(Under) Investment in Cyber Skills and Data Protection Enforcement: Evidence from Activity Logs of the UK Information Commissioner's Office*

Pantelis Koutroumpis† Farshad Ravasan‡ Taheya Tarannum§

Abstract

Data breaches account for a significant share of cyber attacks. While they severely impact customers, who lose valuable personal data, they often have a limited effect on the operations of data-holding companies. This might lead firms to underinvest in cybersecurity. Does stronger data protection alleviate the effects of these misaligned incentives? We address this question by examining the link between firms' cybersecurity hiring and stronger data protection laws and enforcement. We study two institutional changes that affect data protection enforcement by the Information Commissioner's Office (ICO) in the UK. The first is the removal of the requirement to prove substantial damage and distress in 2015 that gives greater discretion to the ICO to issue monetary penalties. The second is the enactment of the Data Protection Act 2018, which significantly raises the ceiling of monetary penalties. To examine the effects of these legal changes, we assemble a novel dataset from ICO activity logs that entails more than 5,000 supervisory actions. We construct an index for exposure to ICO enforcement at the three-digit industry level. Combining the sectoral variation with the timing of the legal changes, we show that while stronger data protection enforcement significantly increases the investment in cybersecurity skills by up to 52%, it has a negative impact on firm dynamics, reducing the entry rate by up to 12% and increasing the exit rate by up to 13%.

JEL Classifications: G31, G38, J23, J24, K20, K24
Keywords: Cybersecurity, Skill Acquisition, Data Protection, GDPR, Law Enforcement, Data Intensity, Cash Holding, Firm Dynamics

*We wish to acknowledge funding support from Citi group.
†Oxford Martin School, University of Oxford, 34 Broad Street, Oxford, OX1 3BD, United Kingdom. E-mail: pantelis.koutroumpis@oxfordmartin.ox.ac.uk
‡Oxford Martin School, University of Oxford, 34 Broad Street, Oxford, OX1 3BD, United Kingdom. E-mail: farshad.ravasan@oxfordmartin.ox.ac.uk
§Oxford Martin School, University of Oxford, 34 Broad Street, Oxford, OX1 3BD, United Kingdom. E-mail: taheya.tarannum@oxfordmartin.ox.ac.uk

1 Introduction

Cybercrimes have become a significant source of risk for corporations. The frequency of cybercrimes and their costs have rapidly increased over the past few years. Yet the growth in firm expenditure for cybersecurity has been slow. One explanation for this lack of spending lies in the agency problem faced by firms investing in cybersecurity. Firms have a complex nexus of stakeholders including customers, employees, and suppliers who are affected by their exposure to cyber risk. However, the cost of a cyberattack is not evenly distributed between firms and their stakeholders. Data breaches, which account for a significant share of cyberattacks, gravely affect customers and employees who are losing valuable personal data. However, these incidents often have a limited impact on business operations. This misalignment of incentives can lead to underinvestment in cybersecurity.

Data protection regulations provide critical legal frameworks in redistributing the cost of cyberattacks between firms and their stakeholders. Stronger data protection laws or increased enforcement could compel firms to internalize cyber-related social costs. As a result, firms will improve their cyber-defense and reach closer to the socially optimum level of cybersecurity investment.

In this paper, we examine whether stronger data protection laws and their enforcement lead firms to invest more in cyber skills. Acquiring IT security talents is considered to be the most effective strategy by corporations to protect themselves against cyberattacks.[1] However, a large proportion of UK firms have been lacking the right skills (technical, incident-response, and governance) that are necessary to manage cybersecurity risks.[2] To draw a causal link between data protection enforcement and firms' cybersecurity hirings, we study two legal reforms affecting the code of practice under which the UK Information Commissioners' Office (ICO) operates. These legal changes significantly increased both the frequency and amount of monetary penalties that have been issued by the ICO in data protection cases. The first legal reform we study is the removal of the requirement to prove substantial damage

[1] According to the Cyber Security Breaches Survey 2020, 90% of the cybersecurity incidents can be attributed to human errors and lack of cyber skills. Moreover, Sophos surveyed 119 financial services establishments that were not hit by ransomware in the previous year and do not expect to be hit in the future, and asked their IT managers, "why do you not expect your organization to be hit by an attack in the future?" The top reasons for this confidence are having trained IT staff capable of preventing attacks (66%), followed by running a full Security Operations Center (SOC) (55%).

[2] The DCMS Cyber Security Breaches Survey 2020 shows 653,000 businesses (48%) have a basic skills gap. That is, people in charge of cybersecurity in those businesses lack the confidence to carry out the basic tasks, such as setting up configured firewalls, storing or transferring personal data, and detecting and removing malware. Around 30% of the businesses have more advanced skill gaps in areas such as penetration testing, forensic analysis, and security architecture. Around 27% of businesses have a skill gap when it comes to incident responses. Moreover, the survey highlights that the firms in the cyber sector report that a third (35%) of their vacancies are hard to fill.

and distress (SDD hereafter) in 2015. The reform, which is an amendment of the Privacy and Electronic Communications Regulations (PECR), gave greater discretion to the ICO to issue monetary penalties. The second is the enactment of the Data Protection Act (DPA) 2018, implementing the General Data Protection Regulation (GDPR) in the UK. The act significantly raised the ceiling of monetary penalties from half a million to £17.5 million or 4% of annual global turnover, whichever is greater.

It is empirically challenging to examine the economy-wide effect of data protection regulations on firms' hiring. Previous literature mainly focuses on specific firms or sectors that are presumably more sensitive to consumer data regulations. However, there has been no systematic approach to define the exposure to data protection laws across the economy. To overcome this challenge, we assemble a novel dataset of ICO activity logs between 2012 and 2018. Our data includes a wide range of supervisory actions from soft enforcement, such as advice, to more severe sanctions, such as civil monetary penalties. This helps us to construct an index at SIC three-digit sector classification revealing a wide variation in firm exposure to ICO enforcement across different industries. Combining this sectoral variation with the timing of the legal changes allows us to isolate the effect of the legal reforms from other contemporaneous shocks.

We begin by documenting the macro trend that demonstrates the divergence in the share of job postings requiring cyber skills across sectors with high and low exposure to ICO enforcement. The divergence has been triggered by a surge in cyber skill demand among highly exposed industries. It starts after the removal of the SDD clause and notably accelerates after the DPA 2018 comes into effect. We then examine this link within the boundary of local labor markets using 228 travel to work areas (TTWAs) across the UK.[3] The TTWA-level analysis allows us to control for region-specific temporal shocks – by adding TTWA-year fixed effects – and regional industrial characteristics – by adding TTWA-sector fixed effects. We draw inferences from the difference-in-differences (DID) design where we compare the sectors in the same TTWA with high and low exposure to ICO enforcement before and after the legal changes. We estimate a rise in the share of cyber job postings by 26% between 2016 and 2018 after SDD removal for highly exposed sectors. The effect increases up to 52% during 2019 and 2020 after the enactment of DPA 2018.

Thus far, our results highlight the aggregate impact of stronger data protection law and enforcement on local labor markets. We next examine how these legal changes affect the

[3]TTWAs are geographical units that reflect self-contained commuting areas in which at least 75% of the population both work and live.

demand for cyber skills at the firm level. We study the skill requirements for more than 4,000 firms between 2012 and 2020. We exploit a panel at the firm-TTWA-year level with 273,488 observations. The firm-level data allows us to control for local shocks when firms hire across different regions. Our results show a rise in cyber skill demand across different TTWAs by 37% after the SDD removal and demand increases up to 73% after the DPA 2018 comes into effect.

We further study the heterogeneous effects on firms according to their data intensity and digital technology portfolios. We exploit the granular information on required skills in job postings. This allows us to evaluate firms' investment strategy across different technologies before the legal changes (Tambe and Hitt, 2012a,b; Goldfarb et al., 2022).[4] We classify digital technologies into two broad groups. The first group is the technologies that support firms' digital organization. Investing in skills related to these technologies implies that firms hold valuable data and actively spend on data management and storage. The second group is data harvesting technologies ranging from artificial intelligence (AI) and big data to business intelligence and data mining. This group of technologies characterizes the firms which actively invest to generate value from data processing (Abis and Veldkamp, 2020).[5]

Using principal component analysis, we aggregate orthogonal variations in firms' demand over the aforementioned digital skills. This allows us to construct an index of data intensity at the firm level. We find that the rise in cyber skill demand is much stronger among the subsample of data-intensive firms that are above the median value of our index. The results become slightly stronger when we construct our index using only data harvesting skills. Importantly, we do not observe any significant changes in skill demand for less data-intensive firms following the SDD removal or the DPA 2018. We also compare the investment in different subclasses of digital skills. We find that firms investing in cloud technologies before the legal changes are more likely to acquire cyber talents afterward. Our analysis of labor market adjustments to digital technology adoption also reveals that cloud solutions and cybersecurity show a strong similarity along their principal components.

We also study how firms' liquidity affects the response to the legal changes. The availability of internal liquidity is crucial to invest in firms' key human capital (Oi, 1962; Shapiro, 1986;

[4]This literature suggests that skill demand provides an early indicator of firms' intention to engage with technologies compared with other measures such as patents.

[5]These two broad classes of digital technologies roughly mirror the two types of data-related skills in Abis and Veldkamp (2020). The first group of skills belongs to data managers who make raw data usable and the second group belongs to data analysts who produce knowledge.

Dixit, 1997; Baghai et al., 2021; Brown and Matsa, 2016; He, 2018). Hence, we expect firms' ability to acquire new skills to vary by its liquidity constraint. Our results highlight a pronounced and steady rise in cyber skills demand among firms that have high cash holding before the legal changes, leading to the divergence in cybersecurity hiring between firms with a high and low level of cash holding.

Furthermore, we examine the effect of data protection laws over a firm's life cycle. We find that the increase in demand for cyber professionals is stronger among younger firms that often face more difficulties in recruiting skilled employees (Israelsen and Yonker, 2017). This result implies that data protection enforcement might have unintended negative effects on firm dynamics. To test this hypothesis, we use the UK business register which provides the universe of firms that entered the market between 2013 and 2020. We compare the entry and exit rate to the industries with high and low exposure to ICO enforcement. We find that the entry to high exposure industries slowed down significantly by up to 1.4 percentage points after the legal changes. Similarly, the exit rate increased by up to 0.9 percentage points. Considering that the average entry and exit rates between 2013 and 2020 respectively were 12% and 7%, these effects roughly amount to 12% change in firms' entry and 13% change in exit rates. This highlights the trade-off of stronger data protection laws and enforcement – although they are effective in increasing investment in cybersecurity, they come at the cost of slowing down the firm creation and business dynamics.

Contributions to the Literature

Our paper answers a question at the intersection of two main strands of literature, namely corporate investment in cybersecurity and the economic impact of data protection. The first strand of literature that our paper contributes to is the study of underinvestment in cybersecurity. This literature shows that, despite its importance, firms' cybersecurity expenditure is often either insufficient or inefficient (Kankanhalli et al., 2003; Gordon et al., 2015a,b). Recent papers point to the potential market failure in the provision of cybersecurity at a socially optimum level (Kopp et al., 2017). This arises from a misalignment between social cost and firms' private cost of cyberattacks (De Cornière and Taylor, 2021). The misalignment is more pronounced in cases of data breach incidents. The reason is that data breaches severely impact customers, who lose valuable personal data, but often have a limited effect on the operations of data holding companies. (Crosignani et al., 2021; Bana et al., 2021).

An extensive line of research has examined the effect of data breach incidents on firm value (Hilary et al., 2016; Johnson et al., 2017; Amir et al., 2018; Richardson et al., 2019;

Florackis et al., 2020; Kamiya et al., 2021). The empirical evidence indicates that the economic consequence of data breaches substantially varies and is often small. The paper closest to ours is Bana et al. (2021), which studies firm hiring response to data breaches in the US. They provide evidence that firms' investment in cybersecurity skills after data breach incidents is limited. That further lends support to the idea that firms lack incentives to provide the socially optimum level of cybersecurity.

We contribute to this literature in three ways. First, our paper is the first that examines the effect of stronger data protection laws and enforcement on firms' investment in cyber skills. We evaluate the effectiveness of two policy instruments of the UK data protection laws, namely more frequent monetary penalties and mega-fines. We find both instruments are effective in increasing investment in cyber skills. While the impact of mega-fines is substantially stronger, it is less homogeneous across firms with different characteristics such as age, liquidity, and digital technology portfolios. In this regard, we also contribute to a broader literature that compares the response to the extensive margin of monetary penalties, that is, receiving a fine, with the intensive margin or receiving a higher fine (Dušek and Traxler, 2022).

Second, there are recent papers that study the effect of data breaches on firms' liquidity management (Boasiako and Keefe, 2021; Garg, 2020). This literature reveals an important insight that firms keep a high balance of cash holding after they (or their peers) are targeted by cyberattacks. Boasiako and Keefe (2021) argues that the precautionary behavior of firms stems from covering potential costs related to litigation, reputation damage, and customer loss. Our results shed light on an alternative channel. We show that the availability of cash is a prerequisite for firms' ability to invest in cyber skills. Thus, the cost of and recruitment difficulties in acquiring cyber talents might be another reason that leads firms to keep a higher balance of cash holding.

Third, despite the effectiveness of a stronger data protection environment to increase investment in cybersecurity, it comes at a cost. We show that data protection laws and their enforcement drastically hamper firm creation and slow down business dynamics. In this regard, our paper also contributes to the second strand of literature on the unintended economic effects of data protection (Goldberg et al., 2019; Abis et al., 2022; Schmitt et al., 2020; Zhao et al., 2021; Chen et al., 2022; Buckman et al., 2021; Mayya and Viswanathan, 2021; Johnson et al., 2021).

Our contribution to this literature is fourfold. First, prior works focused on the GDPR[6] and CCPA[7] that alter the content of law strengthening data protection rights. Our paper is the first to study a legal change – removal of the SDD clause – that enhances law enforcement for data protection cases.[8] This allows us to disentangle the effects of changes in enforcement and the content of the law. Second, the previous literature focuses on data privacy and mandatory consent as the main channel. According to this literature, GDPR or CCPA adversely affects firms by limiting their access to valuable personal data. We instead shed light on data security as an important channel. Data protection laws impact a broader set of companies that hold personal data by increasing the cost of data breaches. Considering that the highest data protection fines issued in the UK (and elsewhere) are often related to data breaches, it is surprising that this channel has been overlooked by the literature. Third, using the granular skill-level information from job postings, we construct a comprehensive index for firm-level data intensity. Fourth, there has been no systematic approach in previous literature to define the exposure to data protection enforcement across different sectors. In this regard, we make an important contribution by assembling a novel dataset based on more than 5,000 ICO supervisory actions, such as guidance, improvement action plans, and monetary penalties. This data allows us to construct a comprehensive data protection exposure index at the three-digit industry level and track which firms are more exposed to ICO enforcement.

The remainder of this paper proceeds as follows: Section 2 discusses the institutional details and the legal changes related to ICO's operation, section 3 describes the data, section 4 discusses our estimation strategy, Section 5 presents the results, and Section 6 concludes.

2 The ICO: Institutional Setup

This section provides an overview of the institutional setting of the ICO which is the data protection authority of the UK. Section 2.1 discusses the legal framework which gives the ICO regulatory power. Section 2.2 provides details on what regulatory actions they can take. The last section highlights the major changes in the data protection regulatory regime.

2.1 Legal Framework

The ICO is the regulatory body in the UK that oversees the secure use of data by organizations. It was established in 1984 to uphold the information rights of the public. The main

[6] GDPR refers to General Data Protection Regulation.
[7] CCPA refers to California Consumer Privacy Act.
[8] SDD clause refers to the requirement to prove substantial damage and distress.

functions of the ICO are to provide compliance guidance to the organizations, process data protection and freedom of information related complaints, and take regulatory actions in case of a violation of the laws. While the ICO enforces a number of legislations[9], the regulatory power to deal with the data protection complaints mainly derives from the following ones:

Data Protection Act (DPA) 1998 and 2018: The Data Protection Act (DPA) is the main legislation that upholds information rights in the UK. It was originally introduced by the DPA 1984 and consequently had been replaced by the DPA 1998 and DPA 2018. It establishes a balance between the use of personal data by businesses (or other organizations) and the privacy rights of individuals (i.e., data subjects). The law states eight data protection principles to ensure good information handling practices.[10] It gives statutory power to the Information Commissioner to enforce the data protection legislation. It also specifies the responsibility of the data controllers, i.e, those in charge of processing personal data, and requires them to register with the ICO (ICO, 2012).

The DPA 2018 updates the data protection framework to make it more relevant to the digital era and implements the European General Data Protection Regulation (GDPR). The law was passed in April 2016 and came into effect on May 25, 2018. Like the previous data protection laws, the DPA 2018 governs the use of personal data, gives certain rights to individuals, and dictates how organizations should process personal data. The act reinforces the Information Commissioner's role as the data protection supervisory authority and extends the ceiling of maximum fine (see Section 2.3). Following its exit from the EU, the UK has also enacted the UK GDPR, which is in effect since January 1, 2021 and is based on the EU GDPR. Together the DPA 2018 and the UK GDPR provide the data protection framework for UK organizations.

The Privacy and Electronic Communications Regulations (PECR) 2003: The

[9]In addition to the laws discussed in the text, the ICO is also responsible for upholding The Network and Information Systems (NIS) Regulations 2018, The Investigatory Powers Act (IPA) 2016, The Electronic Identification and Trust Services for Electronic Regulations (eIDAS) 2016, The Re-use of Public Sector Information (RPSI) Regulations 2015, The Infrastructure for Spatial Information in the European Community Regulations (INSPIRE) 2009, The Environmental Information Regulations (EIR) 2004, The Enterprise Act 2002, and The Freedom of Information Act (FOIA) 2000.

[10]These principles are: (1) Personal data must be fairly and lawfully processed; (2) It should be processed only for the lawful purposes it is collected; (3) Personal data collected should be adequate, relevant, and not excessive; (4) It should be accurate and up to date (5) It should not be kept longer than necessary; (6) It should be processed in accordance with the rights of individuals; (7) Appropriate measures should be taken against unauthorized or unlawful processing and accidental loss or destruction of personal data; (8) It should not be transferred to countries outside the European Economic Area unless there is an adequate level of protection.

PECR protects the privacy of individuals regarding electronic communication. It follows the EU Privacy and Electronic Communications Directive (the ePrivacy Directive) 2002. The law applies to marketing communications (calls, texts, faxes, or emails) and the use of cookies for tracking purposes. It complements the DPA and other data protection laws by giving rights to individuals to refuse unsolicited marketing communications. For example, it dictates that the recipients of unsolicited marketing emails should be able to opt out of receiving such emails, and individuals or businesses can register with the Telephone Preference Service (TPS) to stop receiving unsolicited marketing calls. It also requires the service providers to safeguard the public electronic communications services (ICO, 2018). The law was amended multiple times including as recently as 2018.

2.2 How the ICO Handles Data Protection Complaints

Figure 1 shows a simplified flow of actions that the ICO follows to process data protection complaints. When the ICO receives a complaint, it launches an investigation to see if any violation has occurred. The initial concern could come from the public as anyone could raise a concern with the ICO about their information rights. It could also come from whistle-blowers, media reports, or self-reporting. When the ICO receives a complaint, it assigns a case officer to oversee it. The case officer launches an investigation to determine (1) whether any data breach has occurred and (2) in the case of a breach, what actions to take (ICO, 2014). If the case officer concludes that no law was breached, they will close the case. If a breach has occurred, the case officer will decide on actions depending on the severity of the breaches. They can choose from a wide range of regulatory tools as discussed below.

2.2.1 Supervisory Actions through the ICO Activity Logs

Every complaint that the ICO receives enters into its activity logs. Complaints related to data protection usually fall either under the DPA or PECR. After the case is complete, a case outcome will be filed in the record. The ICO has a number of regulatory tools at its disposal as well as some non-statutory measures that apply to these cases.

The ICO can exercise its power through different types of notices. Upon launching an investigation, it can issue an 'information notice' asking for more information regarding the data protection practices of an organization. It can also serve an 'assessment notice' which requires the data controllers to provide an assessment of their compliance practices. The information notice or assessment notice by itself is not the outcome of a case. The ICO uses them to determine whether there is any lack of compliance. However, any failure to

comply with these notices is subject to fixed penalties. For serious or repeated violations of the law or for the failure to comply with its prior notices, the ICO can serve a penalty notice which charges a person or organization with a specific amount of monetary penalty.[11] In these cases, the ICO often serves an 'enforcement notice' that outlines what steps should or should not be taken by the responsible parties. Organizations' failure to comply will be followed by monetary penalties.

The ICO can also engage with data controllers by asking them to produce plans to improve their information rights practices. This is not an application of the statutory regulatory power of the ICO but is done to promote compliance. To observe compliance with data protection principles and good practices, the ICO can also audit an organization. A compulsory audit usually follows an assessment notice. An audit can also be consensual in which case it comes as a recommendation from the ICO to ensure good practices and is voluntarily agreed by the organization. Most of the ICO cases ending up in audit are against public organizations.

The ICO may also decide to limit its response to very soft actions. For example, the ICO case officers may find that the organization has come short in its dealing with the customers' data complaints and could improve its practices. In these cases, they may end up offering advice or suggesting that the data controller takes a one-off action.

In Figure 1, we show these outcomes in three groups: strict, soft, and very soft outcomes. The first one refers to monetary penalties and enforcement notices, which apply to severe cases of violation. We refer to improvement action plans as soft outcomes because these are non-punitive measures and designed to promote compliance through working with the organization. We categorize advice and one-off actions as very soft outcomes which apply to minor violations.

2.3 A Timeline of the Changes in the Data Protection Regime

Figure 2 shows the timeline of the major changes in the rules and regulations enforced by the ICO. Prior to 2010, the ICO had limited power – its actions were limited to serving enforcement notices and undertakings. On April 6, 2010, the ICO gained the power to issue a civil monetary penalty (CMP) of up to £500,000 for serious breaches of the DPA. The 2011 amendment extended the same power to the PECR-related breaches. In the post-2010

[11] Note that in this paper, we use the terms monetary penalty, civil monetary penalty (CMP), or fine interchangeably.

era, we focus on two periods that are distinct in terms of the regulatory environment. The first period follows the 2015 PECR amendment, which demonstrates an increased number of CMPs imposed by the ICO, and the second is the period following the introduction of the DPA 2018.

(i) Increased Number of Penalties (Post-SDD): Between April 2010 and May 2018, the maximum fine for DPA or PECR cases was £500,000. However, two changes occurred that increased the intensity of enforcement since 2015. The first is that the requirement to prove 'substantial damages and distress' (SDD hereafter) for PECR cases was relaxed by a 2015 amendment. Before that, the ICO had to prove the violation caused substantial damage and distress to levy fines. But the removal of the SDD requirement means that the ICO only needs to be convinced that there has been a serious breach of PECR to serve a monetary penalty notice. The second change during this period is the passage of the DPA 2018 (also EU GDPR) in 2016 (Figure 2).

(ii) Increased Penalty Ceiling (Post-2018): The maximum fine increased substantially under the DPA 2018. Under the latest rule, an organization could be charged for a DPA violation an amount of up to £17.5 million or 4% of its annual global turnover, whichever is higher. The law also makes it mandatory for the data controllers to notify the ICO as well as affected individuals about personal data breaches within 72 hours of becoming aware of a breach.

These two periods provide us with different quasi-experimental setups to evaluate the hiring response of firms due to increased regulatory actions and the strengthening of the data protection laws. While the first increases the ability of the ICO to issue CMPs, the second relies on mega-fines to achieve compliance. Figure 3 shows how ICO enforcement for business organizations changed over these periods. We plot the number of firms per quarter that were targeted by different types of ICO actions. Before 2015, improvement action plans were a dominant supervisory instrument through which the ICO worked with organizations to bring them closer to compliance. Following the removal of the SDD requirement in 2015, there was a surge in monetary penalties and enforcement notices which are the most strict types of enforcement. Nevertheless, they gradually declined after the passage of the DPA 2018 which substantially increased the maximum fine in data protection cases. During the last period, the ICO began to rely more on one-off actions and advice to deal with data protection compliance.

3 Data

In this section, we describe how we construct our local and firm-level datasets to study the demand for cyber skills.

3.1 Job Postings Data

We use the Burning Glass Technologies (BGT) data for the UK from 2012-2020. BGT provides the near universe of online job postings with 6-9 million observations annually. These job postings are scraped on a daily basis from online job boards or company websites. Each entry is parsed and deduplicated. The resulting dataset provides a number of variables for each job posting, such as date of a job posting, location, Standard Occupation Classification (SOC) code 2010, UK Standard Industry Classification (SIC) 2007, employer's name, and education level. Additionally, BGT also maps each job description to a number of skill categories developed by its own skill taxonomy. We use these skills to identify cyber job postings. Figure 4 shows the full list of cyber skills that we consider. Specifically, we flag a job as a cyber job if it requires at least one skill from the following skill clusters: information security management, anti-malware software, web security, application security, cyber engineering, and network security.

For the first part of our analysis, we use job postings with non-missing industry codes and locations. We use travel to work areas (TTWA) as a geographic unit to study how skill demand changes across labor markets in reaction to the laws. In the UK, TTWAs are self-contained commuting areas within which at least 75% of the population both work and live. We aggregate the number of cyber jobs by three-digit SIC industry codes (i.e., SIC group) and match them with the sectoral index of data protection intensity (see Section 3.2). Our final sample includes 143 3-digit industries from 224 TTWAs between 2013 and 2020.

For our firm-level analysis, we work with a smaller subset of job postings data with 28% observations that have non-missing employer name.[12] We keep firms with at least 100 job postings during the period 2012-2020. We standardize the firm names and use a fuzzy matching to link with other firm-level datasets discussed in the next section. We manually validate the matched firms and cross-check industries. This gives us around 4,000 firms. We profile these firms based on their digital technology portfolio and data use as reflected in their skill demand prior to the legal changes. We consider two broad groups of digital technologies: digital organization technologies and data harvesting technologies.

[12] This is around 18.3 million job postings.

A number of papers have used the BGT data to study skill demand and labor market dynamics.[13] The main limitation of the data, as noted by Hershbein and Kahn (2018), is the over-representation of occupations and industries with higher skill requirements. But the authors find that the distribution of job postings is stable over time and the industry trends closely track other official sources of US labor market statistics. For the UK BGT data, Javorcik et al. (2020) compares the monthly job postings with the estimated number of vacancies from the UK Vacancy Survey (UKVS). They conclude that BGT covers around 86% of the vacancies in the UK between 2012 and 2019. For our purpose, the over-representation of high-skilled jobs in the BGT data proves to be useful since it provides sufficient observations to capture cyber jobs, which constitute a small share of all job postings.

For our firm-level analysis, we supplement the job posting data with other firm-level information. We use the Orbis database for the UK from Bureau van Dijk (BvD). Orbis provides financial information from firms' balance sheets or income statements. We use firms' 2015 level of cash holdings as a measure of the liquidity constraint faced by firms prior to the legal changes. We also use the age of these firms to study how older vs. younger firms behave. To examine the impact on a firm's life cycle, we construct measures of entry and exit by firms at the three-digit industry level. Using the universe of UK businesses from the Companies House data, we identify all the firms that started or ended operation between 2013 and 2020.

3.2 ICO Data Protection Case Outcomes

We use data protection complaints from the ICO activity logs to create a sectoral index of ICO enforcement exposure. We use the cases between 2012 and the second quarter of 2018, which is the period before the DPA 2018 came into effect. We consider the cases brought against business organizations and match their names with the UK business register to retrieve industry codes.[14] In total, we have 5,783 cases that ended in regulatory actions against 1,819 firms. Table 1 shows the average number of cases per year for each type of action. Among statutory actions, we have monetary penalties (5 cases per year) and enforcement

[13] Hershbein and Kahn (2018) is the first paper to use the US data to examine whether the Great Recession accelerated the adoption of labor-saving technologies. Among other papers using the data, Deming and Kahn (2018) study heterogeneity in skill demand and returns to high-level skills, Deming and Noray (2020) examines how the changing skill requirement of STEM jobs affect earning dynamics, Azar et al. (2020) measures labor market concentration, and Forsythe et al. (2020) studies the impact of COVID-19 on vacancies. Among the papers using the BGT data at the firm level, Babina et al. (2020) match the BGT data with Compustat to measure firm-level investment in AI-related human capital. Alekseeva et al. (2021) also uses the Compustat matched data to study the characteristics of AI-adopting firms. On the other hand, the UK data has been used by Javorcik et al. (2020) to study the impact of Brexit and Adams-Prassl et al. (2020) to study flexible work arrangements in high- and low-wage jobs.

[14] We use information from Companies House which is the registrar of all UK companies.

notices (3 cases per year). We also have soft enforcement such as improvement action plans (11 cases per year) and very soft outcomes: one-off actions (478 cases per year) and advice (329 cases per year). In our data, the largest number of firms facing ICO enforcement is in the financial sector, followed by wholesale and retail trade and administrative and support services activities. Note that for the health sector we only include the private firms in our sample.

To capture the sectoral variation in ICO regulation, we count the number of firms that are subject to ICO enforcement per three-digit industry code. This gives us an industry-level measure of exposure to ICO enforcement. We match the exposure index with our job posting data and classify the three-digit industries by the intensity of exposure. Figure 5 shows the percentage of three-digit industries within each sector that falls into different quartiles of the exposure index. The top quartile, shown in red, shows the share of three-digit industries with high exposure to ICO enforcement. As the figure shows, most of the service sectors, as well as trade, have industries both at the high and low end of the exposure index. A large part of the finance and insurance sector and real estate sector is heavily regulated while the manufacturing sector is least regulated.

Figure 6 shows how the demand for cyber jobs changes nationally for the high- and low-exposure industries. The top panel reports the share of cyber jobs normalized by their values in 2016. The dashed lines mark the periods of interest: SDD removal in 2015 and the enactment of the DPA 2018 (GDPR) in 2018. Until 2015, the demand for cyber skills was growing at the same pace across both types of industries. In 2016 and 2017, after the removal of the SDD requirement, the high exposure industries show higher demand for cyber skills which grows exponentially after 2018. The bottom panel (Panel b) shows the share of cyber jobs without normalization. This figure shows that the low-exposure industries initially have a higher demand for cyber skills compared to the high exposure sectors. The gap narrows after 2016 with a complete reversal of the trend at the end of the period. The next section outlines our strategy to estimate the skill response to legal changes using local and firm-level data.

4 Empirical Strategy

We measure how increased enforcement and a higher penalty ceiling affect firms' investment in cybersecurity skills. We first examine the effects on cyber skill demand across labor markets (i.e., TTWAs). We apply a difference-in-differences (DID) design combining the

sectoral variation in ICO exposure with the timings of legal changes. We calculate cyber job postings as a share of all job postings annually for three-digit SIC industries of each TTWA in our sample. This gives us a measure of industry-specific demand for cybersecurity talents in local markets. The following equation shows our empirical specification:

$$cyber\_share_{cjt} = \beta_1 high\ ico\ exposure_j \times increased\ enforcement_t + \beta_2 high\ ico\ exposure_j \times increased\ penalty_t + \delta_{ct} + \rho_{cj} + \epsilon_{cjt}$$

where c denotes TTWA, j denotes three-digit SIC industries, and t denotes the year. The variable *increased enforcement* is a binary indicator for the period 2016-2018 which follows the removal of the SDD requirement in 2015. Similarly, *increased penalty* is an indicator for 2019-2020, which is a period following the DPA 2018 raising the penalty ceiling higher. The variable *high ico exposure* denotes the three-digit industries at the top quartile of ICO enforcement exposure index (see Figure 5). For our local labor market regressions, we include TTWA times year fixed effects (FEs) (δ_{ct}) which controls for TTWA-specific temporal shocks in cyber skill demand. We also include TTWA times three-digit industries FEs (ρ_{cj}) which controls for TTWA-specific industry characteristics. We run a weighted least square regression using the share of TTWA level job postings as weights. We use a two-way clustering of standard error at the three-digit industry and year level to account for possible correlation of residuals across both dimensions (Miller et al., 2009).

Our identification strategy relies on the high exposure and low exposure industries being comparable. To check whether this assumption holds prior to the legal changes, we estimate the dynamic model which includes interaction terms of the high exposure industries and yearly dummies. Another threat to our identification could come from other temporal or sectoral shocks that might coincide with the legal changes. One such known shock is the Brexit referendum which affects the total volume of online job postings in certain local labor markets (Javorcik et al., 2020). To deal with the Brexit-related contraction of the UK job market, we use the share of cyber jobs (out of all TTWA-industry job postings or TTWA-firm job postings) for all our specifications. Also, our fixed effect models control for labor market specific temporal shocks. Finally, we report our results only for the service sectors which gives us a more homogeneous sample.

We check the robustness of our results using alternative exposure indices. Our main model uses the exposure index based on the count of firms subject to ICO actions across

three-digit industries. Our second measure of the exposure index uses the count of ICO actions. If firms in some industries face multiple actions within our considered period, those industries will get higher weight by this approach. Our third measure of exposure index weighs different types of ICO cases (e.g., monetary penalty, enforcement notice, or one-off actions) by the inverse of their share among all types of actions. This approach puts more weight on monetary penalties and enforcement notices which are severe in nature but fewer in numbers. On the other hand, advice or one-off actions, which are soft enforcement, get less weight because of their higher share among all types of actions. We extend our firm-level analysis to examine the heterogeneous in skill demand.

We estimate the following DID regression using our firm-level matched data.

$$cyber\_share_{icjt} = \beta_1 high\ ico\_exposure_j \times increased\ enforcement_t + \beta_2 high\ ico\_exposure_j \times increased\ penalty_t + \delta_{ct} + \mu_i + \varepsilon_{icjt}$$

In this regression, i stands for firms and firms have multiple plants across different TTWAs. The dependent variable shows the share of cyber job postings out of all job postings advertised by a specific plant in a year. Our model controls for firm-specific shocks (μ_i) as well as TTWA times year FEs. Standard errors are two-way at the three-digit industry and year level. Similar to our local labor market regressions, we also estimate dynamic models to see how firms' responses change over the years. We compare our results across subsamples of firms with different technology portfolios, cash holding, and demographics (i.e., age). Table 1 reports summary statistics for the key variables of our analysis. The weighted mean of the share of cyber job postings per TTWA-industry-year is 1.15%. The share for the firm is smaller with an average of 0.14% cyber job postings annually. In the table, we report firm characteristics for the year prior to the first change we study. We calculate cash and cash equivalent securities as a share of the total assets of the firm. The average firm in our sample has 24% of total assets as cash or cash equivalent and is 19 years old. The table also reports birth and death rates of firms per TTWA-industry-year. The average TTWA-industry in our sample has 12% birth rate 7% death rate per year.

5 Results

5.1 Local Labor Market Results

We first examine how the legal changes affect cybersecurity hirings across local labor markets. Table 2 reports the results using our TTWA-industry (3-digit)-year panel. Since the size of local labor markets as well as industries varies widely, we use the share of cyber job postings out of all job postings as our dependent variable. In Table 2, we report the average response to the removal of the SDD requirement and increased penalty. Our baseline specification (Column 1) shows that the share of cyber job postings increases by 0.26 percentage points in the high exposure industries following the SDD removal. Compared with the sample mean in 2015, the coefficient translates to a 26% increase. For the increased penalty ceiling, the response is double in magnitude with the share of cyber job postings rising by 52%. We estimate these effects by controlling for TTWA-year fixed effects and TTWA-industry fixed effects. In Figure 5, we see that the manufacturing sector has low exposure to ICO actions, while the service sectors seem to have a more uniform distribution across high- and low-exposure industries. Hence, we produce our results using only the service industries. Column 3 shows that the effect of the SDD removal is less robust to the exclusion of the manufacturing sector. Although the effect of an increased penalty is smaller (42%), it appears to be robust. The last two columns of the Table check our results using alternative measures of the exposure index. Our estimates are robust across both measures. We find a larger response for both increased enforcement and increased penalty when we put more weight on monetary penalties and other strict actions. We conclude from these results that while both the SDD removal and mega-fines induce firms to invest more in cyber skills, the response appears to be stronger for the latter one.

Figure 7 plots the yearly response of cybersecurity hirings before and after the legal changes. The dashed lines show the removal of the SDD clause in 2015 and the enactment of the DPA in 2018. Panel (a) shows the dynamic response for our baseline results. Compared to the low-exposure industries, the high exposure industries show significantly higher demand for cyber skills starting from 2017. Panel (b) and (c) plot the results for alternative definitions of the exposure index, whereas Panel (d) shows the response for service industries. Across all panels, we do not see any differential response prior to the legal changes. Following the removal of the SDD clause, we see a somewhat lagged response but the effect grows stronger following the DPA 2018.

5.2 Cyber Skill Demand across Firms

We also check our results using the subset of job postings with non-missing employer names. Our firm-level regressions control for firm fixed effects and time fixed effects by TTWA. In our data, a firm can have multiple plants across different TTWAs. We use the share of job postings per firm-TTWA (i.e., plant) as our dependent variable. Table 4 shows how the demand for cyber skills changes for firms. Column 1 shows the yearly response with the year 2015 used as a baseline year (also see Figure 8). We can see that prior to 2015 the high exposure industries appear to be the same as the low exposure industries in terms of their demand for cyber skills. The response is significantly higher from 2016 onwards. Column 2 shows the average effect for the two periods of interest. Between 2016 and 2018, the hiring response jumps by 37% compared to the sample mean in 2015. Similar to our local labor market results, the effect is higher for the DPA 2018 with a 73% rise in the share of cyber job postings.

5.3 Firm's Digital Technology Profile and Data Intensity

Data-intensive firms are more exposed to the reforms affecting the data protection codes and practices. In this section, we examine the differential effects of our two legal changes according to firms' data intensity. To this end, we face an empirical challenge to define the importance of the data as a factor of production. Traditionally, the intensity of production factors, such as labor and capital, have been interpreted as technological parameters. Thus, we began by defining a set of digital technology profiles that strongly signal firms' dependence on data. In this paper, we rely on firms' skill acquisition patterns to delineate their engagement with different classes of technologies. The reason is that human capital is a key input into technology adoption. Moreover, it can capture firms' engagement with technologies earlier than other sources such as patent data (Tambe and Hitt, 2012a,b). This is particularly important in the case of more novel technologies such as AI and big data (Goldfarb et al., 2022). Thus, there is a growing number of papers that have used job posting data to infer technology diffusion among firms.

Following this literature, we exploit the granular information on advertised skills in job postings. We examine whether firms require any digital skills before the legal changes take place. We tag skills as digital if they fall in one of the two broad classes of technologies, namely digital organization and data harvesting technologies. These types roughly mirror the two types of data-related skills in Abis and Veldkamp (2020) that estimate a model of the information production function.

Figure 9 visualizes these two sets of digital skills. The data harvesting category entails five major skill clusters: data mining, business intelligence (BI), extract load and transform (ETL), artificial intelligence (AI), and big data. Acquiring skills from this category implies that firms actively invest to generate value from data processing. The second group includes digital organization skills that are nested under three major clusters: SQL data management, enterprise resource planning (ERP), and cloud computing and storage. Investing in skills related to these technologies shows that firms hold valuable data and actively spend on their management and storage. In particular, we narrow our focus to those technologies that are essential to support the digital organization of the firms, but whose vulnerability increases the risk of a cyberattack.

We define eight variables for each of these major clusters. They show whether a firm required such skills between 2012 and 2015. These variables represent the demand for various types of digital skills and thus indicate different aspects of firms' data intensity. Similarly, we further define a dummy variable for skill acquisition from the cybersecurity cluster before the legal changes.

Using the first two principal components, we generate a two-dimensional summary of the observed variation in firms' demand across these skill clusters. This is visualized in Figure 10. The figure provides two important insights. First, the factor loadings of the first principal component (PC1) are positive across all clusters. This implies that the PC1 captures the common feature of firms that invest across cybersecurity and digital skill clusters. We infer that this common feature can be a decent proxy of data intensity. Excluding the cybersecurity cluster also yields very similar results. Thus we use the first principal component of demand for eight major digital skill clusters as our main index of data intensity.

Columns (1) and (2) in Table 4 report the results when we run our baseline regression over two subsamples of firms that are above and below the median level of the data intensity index. The results indicate that the share of job postings that required any cyber skill increased by 0.04% among data-intensive firms after the removal of the SDD clause and further surged up to 0.13% after the enactment of the DPA 2018. However, the change in cybersecurity hiring among firms with low data intensity remained insignificant. In columns (3) and (4), we use an alternative data intensity index that is only based on digital organization skills. Similarly, columns (5) and (6) report the results for an index only based on data harvesting skills. The two indices, based on the components of our main data intensity measure, yield similar results albeit the effects for the data harvesting skills are slightly stronger.

The second insight from Figure 10 comes from a comparison of the PC-based coordination of a digital skill cluster with cybersecurity. The comparison reveals that cloud and cybersecurity show a strong similarity along their principal components. This implies that firms that invested in cloud technologies prior to the legal changes are very likely to have a demand for cyber skills as well. Columns (7) and (8) show the results for subsamples of firms that acquired cloud-related skills prior to the legal changes and those that did not. Our estimates highlight a strong rise of 0.28% in cyber hirings among firms that adopted cloud technologies after the passage of the DPA 2018. The dynamic effects in Figure 11 also show the divergence in cyber skills demand after the legal changes across the firms with and without data-intensive technologies or cloud technologies.

5.4 Firm's Cash Holding before the Legal Changes

There has been a large increase in corporate cash holding over the past few decades that partly stems from precautionary motives among managers (Bates et al., 2009) to be competitive in the product market (Fresard, 2010; Haushalter et al., 2007; Boutin et al., 2013), or to hedge against the volatility in the stock market (Farre-Mensa, 2014; Campello et al., 2018), and weather potential financial stress after a rise in the leverage (Subrahmanyam et al., 2017). In particular, recent evidence indicates that corporations hold cash to stay competitive when they acquire talents and key human capital (Brown and Matsa, 2016; He, 2018; Baghai et al., 2021).

Following this literature, we expect a firm's ability to acquire new skills to depend on its liquidity constraint. Thus, we divide firms into two subsamples according to their cash ratio in 2015, as measured by the firm's cash and cash equivalent securities over its total assets. We refer to the group above the median level of the cash ratio as high cash holding firms. Columns (1) and (2) of Table 5 report the results of our baseline regression for the subsample of high and low cash holding firms. Both groups increase their cyber hirings after the removal of the SDD clause. Although the demand for cyber skills doubles among high cash holding firms following the DPA 2018 coming into effect, it stagnates among firms with limited access to internal liquidity — a pattern we also observe limiting our sample to service sectors in columns (3) and (4). Furthermore, the liquidity management practices might significantly differ across sectors creating undesirable bias in our estimates. In columns (5) and (6), we divide firms according to the median level of cash holding for each two-digit SIC industry. The results show a pattern similar to before.

Figure 13 plots the dynamic effects, which reveal a pronounced and steady rise in cyber skills demand among firms with a high cash reserve. The results provide two interesting insights. First, they highlight the importance of cash holding in firms' adaptation to a more stringent data protection environment. This can stem from the precautionary motive arising from recruiting difficulties of key human capital. Second, the divergence in cybersecurity hiring becomes particularly evident after the enactment of the DPA 2018. This pattern is consistent with our other results when we examine the differential effects according to firm characteristics. These results shed light on the differences in how increased enforcement (i.e., removal of the SDD clause) and increased penalties (i.e., the passage of DPA 2018) can affect cyber hirings. Although the impact of increased penalties is substantially stronger, it is less homogeneous across firms with different characteristics.

5.5 The Adverse Effect on Firm Dynamics

So far we have established that increased enforcement and penalties in data protection cases are effective instruments to increase the investment in cyber skills. This can close the gap between the socially optimum level and firms' provision of cybersecurity. Nevertheless, a stronger data protection environment can have unintended negative impacts on firms' performances and industry concentrations. (Goldberg et al., 2019; Abis et al., 2022; Schmitt et al., 2020; Zhao et al., 2021; Chen et al., 2022; Buckman et al., 2021; Mayya and Viswanathan, 2021; Johnson et al., 2021). Our last section of results provide evidence on the negative effects of the legal changes on firm dynamics in industries exposed to ICO's data protection enforcement.

We begin by studying the effect of the legal changes on cyber security hirings over firms' life cycles. To this end, we compare the response of old and young incumbents that were above and below the median age in 2015. Column (1) in table 6 shows a strong rise in demand for cyber professionals among young incumbents. The cyber skill demand rose by 0.06% after the removal of the SDD clause. The effect doubled after the enactment of the DPA 2018. The old incumbents show the same surge in demand for cyber skills but the response stayed flat afterward. Considering that younger firms often face more difficulties in recruiting skilled employees (Israelsen and Yonker (2017)), this might have disruptive effects on firm dynamics.

Using the universe of firm turnover from the UK business register, columns (3) and (4) examine this hypothesis. Column (3) reports the effect of legal changes on the birth rate

which is measured by the number of new firms as a share of registered firms for each year-industry-TTWA. Our estimates indicate that the birth rate in highly exposed industries to ICO enforcement dropped by 0.6% after the removal of the SDD clause although the effect is not statistically significant. Nevertheless, after the DPA 2018, we estimate the birth rate declines by 1.4% which is economically and statistically significant. Column (4) reports the results for the death rate which is measured by the number of exiting firms to the registered firm at each year-industry-TTWA level. It indicates that the exit rate for high exposure industries increased by 0.9% after removing the SDD clause. Moreover, the death rate of firms operating in highly exposed industries stayed 0.7% higher on average following the DPA 2018 compared with the period before the legal change occurred.

6 Conclusion

Each major cyberattack brings a renewed focus on the vulnerabilities of firms and the importance of investing in cybersecurity. However, the current level of expenditure on cybersecurity is not adequate. In this paper, we delve into the role of data protection regulation in rectifying the lack of investment in cyber skills. The UK has a strong data protection regulatory environment with the Information Commissioner overseeing compliance with the DPA, GDPR, and other relevant regulations. Using information from the ICO's activity log, we consider around 5,000 cases where the ICO has taken regulatory actions against business organizations. We construct a score at the three-digit industry level and check how the industries with high exposure to ICO enforcement respond to legal changes. We consider two different regulatory changes — the first one makes it easier to issue fines for violating privacy rights related to electronic communications and the other imposes a mega-fine following the DPA 2018 coming into effect. Together the timing of the legal changes and sectoral variations in the exposure to ICO actions help us isolate the effect on the demand for cyber skills.

Our local market and firm-level analysis show that both the increased number of fines and mega-fines induce firms to invest more in cyber personnel. Studying the heterogeneous effects across firms, we show which types of firms have a stronger response. Our results suggest that firms with prior investment in digital technologies and more liquidity, as well as relatively young firms, have a stronger demand for cyber skills. Our paper provides the first empirical evidence of how data protection frameworks can address the agency problem that arises due to the gap between the private costs (of firms) and the social costs of cybercrimes. However, strengthening the data protection laws also comes at a cost. We find a significant decline in entry rate in the industries highly exposed to ICO enforcement. These findings outline the

trade-off faced by policymakers or data protection authorities who want to mitigate cyber risks without slowing down the growth of the business sector.

References

Abis, S., Canayaz, M., Kantorovitch, I., Mihet, R., and Tang, H. (2022). Privacy laws and value of personal data. Technical report, EPFL.

Abis, S. and Veldkamp, L. (2020). The changing economics of knowledge production. *Available at SSRN 3570130*.

Adams-Prassl, A., Balgova, M., and Qian, M. (2020). Flexible work arrangements in low wage jobs: Evidence from job vacancy data.

Alekseeva, L., Azar, J., Gine, M., Samila, S., and Taska, B. (2021). The demand for ai skills in the labor market. *Labour economics*, 71:102002.

Amir, E., Levi, S., and Livne, T. (2018). Do firms underreport information on cyber-attacks? evidence from capital markets. *Review of Accounting Studies*, 23(3):1177–1206.

Azar, J., Marinescu, I., Steinbaum, M., and Taska, B. (2020). Concentration in us labor markets: Evidence from online vacancy data. *Labour Economics*, 66:101886.

Babina, T., Fedyk, A., He, A. X., and Hodson, J. (2020). Artificial intelligence, firm growth, and industry concentration. *Firm Growth, and Industry Concentration (November*, 22:2020.

Baghai, R. P., Silva, R. C., Thell, V., and Vig, V. (2021). Talent in distressed firms: Investigating the labor costs of financial distress. *The Journal of Finance*, 76(6):2907–2961.

Bana, S., Brynjolfsson, E., Jin, W., Steffen, S., and Wang, X. (2021). Cybersecurity hiring in response to data breaches. *Available at SSRN 3806060*.

Bates, T. W., Kahle, K. M., and Stulz, R. M. (2009). Why do us firms hold so much more cash than they used to? *The journal of finance*, 64(5):1985–2021.

Boasiako, K. A. and Keefe, M. O. (2021). Data breaches and corporate liquidity management. *European Financial Management*, 27(3):528–551.

Boutin, X., Cestone, G., Fumagalli, C., Pica, G., and Serrano-Velarde, N. (2013). The deep-pocket effect of internal capital markets. *Journal of Financial Economics*, 109(1):122–145.

Brown, J. and Matsa, D. A. (2016). Boarding a sinking ship? an investigation of job applications to distressed firms. *The Journal of Finance*, 71(2):507–550.

Buckman, J., Adjerid, I., and Tucker, C. E. (2021). Privacy regulation and barriers to public health. *Available at SSRN 3983334*.

Campello, M., Matta, R., and Saffi, P. (2018). The rise of the equity lending market: implications for corporate policies. *Available at SSRN*, 2703318.

Chen, C., Frey, C. B., and Presidente, G. (2022). Privacy regulation and firm performance: Estimating the gdpr effect globally. Technical report, The Oxford Martin Working Paper Series on Technological and Economic Change.

Crosignani, M., Macchiavelli, M., and Silva, A. F. (2021). Pirates without borders: The propagation of cyberattacks through firms' supply chains. *FRB of New York Staff Report*, (937).

De Cornière, A. and Taylor, G. (2021). A model of information security and competition.

Deming, D. and Kahn, L. B. (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, 36(S1):S337–S369.

Deming, D. J. and Noray, K. (2020). Earnings dynamics, changing job skills, and stem careers. *The Quarterly Journal of Economics*, 135(4):1965–2005.

Dixit, A. (1997). Investment and employment dynamics in the short run and the long run. *Oxford Economic Papers*, 49(1):1–20.

Dušek, L. and Traxler, C. (2022). Learning from law enforcement. *Journal of the European Economic Association*, 20(2):739–777.

Farre-Mensa, J. (2014). Comparing the cash policies of public and private firms. *Harvard University working*.

Florackis, C., Louca, C., Michaely, R., and Weber, M. (2020). Cybersecurity risk. Technical report, National Bureau of Economic Research.

Forsythe, E., Kahn, L. B., Lange, F., and Wiczer, D. (2020). Labor demand in the time of covid-19: Evidence from vacancy postings and ui claims. *Journal of public economics*, 189:104238.

Fresard, L. (2010). Financial strength and product market behavior: The real effects of corporate cash holdings. *The Journal of finance*, 65(3):1097–1122.

Garg, P. (2020). Cybersecurity breaches and cash holdings: Spillover effect. *Financial Management*, 49(2):503–519.

Goldberg, S., Johnson, G., and Shriver, S. (2019). Regulating privacy online: An economic evaluation of the gdpr. *Available at SSRN 3421731*.

Goldfarb, A., Taska, B., and Teodoridis, F. (2022). Could machine learning be a general purpose technology? a comparison of emerging technologies using data from online job postings. Technical report, National Bureau of Economic Research.

Gordon, L. A., Loeb, M. P., Lucyshyn, W., and Zhou, L. (2015a). The impact of information sharing on cybersecurity underinvestment: A real options perspective. *Journal of Accounting and Public Policy*, 34(5):509–519.

Gordon, L. A., Loeb, M. P., Lucyshyn, W., and Zhou, L. (2015b). Increasing cybersecurity investments in private sector firms. *Journal of Cybersecurity*, 1(1):3–17.

Haushalter, D., Klasa, S., and Maxwell, W. F. (2007). The influence of product market dynamics on a firm's cash holdings and hedging behavior. *Journal of Financial Economics*, 84(3):797–825.

He, Z. (2018). Money held for moving stars: Talent competition and corporate cash holdings. *Journal of Corporate Finance*, 51:210–234.

Hershbein, B. and Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7):1737–72.

Hilary, G., Segal, B., and Zhang, M. H. (2016). Cyber-risk disclosure: who cares? *Georgetown McDonough School of Business Research Paper*, (2852519).

ICO (2012). A guide to the legislation the ICO regulates. upholding information rights for all.

ICO (2014). How we deal with complaints and concerns.

ICO (2018). Guide to the privacy and electronic communications regulations.

Israelsen, R. D. and Yonker, S. E. (2017). Key human capital. *Journal of Financial and Quantitative analysis*, 52(1):175–214.

Javorcik, B., Stapleton, K., Kett, B., and O'Kane, L. (2020). Unravelling deep integration: Local labour market effects of the brexit vote. Technical report, CEPR Discussion Paper 14222.

Johnson, G., Shriver, S., and Goldberg, S. (2021). Privacy & market concentration: Intended & unintended consequences of the gdpr. *Available at SSRN 3477686*.

Johnson, M., Kang, M. J., and Lawson, T. (2017). Stock price reaction to data breaches. *Journal of Finance Issues*, 16(2):1–13.

Kamiya, S., Kang, J.-K., Kim, J., Milidonis, A., and Stulz, R. M. (2021). Risk management, firm reputation, and the impact of successful cyberattacks on target firms. *Journal of Financial Economics*, 139(3):719–749.

Kankanhalli, A., Teo, H.-H., Tan, B. C., and Wei, K.-K. (2003). An integrative study of information systems security effectiveness. *International journal of information management*, 23(2):139–154.

Kopp, E., Kaffenberger, L., and Wilson, C. (2017). *Cyber risk, market failures, and financial stability*. International Monetary Fund.

Mayya, R. and Viswanathan, S. (2021). Delaying informed consent: An empirical investigation of mobile apps' upgrade decisions. *Available at SSRN 3457018*.

Miller, D. L., Cameron, A. C., and Gelbach, J. (2009). Robust inference with multi-way clustering. Technical report, Working Paper.

Oi, W. Y. (1962). Labor as a quasi-fixed factor. *Journal of political economy*, 70(6):538–555.

Richardson, V. J., Smith, R. E., and Watson, M. W. (2019). Much ado about nothing: The (lack of) economic impact of data privacy breaches. *Journal of Information Systems*, 33(3):227–265.

Schmitt, J., Miller, K. M., and Skiera, B. (2020). The impact of privacy laws on online user behavior.

Shapiro, M. D. (1986). The dynamic demand for capital and labor. *The Quarterly Journal of Economics*, 101(3):513–542.

Subrahmanyam, M. G., Tang, D. Y., and Wang, S. Q. (2017). Credit default swaps, exacting creditors and corporate liquidity management. *Journal of Financial Economics*, 124(2):395–414.

Tambe, P. and Hitt, L. M. (2012a). Now it's personal: Offshoring and the shifting skill composition of the us information technology workforce. *Management Science*, 58(4):678–695.

Tambe, P. and Hitt, L. M. (2012b). The productivity of information technology investments: New evidence from it labor data. *Information systems research*, 23(3-part-1):599–617.

Zhao, Y., Yildirim, P., and Chintagunta, P. K. (2021). Privacy regulations and online search friction: Evidence from gdpr. *Available at SSRN 3903599*.

A Figures and Tables

Figure 1: How the ICO Processes Complaints

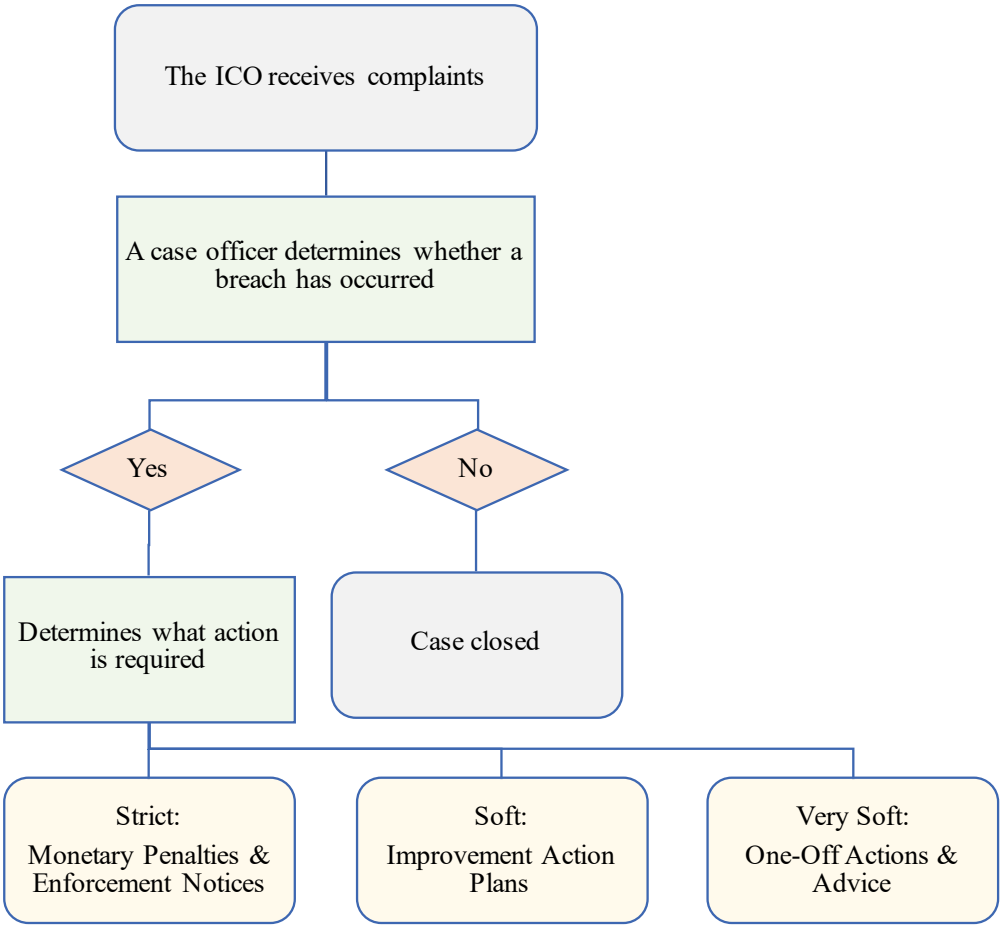

Note: This figure shows a simplified flow of actions that explains how the ICO deals with data protection complaints. The process starts with assigning a case officer who determines whether any law has been violated and what actions to take in case of a violation. The ICO can choose from a number of regulatory tools depending on the severity of breaches. Among the strictest actions are monetary penalties and enforcement notices, which apply to the most severe breaches. The ICO can also choose non-punitive actions such as improvement action plans, one-off actions, or advice. Note that we exclude audit from the list of possible actions as most of the ICO case outcomes involving audit applies to public organizations.

Figure 2: Major Changes in Data Protection Regulations

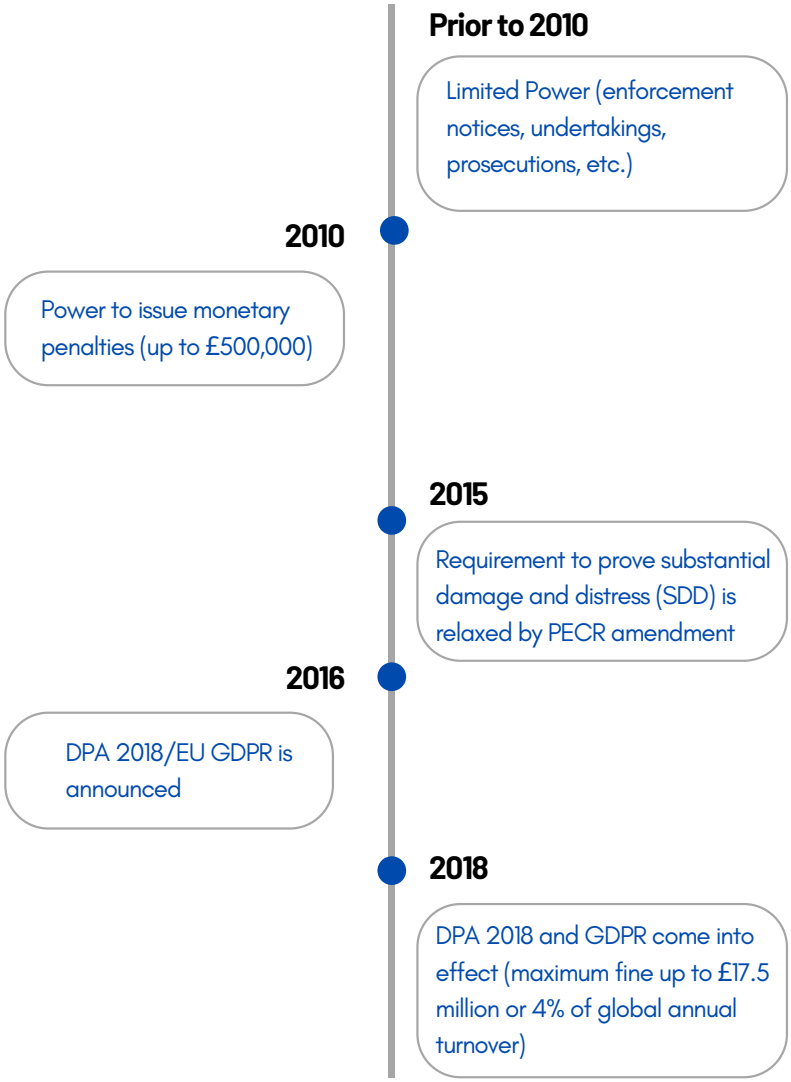

Note: The figure shows major changes in the data protection regulatory environment. From 2010 (2011), the ICO could impose monetary penalties for DPA (PECR) related cases. Other notable changes include the removal of SDD requirement for PECR cases in 2015 and the increased threshold of maximum fine following the DPA 2018 coming into effect (*Source: ICO website*).

Figure 3: Changes in ICO Supervisory Actions

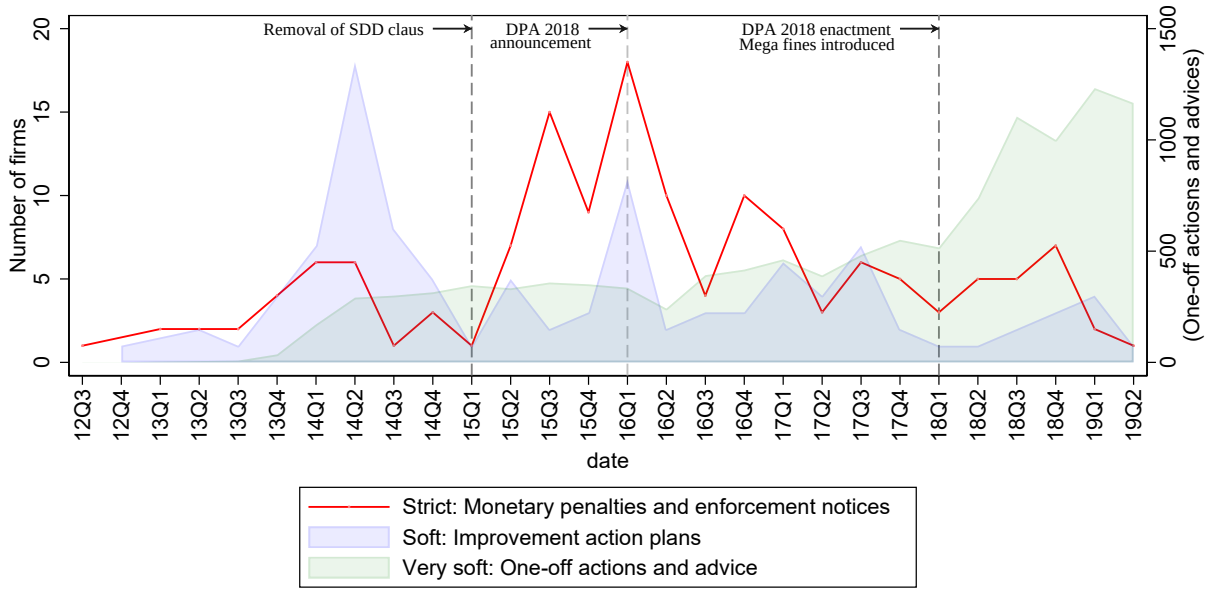

Note: The figure highlights important changes in ICO enforcement strategies after the two legal reforms. It depicts the number of firms per quarter that were targeted by different types of ICO actions. Before 2015, improvement action plans were a dominant supervisory instrument through which the ICO worked with organizations to bring them closer to compliance rather than using its formal enforcement power. However, following the removal of the SDD requirement in 2015, there was a surge in monetary penalties and enforcement notices which are the most strict types of enforcement. Nevertheless, they gradually declined after the passage of the DPA 2018 which substantially increased the maximum fine in data protection cases. During this period, the ICO began to rely more on one-off actions and advice to deal with data protection compliance. Note that the figure reports ICO actions taken against business organizations.Three gray dashed lines respectively indicate the quarter before the following three events take place: 1) the removal of the SDD clause in April 2015, 2) the announcement of the DPA 2018 in May 2016, and 3) the enactment of the DPA 2018 in May 2018.

Figure 4: Required Cyber Skills in Job Postings

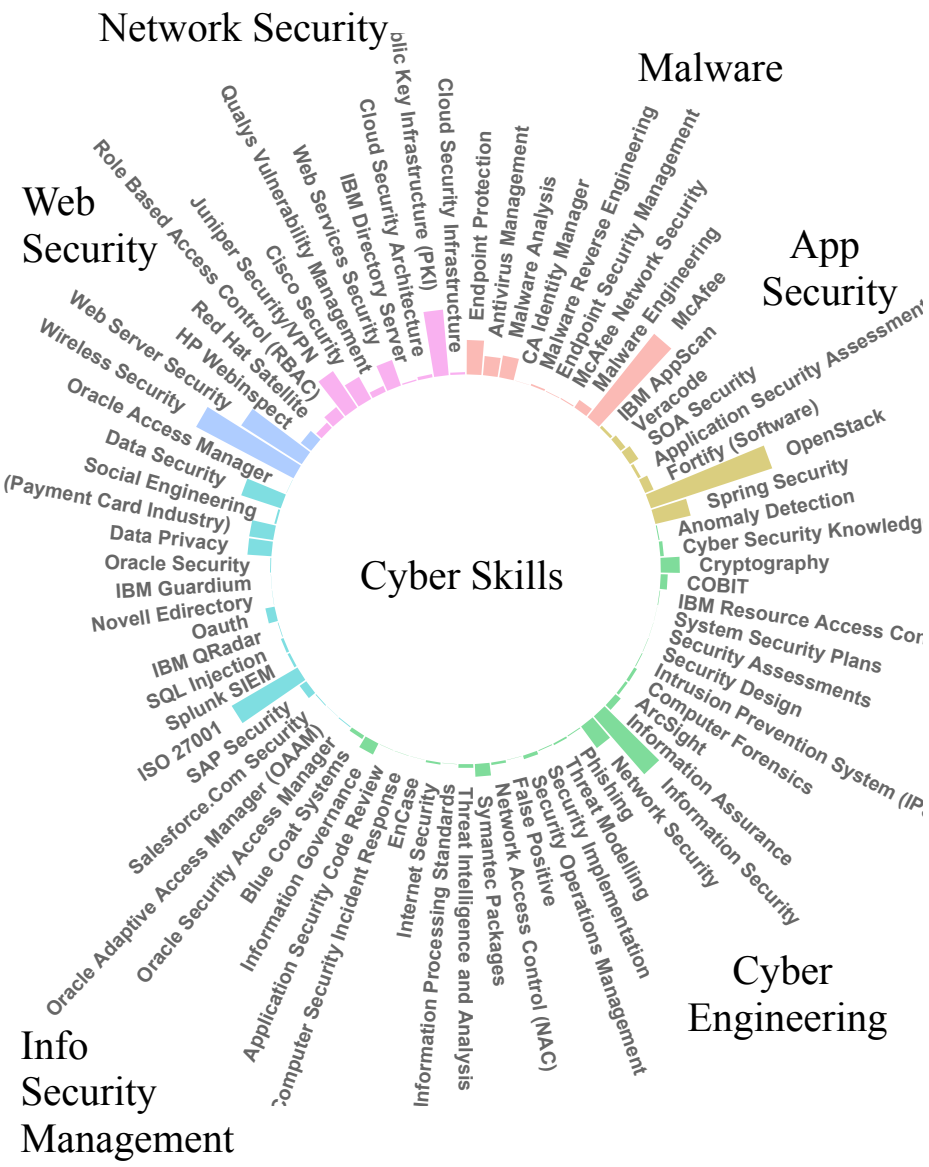

Note: The figure visualizes 87 cyber skills that are nested under six major skill clusters: information security management, anti-malware software, web security, application security, and network security skill clusters. We tag a job posting as a cybersecurity job advertisement if any of these skills is required. Each bar refers to a specific skill. Bars are color-coded and each color represents a cluster. The taller bar shows that the skill is more prevalent within its cluster.

Figure 5: Percentage of Three-Digit Industries by Exposure to ICO Enforcement

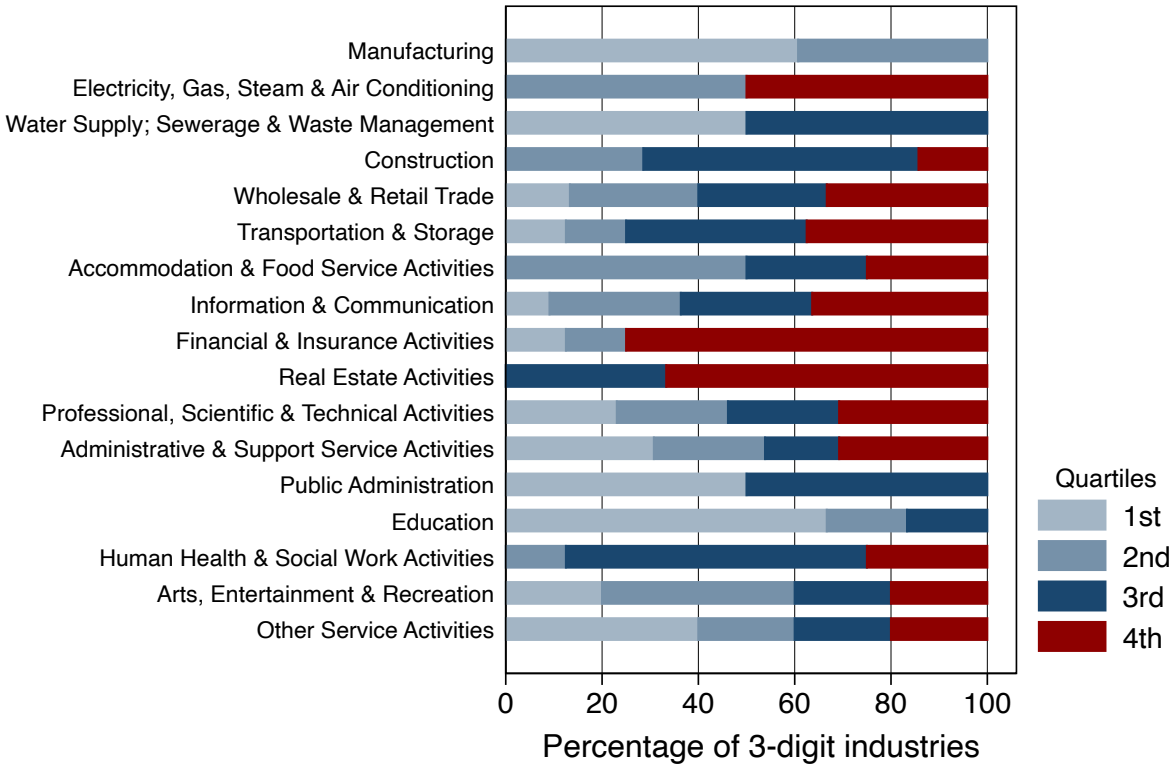

Note: The figure reports the share of three-digit industries within each two-digit sector that fall within different quartiles of the ICO exposure index. We calculate the exposure index by the number of firms against which the ICO takes action(s) between 2012 and 2018:Q2. The red bar (top quartile) represents the industries with high exposure to ICO enforcement.

Figure 6: Demand for Cyber Skills over Time

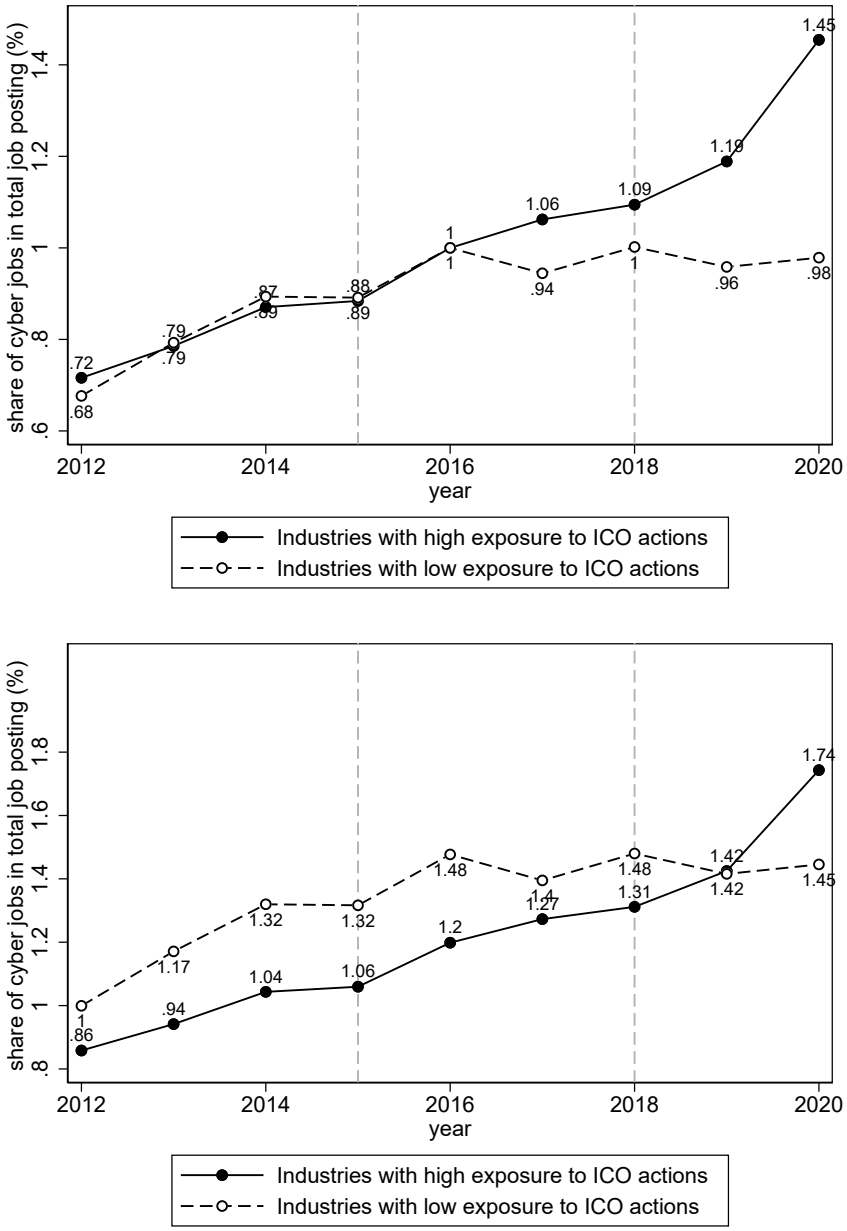

Note: The figure reports the share of cyber job posting over time. It shows the divergence in cyber skill demand across industries with high and low exposure to ICO enforcement. The gray dashed lines show the removal of the SDD clause in 2015 and the enactment of the DPA 2018. 'Industries with high exposure to ICO actions' are industries that are in the top quartile of three-digit industries in terms of the number of firms that are subject to ICO actions in that industry. 'Industries with low exposure to ICO actions' represent the rest of the industries. In panel (a), the shares of cyber job postings are normalized by their 2016 counterparts to highlight the divergence that began in 2016 and magnified in 2018. panel (b) reports the trends without normalization which shows that highly exposed industries initially post fewer cybersecurity jobs compared with the other industries.

Figure 7: Dynamic Effects of Local Labor Markets' Demand for Cyber Skill

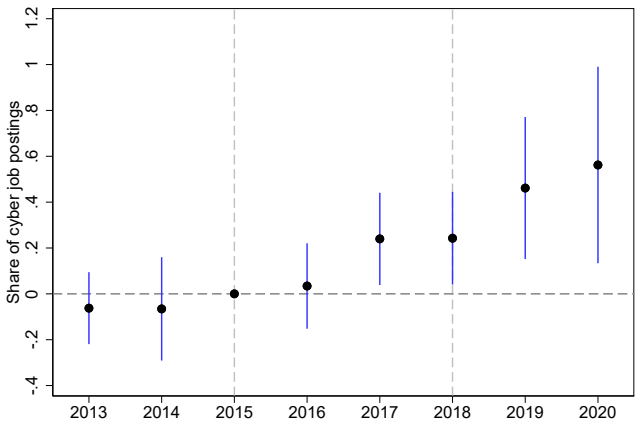

(a)

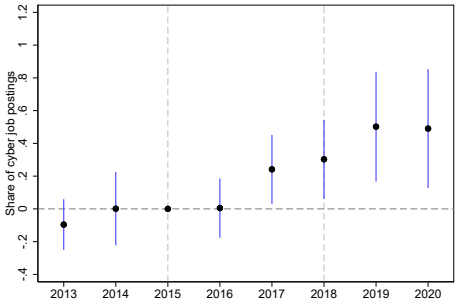

(b)

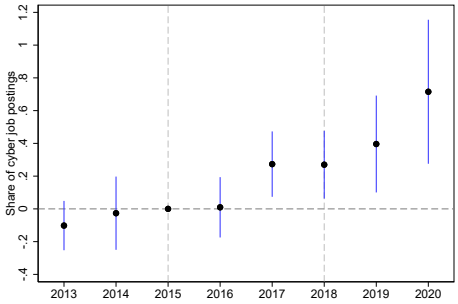

(c)

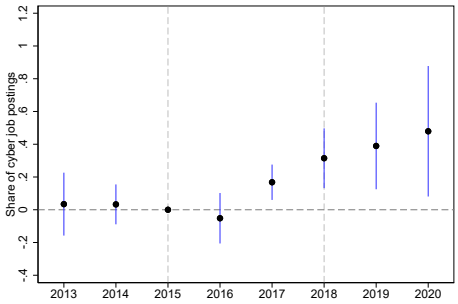

(d)

Note: The figure highlights the changes in cyber skill demand after the legal changes. We compare the demand for cyber skills in industries with high and low exposure to ICO enforcement located in the same travel to work area (TTWA). The gray dashed lines show the removal of the SDD clause in 2015 and the enactment of the DPA 2018. Panel (a) shows our baseline results in which high exposure to ICO enforcement are industries that are in the top quartile of 3-digit industries in terms of the number of firms that are subject to ICO actions. Panel (b) reports the results when the high exposure to ICO enforcement is defined according to the number of cases. Panel (c) reports the results when the high exposure to ICO enforcement is defined according to the number of cases but they are weighted by their severity. In this case, we weigh monetary penalties more than other actions. We use the inverse of occurrence for monetary penalties for weighting monetary penalties cases. Similarly, we weigh other actions by their inverse of occurrence. Panel (d) uses our baseline index for ICO exposure but limits our sample to the services sectors.

Figure 8: Firms' Demand for Cyber Skills in Response to the Legal Changes

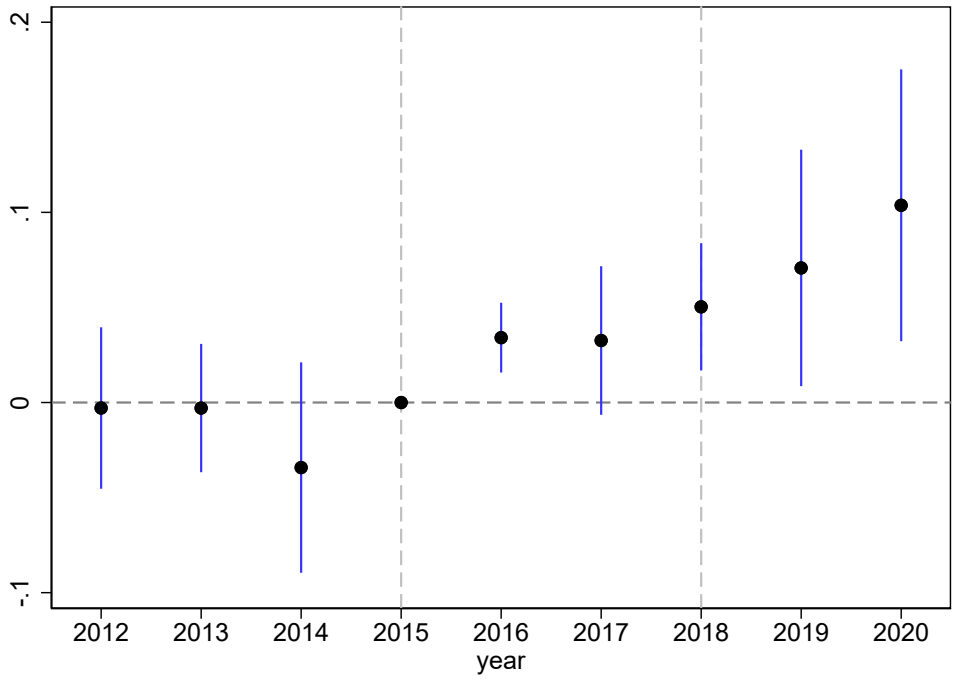

Note: In this figure, the black dots visualize the point estimates from column (1) of Table 3. The blue lines plot the 95% confidence interval. The gray dashed lines show the removal of the SDD clause in 2015 and the enactment of the DPA 2018.

Figure 9: Required Digital Skills in Job Postings

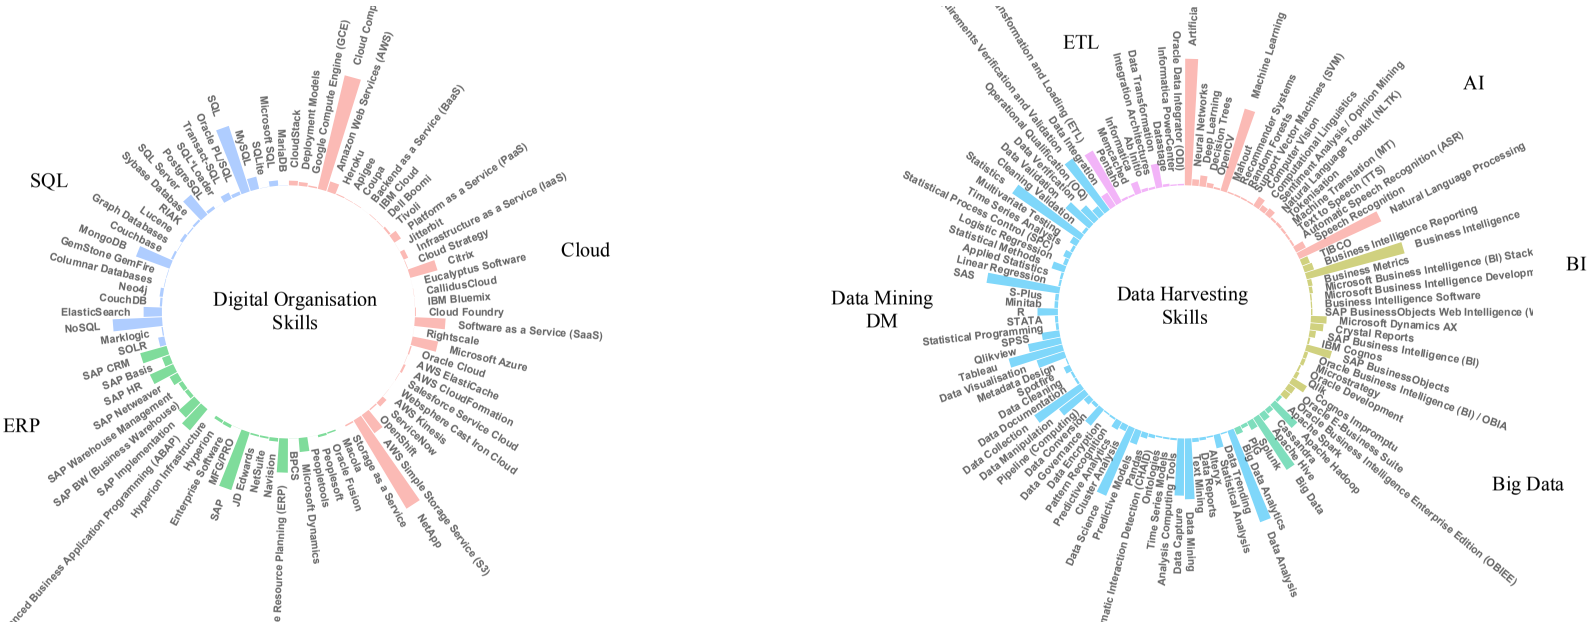

(a) Digital organization skills

(b) Data harvesting skills

The figure visualizes the digital skills required by job postings under two broad categories: digital organization and data harvesting. Panel (a) shows digital organization skills that are nested under three major skill clusters: SQL data management, enterprise resource planning (ERP), and cloud computing and storage. Panel (b) shows the data harvesting skills that are nested under five skill clusters: data mining, business intelligence (BI), extract load and transform (ETL), artificial intelligence (AI), and big data. Note in both panels, each bar refers to a specific skill. Bars are color-coded and each color represents a skill cluster. The taller bar shows that the skill is more prevalent within its cluster.

Figure 10: Principal Components of Firms' Demand for Digital Skills

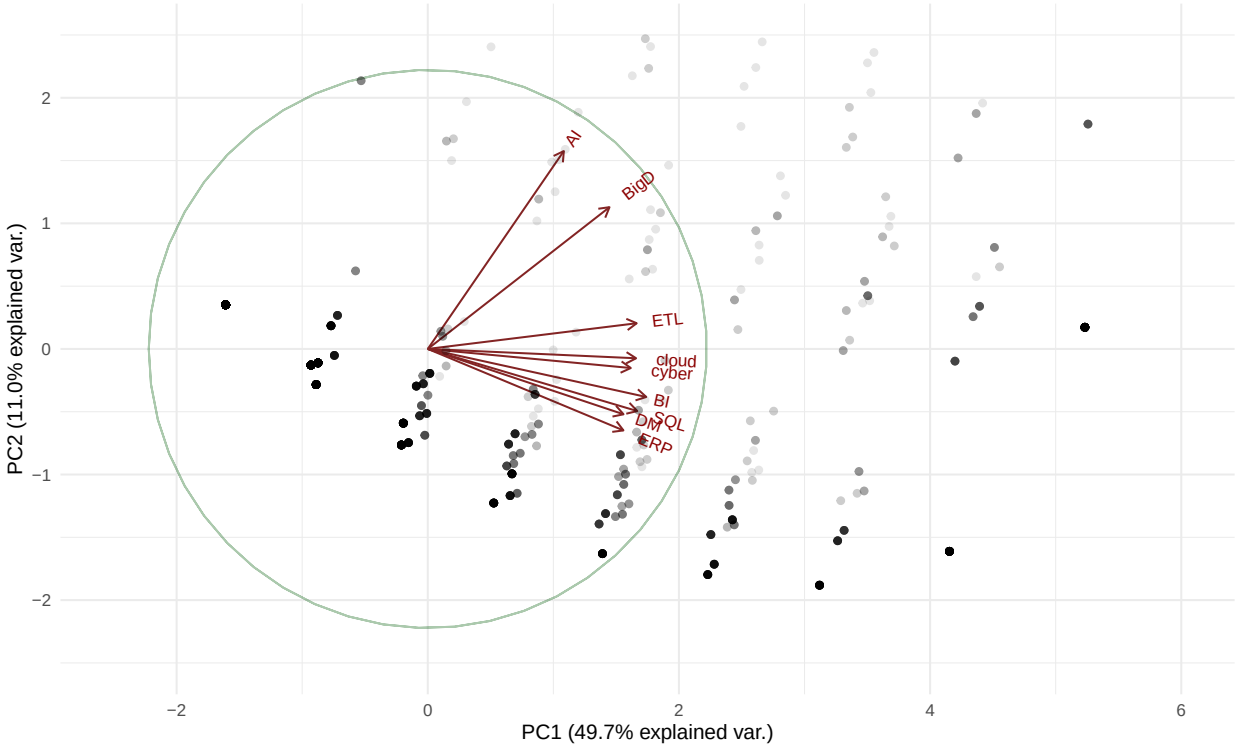

Note: The figure provides a two-dimensional visual summary of the observed variation in digital skill demand between 2012 and 2015 which is the period before the legal changes take place. The principal components are built based on nine dummy variables. Each dummy variable indicates a demand for a specific skill cluster. The dummy takes one if the firm posted a job advertisement requiring a skill belonging to that specific skill cluster. These clusters include the three skill clusters in the digital organization skills (SQL data management, enterprise resource planning, and cloud computing and storage), five skill clusters in data harvesting skills (data mining, business intelligence, extract load and transform, artificial intelligence, and big data) and the cyber skill cluster. Note each arrow indicates the PC-based coordination of a skill cluster. The dots indicate the PC-based locations of firms whereas the darker dots indicate the locations that are more frequent among firms.

Figure 11: Firms' Digital Technology Profile and Data Intensity

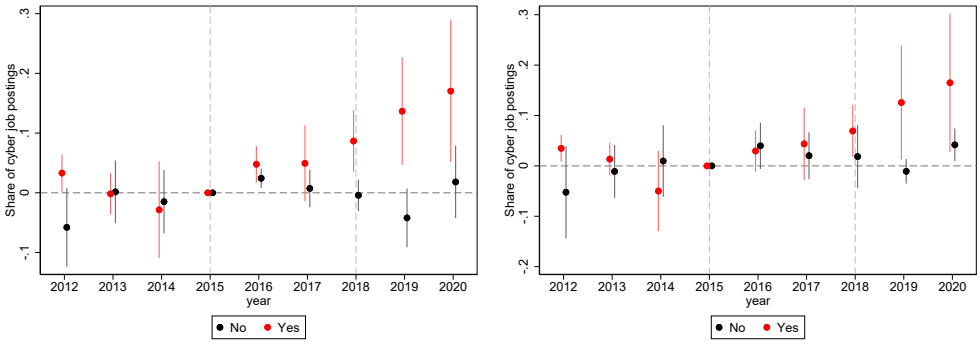

(a) All Digital skills(PC1)

(b) Digital Organization(PC1)

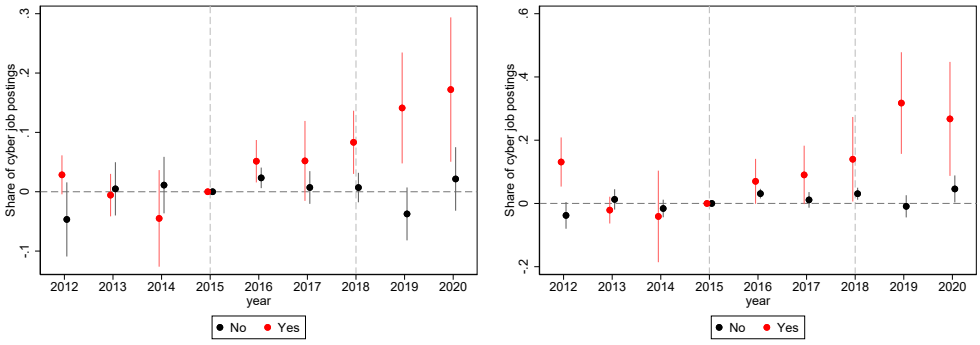

(c) Data Harvesting(PC1)

(d) Cloud

Note: Figure 10 highlights the strong response of data-intensive firms. In panel (a)-(c), we build a data intensity index at the firm level using the first principal component (PC1) of variation in firm demand for digital skills between 2012 and 2015. Low and high in these three panels respectively refer to firms that are below (low data intensity) or above (high data intensity) the median level of PC1. Panel (a) uses the first principal component based on investment in any of the three clusters of digital organization – SQL data management, enterprise resource planning (ERP), and cloud computing and storage and five clusters of data harvesting skills – data mining, business intelligence (BI), extract load and transform (ETL), artificial intelligence (AI), and big data. Panel (b) uses the first principal component based on investment in the three clusters of digital organization. Panel (c) uses the first principal component based on investment in the five clusters of data harvesting skills. Panel (d) indicates a strong surge in demand for cyber skills among firms that invested in cloud-related skills between 2012 and 2015.

Figure 12: Differential Response by Firms' (ex-ante) Cash Holding

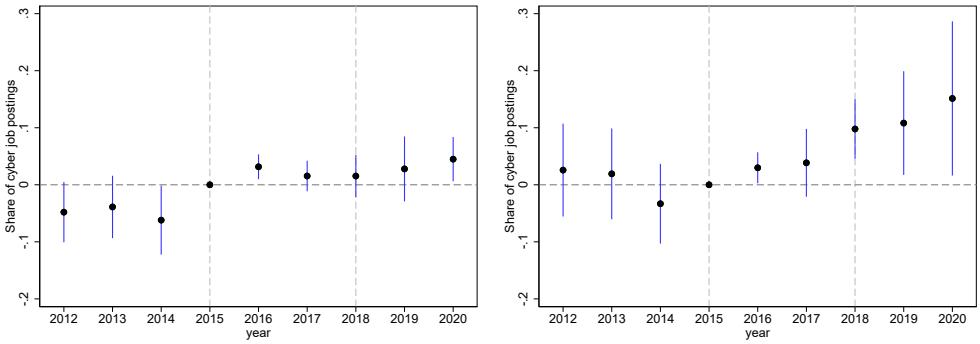

(a) Low cash holding

(b) High cash holding

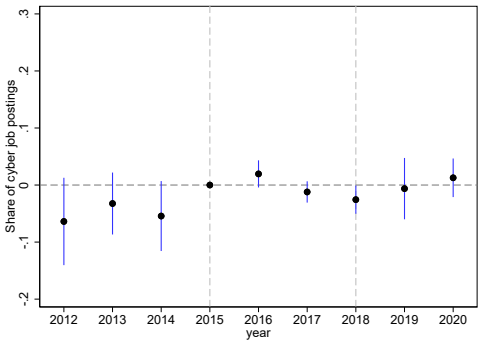

(c) Low cash holding

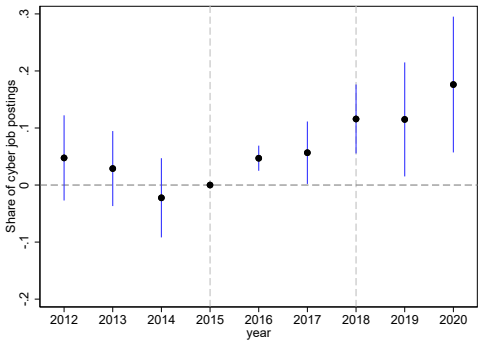

(d) High cash holding

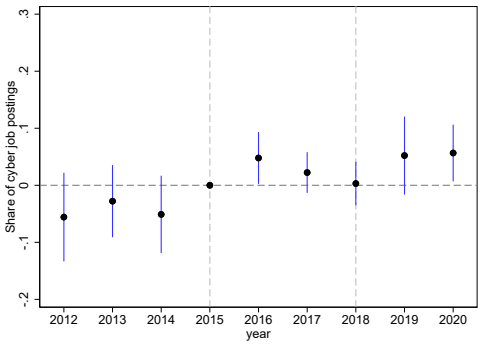

(e) Low cash holding

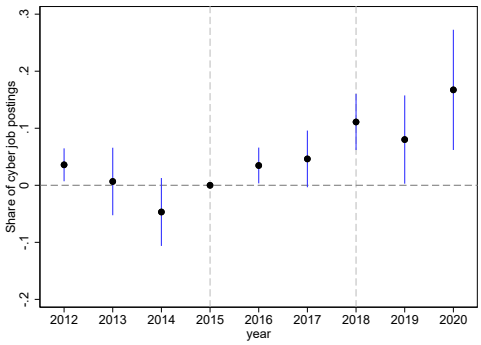

(f) High cash holding

Note: The figure highlights the strong response of firms with higher internal liquidity before the legal changes. In panel (a) and (b), we compare the firms above and below the median level of cash holding in 2015. In panel (c) and (d), we limit our sample to services and thus we divide firms based on the median level of cash holding for services in 2015. In panel (e) and (f), we divide firms according to the median level of cash holding for each two-digit SIC industry.

Figure 13: Dynamic Response (Age and Firm Dynamics)

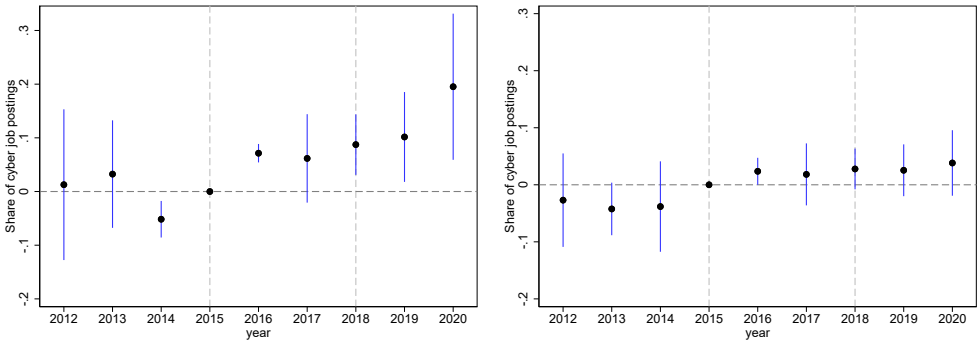

(a) Young incumbent

(b) Old incumbent

Firm dynamic after the legal changes

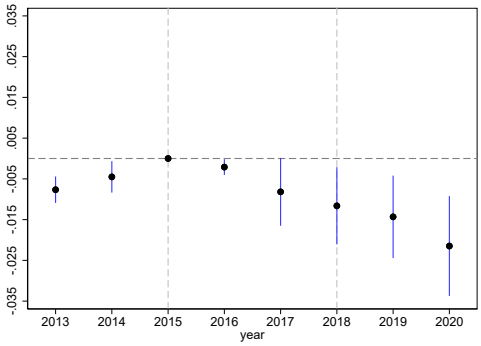

(c) Birth rate

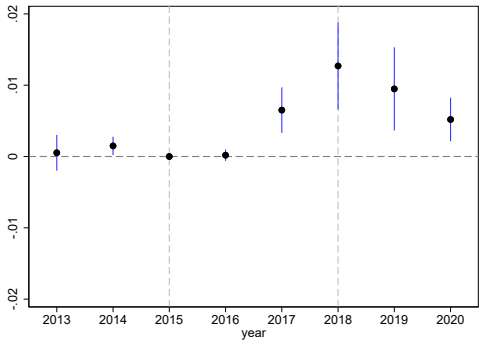

(d) Death rate

Note: In the figure, Panels (a) and (b) highlight the stronger response of young incumbent firms where we compare the firms above and below the median firm's age in 2015 that is fifteen years old. Panels (c) and (d) report the dynamic response for birth and death rates at the TTWA-industry level.

Table 1: Summary Statistics

	Mean	SD	Percentile 1st	Percentile 99th
Panel A: Job postings data				
% Cyber (TTWA panel)	1.15	3.93	0	13.3
% Cyber (firm panel)	0.14	1.92	0	3.44
Panel B: Firm data				
Cash holding (2015)	0.24	0.24	0.0	0.93
Age (2015)	19.16	17.7	1	93
Birth rate (2013-2020)	0.12	0.06	0.02	0.29
Death rate (2013-2020)	0.07	0.04	0.01	0.19
Panel C: ICO enforcement data (per year)				
Monetary Penalty	5.29	6.95		
Enforcement Notice	2.57	2.15		
Improvement Action Plan	11.43	14.2		
One-Off Actions	478.14	334.92		
Advice	328.57	242.04		

Notes: This table shows summary statistics of the key variables in our data. Panel A shows variables from the job posting data. We report the weighted share of cyber job postings at the TTWA level using all job postings in the TTWA as weights. The share of cyber jobs is 0.58% per TTWA-industries-year (TTWA panel). The share is smaller (0.15%) per firm-TTWA-year (firm panel). Panel B reports variables constructed from firm data. We measure cash holding by cash and cash equivalent securities as a share of the total assets of the firm. Panel C shows the average number of ICO actions (by type) for business organizations that occurred between 2012 and 2018:Q2.

Table 2: Baseline Regressions

	Main enforcement exposure index		Alternative indices	
	Baseline	Only Services		
Dependent variable: Share of cyber job postings	(1)	(3)	(4)	(5)
Panel A : Main ICO exposure index: Based on number of firms targeted by the ICO				
High ICO exposure × Increased enforcement (2017-18)	0.264**	0.196*		
	(0.083)	(0.093)		
High ICO exposure × Increased penalty (2019-20)	0.535**	0.424**		
	(0.159)	(0.135)		
Panel B : ICO exposure index: based on number of ICO supervisory cases				
High ICO exposure × Increased enforcement (2017-18)			0.294**	
			(0.094)	
High ICO exposure × Increased penalty (2019-20)			0.519***	
			(0.145)	
Panel C: ICO exposure index: based on number of ICO supervisory cases weighted by severity				
High ICO exposure × Increased enforcement (2017-18)				0.301***
				(0.082)
High ICO exposure × Increased penalty (2019-20)				0.585**
				(0.192)
Industries $\times TTWA$	Yes	Yes	Yes	Yes
TTWA $\times Year$	Yes	Yes	Yes	Yes
Observations	144457	128100	144457	144457

Marginal effects

Table 3: Firm-Level Regressions

	Yearly	Average effects Over the two periods
Dependent variable: Share of cyber job postings	(1)	(2)
High ICO exposure ×2012	-0.003	
	(0.018)	
High ICO exposure ×2013	-0.003	
	(0.015)	
High ICO exposure ×2014	-0.034	
	(0.024)	
High ICO exposure ×2016	0.034***	
	(0.008)	
High ICO exposure ×2017	0.033*	
	(0.017)	
High ICO exposure ×2018	0.050***	
	(0.015)	
High ICO exposure ×2019	0.071**	
	(0.027)	
High ICO exposure ×2020	0.104**	
	(0.031)	
High ICO exposure × Increased enforcement (2016-18)		0.048**
		(0.018)
High ICO exposure × Increased penalty (2019-20)		0.095**
		(0.038)
Firm FE	Yes	Yes
TTWA $\times Year$	Yes	Yes
Observations	273488	273488

Note: Table 3 (also Figure 8) highlights the change in firm demand for cyber skills after the legal changes. Our dependent variable 'share of cyber job postings' calculates the number of job postings that requires any of the cyber skills to the total job postings of a firm in a specific year and local labor market (TTWA). High ICO exposure is a dummy that takes value one for the top quartile of the 3-digit industries in terms of the number of firms that are subject to ICO actions per industry. The standard errors are double clustered at three-digit SIC industry and year. *** and ** denote statistical significance at the 1 and 5 percent levels, respectively.

Table 4: Differential Response by Firms' Digital Technology Profile and Data Intensity

Dependent variable: Share of cyber job postings	All digital skills		Digital Organization		Data harvesting		Cloud	
	(1) Low b/se	(2) High b/se	(3) Low b/se	(4) High b/se	(5) Low b/se	(6) High b/se	(7) No b/se	(8) Yes b/se
High ICO-exposure industries × Increased enforcement (2016-18)	0.023	0.047*	0.038*	0.033	0.018	0.053**	0.023*	0.087
	(0.024)	(0.022)	(0.020)	(0.029)	(0.021)	(0.023)	(0.012)	(0.055)
High ICO-exposure industries × Increased enforcement (2019-20)	0.002	0.133**	0.027	0.124*	-0.003	0.142**	0.012	0.280**
	(0.042)	(0.049)	(0.031)	(0.056)	(0.039)	(0.051)	(0.032)	(0.104)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
TTWA $\times Year$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	105820	153017	117460	141391	113029	145807	195267	63521

Note: Table 4 indicates the effect of legal changes on firms' cybersecurity hiring across a subsample of firms with different digital technologies and data intensity. In column 1-6, we divide firms according to a data intensity index. Low and high respectively refer to firms that are below or above the median level of the data intensity measures. Columns (1) and (2) use the first principal component based on investment in any of three clusters of digital organization – SQL data management, enterprise resource planning (ERP), and cloud computing and storage – as well as five clusters of data harvesting skills – data mining, business intelligence (BI), extract load and transform (ETL), artificial intelligence (AI) and big data. Columns (3) and (4) use the first principal component based on investment in three clusters of digital organization. Columns (5) and (6) use the first principal component based on investment in five clusters of data harvesting skills. Columns (7) and (8) divide the sample into two groups of firms according to whether firms posted any job advertisements that require cloud-related skills between 2012 and 2015. In all columns, the dependent variable 'Share of cyber job postings' indicates the number of job postings that requires any of the cyber skills to the total job posting of a firm in a specific year-TTWA. High ICO exposure is a dummy that varies at the three-digit SIC industry classification. It takes one for the top quartile of industries that are highly exposed to ICO enforcement. The exposure index is built according to the number of firms that are subject to ICO actions. The standard errors are double clustered at the year and three-digit SIC industry level. *** and ** denote statistical significance at the 1 and 5 percent levels, respectively.

Table 5: Differential Response by Firms' Access to Internal Liquidity before Legal Changes

Dependent variable: Share of cyber job postings	Across All firms		Across Service firms		Within Sectors	
	(1) Low	(2) High	(3) Low	(4) High	(5) Low	(6) High
High ICO exposure × Increased enforcement (2016-18)	0.053**	0.065*	0.027	0.057*	0.054*	0.066*
	(0.021)	(0.031)	(0.018)	(0.031)	(0.027)	(0.031)
High ICO exposure × Increased penalty (2019-20)	0.071*	0.135**	0.039	0.126**	0.085*	0.125**
	(0.034)	(0.042)	(0.033)	(0.040)	(0.041)	(0.040)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
TTWA $\times Year$	Yes	Yes	Yes	Yes	Yes	Yes
Observations	149780	124585	131192	130936	143099	131256

Note: Table 5 indicates the effect of legal changes on firms' cybersecurity hirings across a subsample of firms with different levels of cash holding. In Panel A, we examine the differential effects according to cash holding in 2015 – the year before the first legal change takes place. Cash holding is measured by a firm's cash and cash equivalent securities over its total assets. In columns (1) and (2), we compare the firms above and below the median level of cash holding in 2015. In columns (3) and (4), we limit our sample to services and thus we divide firms based on the median level of cash holding for services in 2015. In columns (5) and (6), we divide firms according to the median level of cash holding for each two-digit SIC industry. In all columns, the dependent variable is the share of cyber job postings indicating the number of job postings that required any of the cyber skills to the total job postings of a firm in a specific year-TTWA. High ICO exposure is a dummy that varies at the three-digit SIC industry classification. It takes one for the top quartile of industries that are highly exposed to ICO enforcement. The exposure index is built according to the number of firms that are subject to ICO actions. The standard errors are double clustered at the year and three-digit SIC industry level. *** and ** denote statistical significance at the 1 and 5 percent levels, respectively.

Table 6: Differential Response over Firm Life Cycle and Firm Dynamics

	Cyber jobs share		Birth rate	Death rate
	Young Incumbent	Old Incumbent		
	(1)	(2)	(3)	(4)
High ICO exposure × Increased enforcement (2016-18)	0.067**	0.052**	-0.006	0.009**
	(0.023)	(0.021)	(0.004)	(0.003)
High ICO exposure × Increased penalty (2019-20)	0.138***	0.065	-0.014**	0.007**
	(0.039)	(0.037)	(0.006)	(0.003)
Firm FE	Yes	Yes	No	No
TTWA $\times Year$	Yes	Yes	Yes	Yes
Sector $\times Year$	No	No	Yes	Yes
Observations	121210	153130	133336	133336

Note: Columns (1) and (2) highlight the differential response for young and old incumbents for which we divide the firms according to the median firm age (15 years) in 2015. In these two columns, the dependent variable 'share of cyber job postings' indicates the number of job postings that requires any of the cyber skills to the total job postings of a firm in a specific year and TTWA. Column (3) and (4) indicate the effects of legal changes on sectoral firm turnover. In column (3), the dependent variable is the birth rate which measures the number of entries as a share of all registered firms each year at the three-digit industry-TTWA. In column (4), the dependent variable is death rate which indicates the number of exits as a share of all registered firms each year at three-digit industry-TTWA. High ICO exposure is a dummy that varies at the three-digit SIC industry classification. It takes one for the top quartile of industries that are highly exposed to ICO enforcement. The exposure index is built according to the number of firms that are subject to ICO actions. The standard errors are double clustered at the year and three-digit SIC industry level.*** and ** denote statistical significance at the 1 and 5 percent levels, respectively.