*, 17. ACM. https://doi.org/10.1145/2630069.

Allodi, Luca, Fabio Massacci, and Julian Williams. 2021. "The Work-Averse Cyberattacker Model: Theory and Evidence from Two Million Attack Signatures." *Risk Analysis*. https://doi.org/10.1111/risa.13732.

Baines, Jacob. 2023. "Assessing Potential Exploitation of Grafana's CVE-2021-43789 for Initial Access." 21 February. https://vulncheck.com/blog/grafana-cve-2021-43798.

Brazilian Government. 2018. Lei Geral de Proteção de Dados Pessoais, Article 46. https://www.planalto.gov.br/ccivil_03/_ato2015-2018/2018/lei/l13709.htm.

Bussani, M., ed. 2011. *Pure Economic Loss in Europe*. Cambridge: Cambridge University Press.

California Code Civil Code. 2018. "Civ Division 3 - Obligations Part 4 - Obligations Arising from Particular Transactions Title 1.81.5." California Consumer Privacy Act of 2018 Section 1798.100. (2018). https://law.justia.com/codes/california/2018/code-civ/division-3/part-4/title-1.81.5/section-1798.100/.

China Briefing. 2021. "RC Personal Information Protection Law (Final): A Full Translation". 24 August. https://www.china-briefing.com/news/the-prc-personal-information-protection-law-final-a-full-translation/.

CISA. 2016. "SSL 3.0 Protocol Vulnerability and POODLE Attack." 30 September. https://www.cisa.gov/news-events/alerts/2014/10/17/ssl-30-protocol-vulnerability-and-poodle-attack.

CISA. 2019. BOD 19-02: Vulnerability Remediation Requirements for Internet-Accessible Systems. https://www.cisa.gov/news-events/directives/bod-19-02-vulnerability-remediation-requirements-internet-accessible-systems.

CISA. 2021a. BOD 22-01: Reducing the Significant Risk of Known Exploited Vulnerabilities. https://www.cisa.gov/news-events/directives/bod-22-01-reducing-significant-risk-known-exploited-vulnerabilities.

CISA. 2021b. "CISA Releases Directive on Reducing the Significant Risk of Known Exploited Vulnerabilities." 3 November. https://www.cisa.gov/news-events/news/cisa-releases-directive-reducing-significant-risk-known-exploited-vulnerabilities.

CISA. 2022. "Known Exploited Vulnerabilities Catalog." https://www.cisa.gov/known-exploited-vulnerabilities-catalog.

Coatanroch, et al. 2022. "New Technologies and New Digital Solutions for Improved Safety of Products on the Internal Market." European Parliament, 30 June. https://www.europarl.europa.eu/thinktank/en/document/IPOL_STU(2022)703348.

Common Vulnerability Scoring System Version 3.1: User Guide. n.d. https://www.first.org/cvss/user-guide.

Condon, Caitlin, Jake Baines, Spencer McIntyre, and Brendan Watters. 2021. "Rapid7 2021 Vulnerability Intelligence Report." Rapid7. https://information.rapid7.com/rs/411-NAK-970/images/Rapid7%202021%20Vulnerability%20Intelligence%20Report.pdf.

Council of Europe. 1981. Convention for the Protection of Individuals with Regard to Automatic Processing of Personal Data. https://rm.coe.int/1680078b37.

CVE. 2023. "CVE Metrics." https://www.cve.org/About/Metrics.

CVE. n.d. "CVE® Program Mission." https://www.cve.org/.

Cyber Safety Review Board. 2023. "Review of the December 2021 Log4j Event." 11 June. https://www.cisa.gov/sites/default/files/publications/CSRB-Report-on-Log4-July-11-2022_508.pdf.

Cyentia Institute, Kenna Security. n.d.a.. "Prioritization to Prediction Volume 2: Measuring and Minimizing Exploitability." Prioritization to Prediction. https://library.cyentia.com/report/report_002992.html.

Cyentia Institute, Kenna Security. n.d.b. "Prioritization to Prediction Volume 8: Measuring and Minimizing Exploitability." Prioritization to Prediction. https://library.cyentia.com/report/report_008756.html.

DSG Retail Limited v Information Commissioner. 2022. (UK First Tier Tribunal 5 July).

ENISA. 2009. "Cloud Computing Risk Assessment." 20 November. https://www.enisa.europa.eu/publications/cloud-computing-risk-assessment.

ENISA. n.d. "Vulnerabilities and Exploits". Accessed 4 November 2022. https://www.enisa.europa.eu/topics/csirts-in-europe/glossary/vulnerabilities-and-exploits.
EPSS Model Motivation. n.d. https://www.first.org/epss/model.
European Commission. 2022. "Product Liability Directive - Adapting Liability Rules to the Digital Age, Circular Economy and Global Value Chains." 11 December. https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12979-Product-Liability-Directive-Adapting-liability-rules-to-the-digital-age-circular-economy-and-global-value-chains_en.
European Data Protection Board. 2021. "Guidelines 01/2021 on Examples Regarding Personal Data Breach Notification." January. https://edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-012021-examples-regarding-personal-data-breach_en.
European Parliament. 2022. "Cyber Resilience Act - Impact Assessment." 15 September. https://digital-strategy.ec.europa.eu/en/library/cyber-resilience-act-impact-assessment.
European Union. 1985. Directive 85/374/EEC - Product liability for defective products, § Art 1. https://osha.europa.eu/en/legislation/directives/council-directive-85-374-eec#:~:text=of%2025%20July%201985%20on,concerning%20liability%20for%20defective%20products.&text=The%20Directive%20establishes%20the%20principle,a%20defect%20in%20his%20product.
European Union. 1995. Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data.
European Union. 2001. Directive 2001/95 EC - Product Safety. https://osha.europa.eu/en/legislation/directives/53.
European Union. 2009. Directive 2009/136/EC of the European Parliament and of the Council. https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:337:0011:0036:en:PDF.
European Union. 2014. Directive 2014/53/EU of the European Parliament and of the Council. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32014L0053.
European Union. 2016. "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation)". https://eur-lex.europa.eu/eli/reg/2016/679/oj.
European Union. 2017. Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and repealing Council Directives 90/385/EEC and 93/42/EEC (2017). https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32017R0745.
European Union. 2019. Commission Implementing Decision (EU) 2019/417, § Annex 3. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32019D0417&from=EN.
European Union. 2021. Commission Delegated Regulation (EU) 2022/30 of 29 October 2021 supplementing Directive 2014/53/EU of the European Parliament and of the Council with regard to the application of the essential requirements referred to in Article 3(3), points (d), (e) and (f), of that Directive (2021). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv%3AOJ.L_.2022.007.01.0006.01.ENG.
European Union. 2022a. Proposal for a Regulation of the European Parliament and of the Council on horizontal cybersecurity requirements for products with digital elements and amending Regulation (EU) 2019/1020.
European Union. 2022b. Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector and amending Regulations (EC) No 1060/2009, (EU) No 648/2012, (EU) No 600/2014, (EU) No 909/2014 and (EU) 2016/1011 (2022). https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32022R2554.
Floresca, Lauri. 2014. "Cyber Insurance 101: The Basics of Cyber Coverage." *Woodruf Sawyer and Company*. 19 June. https://conferences.law.stanford.edu/cyberday/wp-content/uploads/sites/10/2016/10/3g_CorpEx-DO-Blog-Cyber101-Lauri-061914-2.pdf.
Fruhlinger, Josh. 2019. "Zero Days Explained: How Unknown Vulnerabilities Become Gateways for Attackers." *CSO*, 12 April. https://www.csoonline.com/article/565704/zero-days-explained-how-unknown-vulnerabilities-become-gateways-for-attackers.html.

GDPR. 2018. (n 65) Article 5(1)(f)., 65 § Article 5(1)(f).
GDPR Enforcement Tracker. n.d. https://www.enforcementtracker.com/?insights.
Graeme Smith & Other Claimants and Talktalk Telecom Group Plc Defendant. n.d.
Hedley, Steve. 2002. *Tort*. Butterworths. https://books.google.co.uk/books/about/Tort.html?id=Gyk5AAAACAAJ&redir_esc=y.
Hessen Datenschutzgesetz. 1970. 7 October. 41 § Part I. https://starweb.hessen.de/cache/GVBL/1970/00041.pdf.
ICO. 2018a. "MPN: The Carphone Warehouse Limited". Monetary Penalty Notice. Information Commissioners Office, 8 January. https://ico.org.uk/media/action-weve-taken/mpns/2172972/carphone-warehouse-mpn-20180110.pdf.
ICO. 2018b. "MPN: Bayswater Medical Centre". Monetary Penalty Notice. Information Commissioners Office, 21 May. https://ico.org.uk/media/action-weve-taken/mpns/2258897/bayswater-medical-centre-mpn-20180523.pdf.
ICO. 2020a. "MPN: DSG Retail Limited". Monetary Penalty Notice. Information Commissioners Office, 7 January. https://ico.org.uk/media/action-weve-taken/mpns/2616891/dsg-mpn-20200107.pdf.
ICO. 2020b. "MPN: Cathay Pacific Airways Limited". Monetary Penalty Notice. Information Commissioners Office, 10 February. https://ico.org.uk/media/action-weve-taken/mpns/2617314/cathay-pacific-mpn-20200210.pdf.
ICO. 2022. "MPN: Tuckers Solicitors LLP". Monetary Penalty Notice. Information Commissioners Office, 28 February. https://ico.org.uk/media/action-weve-taken/mpns/4019746/tuckers-mpn-20220228.pdf.
ICO. n.d. "Information Commissioner's Office: Action We've Taken." https://ico.org.uk/action-weve-taken/enforcement/.
ICO. n.d. "Information Commissioner's Office: How We Are Funded." https://ico.org.uk/about-the-ico/who-we-are/how-we-are-funded/.
Indian Government. 2022. Digital Personal Data Protection Bill. § c.9(4). https://www.meity.gov.in/writereaddata/files/The%20Digital%20Personal%20Data%20Potection%20Bill%2C%202022_0.pdf.
Jacobs, Jay, Sasha Romanosky, Benjamin Edwards, Idris Adjerid, and Michael Roytman. 2021. "Exploit Prediction Scoring System (EPSS)." *Digital Threats: Research and Practice* 2 (3): 1–17. https://doi.org/10.1145/3436242.
Leverett, Éireann, Andrew Coburn, and Gordon Woo. 2019. *Solving Cyber Risk*. Hoboken, NJ: Wiley.
Leverett, Éireann, Art Manion, and Matilda Rhode. n.d. "Vuln4Cast Source Code (Version 1.0.0)." https://github.com/FIRSTdotorg/Vuln4Cast/.
Leverett, Éireann, Matilda Rhode, and Adam Wedgbury. 2022. "Vulnerability Forecasting: Theory and Practice." *Digital Threats: Research and Practice* 3 (4): 1–27. https://doi.org/10.1145/3492328.
Mari, Angelica. 2023. "Brazil Issues First Fine for Data Protection Breach". *Forbes*. 11 July. https://www.forbes.com/sites/angelicamarideoliveira/2023/07/11/brazil-issues-first-fine-for-data-protection-breach/.
Massachusetts Government. 2022. "Data Breach Notification Letters October 2022." https://www.mass.gov/lists/data-breach-notification-letters-october-2022.
Mell, Peter, Karen Scarfone, and Sasha Romanosky. n.d. "A Complete Guide to the Common Vulnerability Scoring System." https://www.first.org/cvss/v2/guide.
Miranda, Lucas, Daniel Vieira, Leandro Pfleger de Aguiar, Miguel Angelo Bicudo, Mateus Schulz Nogueira, Matheus Martins, Leonardo Ventura, Lucas Senos, and Enrico Lovat. 2021. "On the Flow of Software Security Advisories." *IEEE Transactions on Network and Service Management* 18 (2).
Miura-Ko, R. A., and N. Bambos. 2007. SecureRank: A Risk-Based Vulnerability Management Scheme for Computing Infrastructures. IEEE Explore. *IEEE*.
Muvija, M. 2021. "Grand National Sets Record for UK Online Sports Betting." *Reuters*, 12 April.
NCSC. 2014. "Cyber Essentials Scheme: Overview." 7 April. https://www.gov.uk/government/publications/cyber-essentials-scheme-overview.
NCSC. 2022. "Cyber Essentials: Requirements for IT Infrastructure". 1 January. https://www.ncsc.gov.uk/files/Cyber-Essentials-Requirements-for-Infrastructure-v3-0-January-2022.pdf.
NCSC. 2023. "Cyber Security Toolkit for Boards." 30 March. https://www.ncsc.gov.uk/collection/board-toolkit.

NCSC. n.d. "Coordinated Vulnerability Disclosure; The Guideline." Accessed 8 July 2022. https://english.ncsc.nl/binaries/ncsc-en/documenten/publications/2019/juni/01/coordinated-vulnerability-disclosure-the-guideline/WEB_Brochure-NCSC_EN.pdf.

Nigerian Government. 2020. Nigeria Data Protection Regulation 2019: Implementation Framework". 1 November. https://nitda.gov.ng/wp-content/uploads/2021/01/NDPR-Implementation-Framework.pdf.

Ormandy, Tavis. 2021. "Google Project Zero". *This Shouldn't Have Happened: A Vulnerability Postmortem*. (blog) 1 December. https://googleprojectzero.blogspot.com/2021/12/this-shouldnt-have-happened.html.

Overseas Tankship (UK) Ltd v Morts Dock & Engineering Co (The Wagon Mound). 1961. AC 388.

PCI Council. n.d. "PCI DSS: V4.0". https://docs-prv.pcisecuritystandards.org/PCI%20DSS/Standard/PCI-DSS-v4_0.pdf.

Risk Metrics Working Group. 2022. "Reporting Cyber Risk to Boards: CISO Edition." 14 March. https://www.eurocontrol.int/sites/default/files/2022-03/reporting-cyber-risk-to-boards-ce-20220322.pdf.

Romanosky, Sasha, and Alessandro Acquisti. 2009. "Privacy Costs and Personal Data Protection: Economic and Legal Perspectives." 24. *Berkeley Technology and Law Journal*. https://www.jstor.org/stable/24118273.

Roncevich, Tim. 2018. "Why Unpatched Vulnerabilities Will Likely Cause Your Next Breach." *Infosecurity Magazine*, 23 May. https://www.infosecurity-magazine.com/opinions/unpatched-vulnerabilities-cause/.

Royal Courts of Justice. 2015. Google Inc. - and - Judith Vidal-Hall Robert Hann Marc Bradshaw - and - The Information Commissioner. 3 March.

Royal Courts of Justice. 2023. Tulip Trading Limited v Wladimir van der Laan, Jonas Schnelli, Pierer Wuille, Marco Falke, Samuel Dobson, Michael Ford, Cory Fields, George Dombrowski, Matthew Corallo, Peter Todd, Gregory Maxwell, Amaury Séchet, Jason Cox, Bitcoin Association for BSV, Eric Lombrozo, Roger Ver.

Sarabi, Armin, Parinaz Naghizadeh, Yang Liu, and Mingyan Liu. 2016. "Risky Business: Fine-Grained Data Breach Prediction Using Business Profiles." *Journal of Cybersecurity* 2 (1): 15–28. https://doi.org/10.1093/cybsec/tyw004.

Shahzad, Muhammad, Muhammad Zubair Shafiq, and Alex X. Liu. 2012. "A Large Scale Exploratory Analysis for Software Vulnerability Life Cycles." *IEEE*. https://ieeexplore.ieee.org/document/6227141/authors#authors.

Solove, Daniel J., and Woodrow. Hartzog. 2022. *Breached!: Why Data Security Law Fails and How to Improve It*. Oxford: Oxford University Press. https://scholarship.law.gwu.edu/cgi/viewcontent.cgi?article=2857&context=faculty_publications.

Talesh, Shauhin A. T. 2018. "Data Breach, Privacy, and Cyber Insurance: How Insurance Companies Act as 'Compliance Managers' for Businesses." *Law and Social Inquiry*, 27 December. https://www.cambridge.org/core/journals/law-and-social-inquiry/article/abs/data-breach-privacy-and-cyber-insurance-how-insurance-companies-act-as-compliance-managers-for-businesses/1A10E0F87EB1C205EEA43AB4E8270FB2.

UK Government. 2019. "Product Safety Advice for Businesses." 29 March. https://www.gov.uk/guidance/product-safety-advice-for-businesses.

UK Government. 2021. "Product Security and Telecommunications Infrastructure (PSTI) Bill: Factsheets." 24 November. https://www.gov.uk/government/collections/the-product-security-and-telecommunications-infrastructure-psti-bill-factsheets.

UK Parliament. n.d. "UK Parliament: Consolidated Fund." https://www.parliament.uk/site-information/glossary/consolidated-fund/.

US Government. n.d. "National Vulnerability Database." https://nvd.nist.gov/.

Wikipedia. n.d. "Patch Tuesday". https://en.wikipedia.org/wiki/Patch_Tuesday.

Wolf, Nicky. 2016. "DDoS Attack That Disrupted Internet Was Largest of Its Kind in History, Experts Say." *The Guardian*, 26 October. https://www.theguardian.com/technology/2016/oct/26/ddos-attack-dyn-mirai-botnet.

Woods, Daniel, and Andrew Simpson. 2017. "Policy Measures and Cyber Insurance: A Framework." *Journal of Cyber Policy*, https://doi.org/10.1080/23738871.2017.1360927.