Hacking Corporate Reputations*

Pat Akey[†]	Stefan Lewellen[‡]	Inessa Liskovich[§]
Toronto	Penn State	UT Austin

November 13, 2018

Abstract

We exploit unexpected corporate data breaches to study how firms respond to negative reputation events. Data breaches negatively affect firm profitability and firm value for years following the event. In response, firms increase their investment in corporate social responsibility (CSR) by 0.4-0.5 standard deviations, suggesting that a reputation-induced decline in the value of existing CSR investments causes firms to increase their stock of CSR. Our paper represents the first empirical study to directly link CSR to corporate reputations and presents the first evidence in the literature that firms actively invest in CSR as the result of a negative reputation shock.

*We thank Steve Karolyi, Yrjo Koskinen, Souad Lajili Jarjir, Laura Starks, Duane Seppi, seminar participants at Carnegie Mellon, Toronto, and UT Austin, and participants in the 2017 LBS Summer Symposium early ideas session, the 2018 NBER Summer Institute IT & Digitization poster session, the 2018 NFA Meetings, and the 2018 Conference on CSR, the Economy, and Financial Markets for helpful comments. We also thank Avi Schiff for research assistance. Part of this research was conducted while Lewellen was visiting Carnegie Mellon University on leave from London Business School.

[†]University of Toronto. Phone: +1 (647) 545-7800, Email: pat.akey@rotman.utoronto.ca

[‡]Pennsylvania State University. Phone: +1 (814) 441-7151, Email: lewellen@psu.edu

[§]University of Texas at Austin. Phone: +1 (512) 232-6825, Email: Inessa.Liskovich@mccombs.utexas.edu

1 Introduction

How do firms respond to the destruction of capital? A large literature in economics suggests that events such as natural disasters often produce significant asset losses that can spill over to other sectors of the economy (Hallegatte and Vogt-Schlib, 2016; Barrot and Sauvagnat, 2016). In response to these events, anecdotal evidence suggests that firms often attempt to rebuild damaged factories, damaged offices, and damaged stores. For example, after sustaining significant damage to one of its stores during Hurricane Harvey, the Texas grocery chain H-E-B told the *Houston Chronicle* that it expected to open a new store nearby and was already in "active negotiations for another site."[1]

However, while much is known about the destruction of *tangible* capital, less is known about how firms respond to the destruction of *intangible* capital.[2] For example, if an idiosyncratic event negatively harms a company's reputation, can the firm take actions to rehabilitate its reputation? We would expect firms to rebuild intangible capital so long as the marginal benefits of such actions exceed the marginal costs. However, unlike rebuilding a store, which involves specific, concrete actions, little is known about how a firm may rebuild its stock of intangible capital.

In this paper, we focus on a specific type of intangible capital – namely, a company's reputation. We exploit unexpected data breaches as a negative shock to corporate reputations and examine how firms respond to these negative corporate reputation shocks. Anecdotal evidence suggests that data breaches can have a large, negative impact on corporate reputations. For example, Deloitte's 2014 global survey on reputation risk found that executives

[1]Sternitzky-Di Napoli, Daniela and Katherine Blunt,"H-E-B location closes after sustaining significant damage during Hurricane Harvey," *Houston Chronicle*, October 5, 2017.

[2]There is a literature examining how firms respond to the destruction of human capital. For example, Jäger (2016) examines the ability of firms to replace incumbent workers following worker deaths and finds that labor market frictions can prevent these workers from being fully replaced. To our knowledge, however, there is no evidence on how firms respond to the destruction of other types of intangible capital.

believe cyber-security risks are a bigger threat to corporate reputations than product safety or customer service issues.[3] A 2014 survey by the UK fraud prevention firm Semafone found that such fears are not misplaced: more than 86% of consumers surveyed stated that they were not likely to do business with a firm that suffered a data breach involving personal financial information such as credit card numbers.[4] Hence, unexpected data breaches have the potential to destroy a significant fraction of a firm's intangible capital.[5]

To examine firms' responses to the destruction of intangible capital, we focus on firms' corporate social responsibility (CSR) following unexpected data breaches. Our focus on CSR investment is driven by three main considerations. First, unlike investments in new IT systems or new data protection policies, firms' investments in CSR are publicly observable. Second, there is a strong theoretical link between CSR and corporate reputation (Albuquerque, Durnev, and Koskinen, 2013; Heal, 2005; Kitzmueller and Shimshack, 2012), suggesting that a negative reputation event would likely affect the value of a firm's existing stock of CSR and the marginal benefit of new CSR investment. Finally, CSR investment is ubiquitous across industries, allowing us to abstract away from any industry-specific responses that firms might take following an unexpected data breach.

Our main hypothesis is that firms will respond to the destruction of intangible capital caused by unexpected data breaches by increasing their investment in CSR. Intuitively, our argument is that negative reputation shocks reduce the value of firm's existing investments in CSR, thereby leaving the firm below its optimal level of CSR. Under the assumption of decreasing returns, a reduction in the value of the firm's existing stock of CSR will lead to an increase in the marginal benefit of additional investment. Finally, since the marginal

[3] http://deloitte.wsj.com/riskandcompliance/2015/01/21/security-attacks-a-lead-driver-of-reputation-risk/

[4] https://semafone.com/press-releases-us/86-customers-shun-brands-following-data-breach/?lang=us

[5] Both theoretical arguments (Kreps, 1990; Tadelis, 1999) and empirical evidence (Armour, Mayer, and Polo, 2017; Karpoff, Lee, and Martin, 2008; Murphy and Shrieves, 2009) suggest that negative reputation shocks can be costly.

costs of CSR investment are likely to remain unchanged, the increased marginal benefits from investment will cause firms to increase their investment in CSR following negative reputation events. The existing CSR literature has focused on whether firms' *pre-existing* CSR investments can help to mitigate the effects of negative corporate events such as product recalls, oil spills, or financial crises, but this literature does not examine how firms respond to such events. To our knowledge, our paper represents the first attempt to evaluate whether firms invest in rebuilding their intangible capital following an unexpected negative shock.

Our first set of results confirms that a high pre-existing stock of CSR can mitigate the effects of negative corporate events. We find that while firms affected by a material data breach experience cumulative abnormal returns (CARs) of -1.5% to -1.9% in the 30 days following the disclosure of the breach, this effect is much smaller at firms with a greater pre-event stock of CSR. Hence, consistent with the existing literature on CSR, we find that negative reputation shocks are at least partially insurable through mechanisms such as CSR.

We then depart from the existing literature by examining how firms respond to a negative reputation shock. If reputation shocks affecting firm value or firm profitability are short-lived, managers would probably *not* respond to such shocks by investing in reputation-building activities such as CSR, since such investments are typically costly and often take years to materialize. However, if managers expect data breaches to have long-lasting effects on firms' reputations, they might take steps to invest in costly activities like CSR in order to rebuild the firm's reputation among customers, employees, shareholders, and other stakeholders.

Consistent with the latter hypothesis, we find that the negative reputation effects of corporate data breaches have long-lasting effects on firm value and firm profitability. One-year changes in the market-to-book ratio (M/B) for affected firms are approximately -10 to -20% relative to firms that did not experience a data breach. Decomposing the M/B ratio

into ROE and the price-to-earnings (P/E) ratio shows that *both* current profitability and expectations about future profitability are affected following a breach. These reputation shocks still affect ROE and P/E at least four years after the disclosure of a data breach. Hence, while a high pre-existing stock of CSR seems to mitigate the market's initial reaction to data breaches, the long-lasting negative effects of data breaches on firm value suggest that firms might respond to such breaches by making additional investments in reputation-enhancing activities such as CSR.

Indeed, our next set of findings confirms that affected companies significantly increase their CSR investment in the years following a data breach. Relative to unaffected firms, we find that affected firms' CSR scores increase by approximately 0.4-0.5 standard deviations in the four years following the disclosure of a data breach. In addition, following a breach, firms might be expected to incur one-off expenses associated with attempts to repair their reputations (such as the expenses associated with increased investment in CSR). Under normal accounting practices, these types of expenses would be recorded as non-recurring items. Consistent with this hypothesis, we find that affected firms have similar (scaled) sales and EBITDA as unaffected firms, but report significantly higher non-recurring expenses in the four years following the breach.

We also examine different types of data breaches to determine whether certain types of reputation shocks are more costly to firms than others. In particular, most data breaches affect either customer records or employee records. While value declines are more pronounced for data breaches involving customer records, we find that firms subsequently increase CSR in cases involving both types of breaches. Finally, we examine firms' post-breach investments in different forms of CSR. We find that affected firms predominantly increase their investment in the "environment" and "diversity" CSR categories (relative to unaffected firms). Under the assumption that firms are optimizing their CSR investments, this result suggests that

firms believe that better environmental policies and diversity policies represent the CSR investments that provide the largest potential reputation gains conditional on their costs and implementation time.

We focus on data breaches because these events arguably affect corporate reputations while being plausibly unrelated to firms' product quality or financial condition. Anecdotal evidence suggests that firms' investments in cyber-security are driven in part by reputational considerations: for example, according to a 2016 survey by the *Economist Intelligence Unit*, C-level executives listed corporate reputation as the single most important company asset requiring protection from cyberattacks.[6] In addition, data breaches are largely idiosyncratic, the timing of such breaches is plausibly random, and except in rare cases, the breaches themselves do not specifically affect the quality of the products of services offered by the affected company. As such, data breaches offer a number of empirical advantages over other types of negative reputation shocks such as product recalls, health or safety violations, fraud, other types of employee misconduct, fines, legal settlements, or critical product reviews.

Nonetheless, it is still possible that latent firm characteristics could be responsible for both the incidence of data breaches and subsequent value losses or firm responses. For example, the market may interpret news about a data breach as containing new information about managerial competence or firm governance rather than containing information about a company's reputation. However, we find that controlling for the *G*-index of Gompers, Ishii, and Metrick (2003) and the *E*-index of Bebchuk, Cohen, and Ferrell (2009) does not change any of our main results. Furthermore, we find that CEOs of affected firms are *less* likely to leave following a data breach relative to control firms. Finally, we find that the market responds differently to different types of data breaches, which is inconsistent with

[6]The survey sampled C-suite members from firms across 16 countries and a variety of industries. See https://perspectives.eiu.com/sites/default/files/images/EIU-VMware Protectingthebrand_PDF.pdf.

arguments related to governance or managerial competence.[7] Collectively, these findings make it unlikely that latent firm characteristics are driving our results.

One might also be concerned that cyber-attackers select target firms on the basis of other (possibly unobservable) variables that are correlated with the outcome variables that we study. However, we do not find evidence that variables such as corporate governance or firm value are predictors of whether a firm experiences a data breach. We also perform nearest-neighbor matching (within industry, within time) for each affected firm and confirm that our main findings persist even after restricting the example to only include affected firms and their closest substitutes. Hence, it is unlikely that targeted cyber-attacks can fully explain the differences we observe between affected and unaffected firms.[8]

Our paper makes three primary contributions to the existing literature. First, an oft-considered theoretical motivation for CSR investment is that CSR guards against reputation risk, all else equal (Albuquerque, Durnev, and Koskinen, 2013; Heal, 2005; Kitzmueller and Shimshack, 2012).[9] However, we are unaware of any empirical studies that specifically attempt to isolate corporate reputation as a motivation for firms' investment in CSR.[10] To our knowledge, our paper is the first to document direct investment in CSR as a response to a negative reputation shock.

Second, the idea that CSR can provide "insurance" against negative shocks has also been the subject of many empirical studies (Barrage, Chyn, and Hastings, 2014; Godfrey,

[7] In particular, it is not clear why a breach involving (say) customer records would be more indicative of incompetence or poor governance than a breach involving employee records.

[8] However, if selection bias were to exist, the magnitudes of our results would potentially be different than the magnitudes from a fully randomized sample.

[9] A large literature has also established a positive empirical link between financial performance and CSR (Margolis, Elfeinbein, and Walsh, 2009; Edmans, 2011; Flammer, 2015). Researchers have also introduced and investigated a number of theories for why firms may choose to engage in strategic CSR (see, e.g., Bénabou and Tirole (2010)).

[10] In contrast, existing empirical studies have focused on the preferences and political ideologies of investors and managers to explain firms' CSR investments (Cheng, Hong, and Shue, 2013; Dyck, Lins, Roth, and Wagner, Forthcoming; Giuli and Kostovetsky, 2014; Hong and Kostovetsky, 2012). The closest paper to ours is arguably (Servaes and Tamayo, 2013), who find that the benefits of CSR appear to be concentrated among consumer-focused firms.

Merrill, and Hansen, 2009; Hong and Kacperczyk, 2009; Hong and Liskovich, 2014; Lins, Servaes, and Tamayo, 2017; Vanhamme and Grobben, 2009).[11] However, these studies examine whether a firm's *existing* stock of CSR can help the firm when it experiences a negative shock, such as an oil spill, misconduct, regulatory actions, or a financial crisis. In contrast, we exploit a negative shock to firms' reputations and then examine whether firms *replenish* their stock of CSR in response to the event. Relative to the studies above, our paper also exploits a setting in which it is less likely that a negative corporate reputation shock is contaminated with other fundamental news about the firm's current or future business prospects.

Finally, the existing empirical literature on corporate reputations has focused on negative reputation shocks such as financial misconduct (Armour, Mayer, and Polo, 2017; Karpoff, Lee, and Martin, 2008; Murphy and Shrieves, 2009; Chakravarthy, DeHaan, and Rajgopal, 2014), environmental violations (Karpoff, Lott Jr, and Wehrly, 2005) and product recalls (Jarrell and Peltzman, 1985; Liu and Shankar, 2015). The evidence from these papers suggests that reputation losses are large and in some cases overshadow direct legal penalties. Our paper adds to this literature by exploring a new type of reputation shock (data breaches) that is arguably less likely to also contain fundamental information about a company's prospects or the threat of future regulatory actions.

Our paper is not the first to examine corporate data breaches. In the computer science literature, Acquisti, Friedman, and Telang (2006), Campbell, Gordon, Loeb, and Zhou (2003), and Spanos and Angelis (2016) document significant negative short-term stock market reactions to corporate data breaches. In a contemporaneous paper, Kamiya, Kang, Kim, Milidonis, and Stulz (2018) also examine the cross-sectional effects of unexpected cyber-attacks using a wide range of sorting variables including size, firm value, sales growth, in-

[11] Experiments have also shown that consumers respond positively to sellers that engage in charity (Elfenbein, Fisman, and McManus, 2012).

vestment, asset tangibility, governance and risk management, leverage, credit ratings, and cash flow volatility. Kamiya et al. (2018) also examine firm responses to data breaches – in particular, responses to CEO pay and board-level risk management metrics. In contrast, our paper specifically focuses on CSR (not studied by Kamiya et al. (2018)) and finds many results that differ from the results in Kamiya et al. (2018), possibly due to the use of a different sample of data breaches. Finally, Makridis and Dean (2017) also examine whether corporate investment changes following data breaches. We add to this literature by showing that data breaches have *long-term* negative firm value effects and by examining how firms attempt to rebuild their reputations following a data breach through investing in reputation-building activities such as CSR.

2 Data

2.1 Corporate Data Breaches

We obtain data on corporate data breaches from the Privacy Rights Clearinghouse (PRC) website.[12] The PRC is a non-profit foundation that advocates to educate consumers about privacy protection. In addition to providing educational services, it has compiled a database of publicly disclosed data breaches starting in 2005. We download the list of breaches that affected private organizations (as opposed to government agencies or universities) and match these organizations to publicly traded firms.

Panel A of Figure 1 presents the frequency of data breaches over time, while Panel B presents the breakdown of events by two-digit Global Industry Classification Standard (GICS) industries. The PRC data includes information about which firms were affected by the breach, a short description of the breach and, when available, the number of records

[12]https://www.privacyrights.org/

that were affected. We classify hacks as affecting customer records (such as account information or personal details), employee records, or internal company documents when such information is available. Panel A of Table 1 reports summary statistics for these variables. Overall, there are 287 data breaches, of which the vast majority represent breaches involving customer records (66%) or employee records (33%). Panel A of Figure 2 presents a kernel density plot of the natural log of the number of records breached over the sample, while Panel B plots the log of the average number of records over time. These data breaches can take several forms, including external hacks, lost or stolen portable devices, insider employees improperly accessing data, physical theft of documents containing information, and inadvertent disclosure of sensitive information. Figure 3 presents summary statistics on the type of breaches in our sample. Panel A presents the fraction of different beach types, while Panel B presents the proportion of records that are affected by different types of breaches. While hacks, the loss of portable electronics and insider breaches account for the largest fraction of data breaches, the overwhelming majority of records are breached through external hacks.

2.2 Firm Data

We obtain data on firm returns from CRSP and firm fundamentals from COMPUSTAT. To maximize the probability of comparing similar types of firms in our tests, we constrain our main sample to focus on firms in industries that were hacked at least once over our sample period (based on six-digit GICS classifications). We also construct a smaller, matched sample using a nearest-neighbor matching technique.

We measure corporate social responsibility using the widely-used MSCI ESG KLD Stats measure of CSR. These scores are developed to provide an independent assessment of firms' social responsibility, similar to the manner in which credit rating agencies assign credit ratings. To calculate the score, MSCI first determines the presence or absence of a series

of social responsibility "strengths" or "concerns" within a firm. The score itself is an index that sums all the strengths and subtracts all the concerns. Therefore a one point increase in the CSR score requires a firm to change one corporate social responsibility category from a concern to neutral, or from neutral to a strength. The score can be further broken down into several dimensions of CSR: community relations, product characteristics, environmental impact, employee relations, diversity, and governance.

We make several modifications to the CSR score to account for the fact that the calculation of the index has changed over the years. The individual strengths and concerns making up the index have variously been added, deleted, and redefined. Therefore the index itself has not referred to a consistent set of actions over time. This is especially problematic around 2009, which saw a large redefinition in index components. To ensure that we study a consistent measure of CSR, we create a time-consistent index. We take the following three steps: (1) we match indicators that changed names but covered the same concepts over time, (2) we only use the indicators that are covered from 1991 through 2015, and (3) we limit our index to those indicators that were non-missing for the full sample in 2010, following the major redefinition. This leaves us with a time-consistent CSR score that we use in our analysis. It is made up of eighteen strengths and six concerns.[13] For ease of interpretation, we normalize the measure to have a mean of 0 and a standard deviation of 1 throughout the sample. We refer to this measure in our tables as "Norm CSR." Panel B of Table 1 reports summary statistics for fundamentals and standardized CSR scores for the sample of firms that suffered a data breach that impacted at least 1,000 records (which represents the main sample that we use for most of our tests).

[13]We do not include any governance measures as there are none that meet the criteria for time consistency.

3 Empirical Strategy

3.1 Short-term Valuation Effects

We first examine short-term stock market reactions following the disclosure of a data breach. We measure cumulative abnormal returns (CARs) using a 100-day estimation window that ends 50 days before a breach is publicly disclosed. Expected returns are estimated using the Fama-French three-factor model. We estimate CARs for a number of different event windows: [-1,3], [-1,5], [-1,10], and [-1,30], where the numbers refer to trading days relative to the date on which the data breach is disclosed. Our identifying assumption in these tests is that the revelation of a data breach is not correlated with a firm's expected return after controlling for the Fama-French factors.

To better understand heterogeneity within stock market reactions, we also run regressions of the form:

$$c_{it} = \alpha + \beta_b b_{it} + \beta_x x_{it} + \varepsilon_{it} \ ,$$

where c_{it} represents the CAR for that data breach for firm i and time t, b_{it} captures data breach characteristics, and x_{it} represents time-varying firm characteristics.

3.2 Long-term Valuation Effects

To test the effects of data breaches on longer-term firm value and subsequent firm responses, we construct an annual panel of all firms from 1999 to 2015. As stated previously, we limit our sample to firms in six-digit GICS industries that have at some point experienced a data breach. This restriction eliminates 37% of the firms in COMPUSTAT. In order to study firms in the wake of data breaches, we create an indicator variable, *Post*, that identifies the

firm-year observations following the disclosure of breach. Our main specification is:

$$y_{ijt} = \alpha + \gamma Post_{ijt} + \beta x_{ijt} + f_{jt} + f_i + \varepsilon_{ijt} \ .$$

The variable y_{ijt} captures annual firm outcomes such as the market-to-book ratio and CSR for firm i in industry j in year t. Industries are defined using six-digit GICS classifications. Time-varying firm characteristics are captured by x_{ijt} which controls for ln(Assets), ln(Assets)^2, and market leverage. We use two different definitions of $Post_{ijt}$ to capture responses over different time periods. One definition includes the year of the data breach along with the subsequent year. Therefore $Post_{ijt}$ is equal to zero for all firms that were never subject to a data breach as well as firms that did not experience a breach within the previous two years. The second definition includes the breach year plus four years following the breach. Finally, f_{jt} represents industry-by-year fixed effects, and f_i represents firm fixed effects. The inclusion of firm fixed effects ensures that our identification controls for any time-invariant characteristics that differ across affected and unaffected firms. The industry-by-year fixed effects ensure that our comparisons are within a given industry, within a given year, between affected and unaffected firms.

While our preferred empirical specifications include both firm and industry-by-year fixed effects, we also ensure that our results hold without firm fixed effects. To this end, we estimate the specification:

$$y_{ijt} = \alpha + \gamma Post_{ijt} + \delta Treated_{ij} + \beta x_{ijt} + f_{jt} + \varepsilon_{ijt}$$

where $Treated$ identifies whether a firm has ever been subject to a data breach. All other variables are defined as before.

Although our main specifications control for firms' time-invariant unobservable char-

acteristics, there may be time-varying characteristics that correlate with outcomes. To limit these confounding factors, we also estimate a matched specification using nearest-neighbor matching. For every data breach, we match the affected firm to nine firms in the same six-digit GICS industry and year, for a total of ten firms per breach. We choose the firms that are closest in size, CSR, and market-to-book ratios, as measured by the sum of absolute normalized distance along those three variables. We then estimate the stacked difference-in-difference specification:

$$y_{ite} = \alpha + \gamma Post_{ite} \times Treated_{ie} + \delta Treated_{ie} + \theta Post_{ite} + f_e + \varepsilon_{ite}$$

Here, the variable $Post_{ite}$ is defined for treated as well as control firms. It simply identifies the years (either 0-1 or 0-4) following the data breach for all matched firms. The main variable of interest is the interaction term $Post_{ite} \times Treated_{ie}$. For these tests, we include a fixed effect for each data breach, captured by f_e.

Our identifying assumption in these tests is that data breaches constitute an exogenous negative shock to a firm's reputation. In particular, the vast majority of data breaches are arguably unrelated to a firm's products or services. Instead, they draw negative attention to the firm and may influence the firm's reputation among consumers, suppliers, partners, and current (or future) employees. For example, consumers' credit card numbers and passwords are extremely valuable to hackers, but do not directly affect the products or services offered by a firm.

However, it may be the case that the disclosure of a data breach conveys other information to investors, such as information about the competence of the firm's management team. To examine this possibility, we first include controls for firm governance in our tests. We also test for differences in CEO turnover propensities between affected and unaffected firms.

Finally, we perform a series of tests to ascertain whether heterogeneity exists across different types of data breaches (such as employee records versus customer records). In particiular, data breaches involving (say) employee records or customer records should send similar public signals about executive competence. Hence, if we find that investors respond differently to different types of data breaches, this would cast doubt on the idea that data breaches only contain information about variables such as corporate governance or the competence of the firm's management team.

Another concern is that the firms subjected to data breaches differ from unaffected firms, potentially in unobservable ways. For example, firms that under-invest in data security may under-invest in other aspects of their business as well. This explanation is unlikely, however, as interviews with cybersecurity practitioners highlight that all firms are exposed to the risk of data breaches. As explained by Dave DeWalt, the CEO of a cybersecurity firm, data breaches are inevitable for nearly every company: "Even the strongest banks in the world: banks like JPMorgan, retailers like Home Depot, retailers like Target can't spend enough money or hire enough people to solve this problem."The CTO of Vodafone Enterprise Security Services further emphasized that "[companies] need to change their mindset away from just protecting [their data] but enabling their businesses to thrive in a world where they are constantly under attack."

However, we address the above concern by focusing on within-firm variation. By including year fixed effects we ensure that we compare years following a data breach to the same firm at a different point in time. We also include observations on firms that were never hacked in order to better estimate year-by-industry fixed effects. In this specification the necessary identifying assumption is that there are no omitted time-varying firm characteristics that covary with the probability of a data breach. Some firms may be more vulnerable to data breaches than others, but we would not expect their vulnerabilities to vary over

time in a predictable way. We believe this assumption is plausible, particularly given that a firm's information technology infrastructure is difficult to change and requires long-term investment. Moreover, we use nearest neighbor matching to verify that our main results are robust when we compare firms affected by data breaches to observably similar firm that have never been affected.

Finally, we limit our main analysis to data breaches affecting at least 1,000 records. In many cases, the number of records affected by a breach is not reported. While excluding these breaches reduces our sample size by approximately 50%, this restriction allows us to focus on data breaches that are arguably the most similar in nature, helping to reduce the risk that our results are spuriously driven by a few "outlier" breaches that are not representative of data breaches in general.

4 Results

4.1 What Characteristics Correlate with Data Breaches?

We first examine whether certain types of firm characteristics are more likely to be associated with a data breach. In Table 2, the outcome variable of interest is whether a given firm will be subject to a data breach with at least 1,000 records compromised. Column (1) utilizes our full sample, whereas column (2) include only those firms with a non-missing CSR score and column (3) limits the sample to firms with non-missing governance scores. All firm characteristics are measured as of the year prior to the data breach disclosure. We find that firms are more susceptible to data breaches when they are smaller and have a lower market-to-book ratio. A firm's profitability does not seem to be systematically related to the likelihood of experiencing a data breach. Similarly, corporate social responsibility, as measured by normalized CSR score, is not associated with significant differences in the occurrence of data

breaches.[14] Finally, in column (3), we include measures of firms' corporate governance from Gompers, Ishii, and Metrick (2003) and Bebchuk, Cohen, and Ferrell (2009) (the "*G*-index" and "*E*-index," respectively), and find that the likelihood of a data breach is unrelated to a firm's corporate governance practices.

4.2 Data breaches and Firm Value

4.2.1 Short-term returns

We begin our analysis of the impact of data breaches on firm value by examining the stock market reactions to the disclosure of the data breaches. Panel A of Table 3 presents Fama-French three-factor CARs for a variety of windows around the disclosure of the data breach. Panel B presents CARs for breaches where at least 1,000 records were impacted, which constitutes our main sample of interest in the remainder of the paper. Figure 4 plots CARs starting 10 days before the breach event through 30 days following the event for all breaches for which the number of impacted records is known. We do not find any evidence of abnormal returns in the 10 days prior to the disclosure of the hack. However, there are large and growing negative price reactions to the announcement of a data breach.

In the sample of all breaches, CARs range from -0.64% from the day before the breach through the first three days after the event to -1.5% for the window ranging from one day before the event to 30 days after the event. Unsurprisingly, the effects are stronger for more important breaches, which we define as those impacting at least 1,000 records. The CARs for this sample range from -0.9% measured one day before the breach to three days after the breach to -1.92% over a one-month window.

We next explore cross-sectional determinants of the market reactions reported previ-

[14]The coefficient on CSR implies that a one-standard deviation change in CSR has one third of the effect of a one-standard deviation change in market-to-book.

ously. Table 4 reports regressions where the dependent variable is a firm's CAR from one day before to thirty days after the data breach. In columns (2) – (6) we test whether stock market reactions are related to attributes of the data breach. Consistent with our findings above, we find that larger data breaches, as measured by the natural logarithm of the number of records impacted, are associated with a more negative response. However, we do not find robust evidence that there are different short-term reactions for data breaches involving customer records or employee records.

In columns (3) – (6) we introduce firm characteristics as of the year prior to the data breach disclosure. The indicator for "high CSR" identifies firm-years in which a firm's CSR score is above the industry-year average. Consistent with the existing literature on the "reputation insurance" effects of CSR, we find that firms with high CSR scores in the year before the data breach have smaller negative abnormal returns. The magnitude of the effect suggests that stock prices are effectively flat for firms with above-average CSR, while market reactions are negative and large for firms with below-average CSR scores. We next introduce size, size-squared, leverage and governance controls in columns (4) – (6) and find that larger firms were less negatively impacted by data breaches, firms with higher leverage were more strongly impacted, and governance does not appear to affect the size of the market reaction.

4.2.2 Long-term Value Effects

We further explore how data breaches affect firm valuations by examining how firms' M/B ratios change in the years following a data breach. If data breaches only affect firm value in the short run, there would be no need for managers to invest in costly reputation-building activities such as CSR that often take years to implement. However, if declines in firm value are persistent, managers would have a stronger incentive to invest in lengthy projects to improve the firm's reputation.

To better understand the drivers behind firm value changes, we also decompose the M/B ratio into the return on equity (ROE) and the price-to-earnings ratio (P/E). Changes in ROE capture how firms' current performance is impacted by the data breach, while changes in P/E ratios capture how market participants view the impact of the data breaches on firms' longer term growth opportunities and growth options. Hence, this decomposition allows us to examine whether value changes are primarily driven by changes in short-run firm profitability or changes in long-term market expectations.

Table 5 presents the results of this analysis. Specifically, Panel A of Table 5 examines how firms' M/B, ROE, and P/E change in the two years following a data breach, while Panel B presents the results for the four years following a data breach. Columns (1) – (2) of both panels study changes in the log M/B ratio, columns (3) – (4) study changes in ROE, and columns (5) – (6) study changes in firms' P/E ratios. The control group in all tests consists of other firms in industries that at some point suffered a data breach. All columns contain industry-by-year fixed effects, and columns (2), (4), and (6) contain firm fixed effects. Finally, all columns include controls for size, size-squared, and the firm's market leverage.

We find strong evidence that *long-term* firm value declines following unexpected data breaches. In the two years following a data breach, affected firms' M/B ratios decline by nearly 12% relative to unaffected firms in the same industries. We also find that *both* firms' current profitability and firms' expected growth opportunities are negatively impacted by the data breaches. For example, in the two years following the event, ROE declines by 3% to 6% and P/E ratios decline by 3.13 to 3.38, although the ROE results are statistically mixed. The economic magnitudes of these results are substantial: the sample standard deviation of ROE is 0.525, suggesting that the across-firm, two-year coefficient is 12.5% of a standard deviation, whereas the standard deviation of P/E ratios is 19.4, suggesting that the across-firm, two-year coefficient is 18% of a standard deviation. Hence, data breaches appear to

have long-term, lasting effects on firm value.

Panel B extends our analysis to consider four years following each data breach. Again, we find evidence that M/B ratios decline, ROE declines, and P/E ratios decline. The decline in M/B loses its statistical power in our most stringent specification, but the economic magnitude of this effect (-6.5%) is still large. The magnitudes of the declines in ROE and P/E are similar to those reported above, and these results are all statistically significant. Hence, even *four years* after a data breach, firm value, current profitability, and expectations about future profitability are lower for affected firms relative to unaffected firms in the same industries.

4.2.3 Other effects

We next dig deeper to better understand how firms' current profitability is affected by data breaches. Intuitively, data breaches may lead to a reduction in sales or an increase in operating costs. However, ROE could also fall due to the existence of non-recurring expenses associated with increases in firms' one-time investments in technologies such as CSR.

Table 6 reports the results of our analysis. As before, Panel A present shorter-term responses (event years zero and one), while Panel B presents longer-term responses (event years zero through four). Columns (1) and (2) suggest that firms' sales do not seem to be directly impacted by the data breaches. Columns (3) and (4) suggest that corporate EBITDA declines modestly at a the shorter-term horizon, although the results are statistically mixed. Finally, in specifications (5) and (6), we find that firms are more likely to report non-recurring items on their income statements, which we interpret as an increased likelihood to engage in initiatives to rehabilitate their corporate reputation. For example, over the four years following the data breach, firms are 8.7% more likely to report a non-recurring (one-time) item, which compares to an unconditional pre-breach likelihood of 37.4%. We interpret these

results as suggesting that firms are taking discrete, one-time actions to respond to a data breach.

4.3 CSR Investment Following Data Breaches

Our results thus far suggest that data breaches have a major effect on both short-term and long-term value. However, it is an open question whether (and how) firms attempt to respond to these negative shocks.

Our main hypothesis is that conditional on experiencing a negative reputation shock, firms will undertake costly investments in activities such as CSR in order to rebuild their reputations. In particular, while firms may have undertaken CSR investments in the past, the occurrence of a negative reputation shock may effectively reduce the value of the firm's prior investments in CSR. As such, in order to repair their reputations, firms would have to invest *even more* in CSR following a negative shock.

This hypothesis is tested directly in Table 7. While market-based valuations update instantaneously, changes in CSR scores are likely to occur with a lag since firm actions to change their CSR may take a long time to be implemented. For example, in order to change an environmental impact "strength" indicator from a zero to a one, a company would need to implement a "notably strong pollution prevention [program], including both emissions reductions and toxic-use reduction programs." Hence, although a firm may invest in developing such a program immediately after a negative reputation shock, it would likely take a few years for the program to become active.

In specifications (1) – (5), we analyze how firms' CSR scores change in the two years spanning the data breach, while in specifications (6) – (10), we analyze how firms' CSR changes over the four years following the data breach. We find little evidence that firms' CSR scores change in the shorter horizon following the data breaches. However, we find large

increases in CSR scores in the four-year window following a data breach. For example, in specification (10), the most saturated regression model, we find that CSR scores increase by 39% of a standard deviation, which is an economically large effect. Hence, consistent with our main hypothesis, we find that firms respond to a negative reputation event by increasing their investment in reputation-building activities such as CSR.

We further study which components of CSR are most likely to receive additional investment following a data breach. Specifically, we split firms' CSR scores into the five (standardized) components that form our composite measure of CSR — community relations, product characteristics, environmental impact, employee relations, and diversity — and repeat the analysis of our most stringent specification in Table 7. Results of this analysis are reported in Table 8. We find that the increase in CSR scores following data breaches primarily comes from improvements in firms' environmental and diversity scores (specifications (1) and (4)). Indeed, we find that CSR environmental impact scores increase by 60% of a standard deviation, while diversity scores increase by 33.5% of a standard deviation. These results are consistent with the theoretical findings of Albuquerque, Durnev, and Koskinen (2013) that environment and diversity components have the largest impact on firms' subsequent financial statements.[15]

One plausible concern with our CSR results is that firms that are more likely to be hacked might be run by poorly performing or entrenched managers (although Table 2 shows that governance variables do not seem to predict the incidence of data breaches). In this scenario, firms might respond to a data breach by taking actions such as firing their CEO or otherwise changing their governance practices in a manner that might correlate with the firm's CSR score. We examine this possibility in two ways. First, we re-estimate our main CSR regressions after controlling for a firm's corporate governance (using the governance

[15]Similarly, Bouslah, Kryzanowski, and M'Zali (2013) show that within large firms, firm risk is correlated with the employee, diversity, and governance components of CSR.

indices proposed by Gompers, Ishii, and Metrick (2003) and Bebchuk, Cohen, and Ferrell (2009)). Second, we directly examine the likelihood of CEO turnover following data breaches. Columns (1) – (4) of Table 9 presents our augmented CSR regression results. Columns (1) – (2) present the results for normalized CSR over the two years following the data breach, while columns (3) – (4) present the results over the 4 years following the data breach. Our results are very similar to the results we report in Table 7. Hence, the changes we observe in CSR do not seem to be significantly correlated with changes in governance measures.

We next examine CEO turnover using data from ExecuComp. On the one hand, firms might be more likely to fire poorly performing CEOs following a data breach if the breach provides information to the board or shareholders about the competence of the firm's management. However, if breaches do not contain a negative signal about managerial competence, firms might actually be *less* likely to fire a CEO following a data breach since fallout from the breach itself would require a significant amount of managerial time. To examine CEO turnover following data breach disclosures, we construct a binary variable that takes the value of one if there is a CEO turnover event (for any reason) in a given year and zero otherwise. Columns (5) – (8) of Table 9 contain the results of our tests. Columns (5) – (6) examine the likelihood of CEO turnover in the two years following an event, while columns (7) – (8) examine the likelihood of CEO turnover over the four years following an event. We find that CEOs are 7% to 13% *less* likely to leave the firm in the two to four years following a data breach. This result supports the argument that data breaches are idiosyncratic events that are unrelated to the competence of the firm's CEO.

4.4 Heterogeneity in Responses

Firms may also respond differently to different types of data breaches (for example, breaches involving customer records versus employee records). To examine this possibility, we re-

estimate our main results on M/B, ROE, P/E, and standardized CSR scores separately for breaches involving customer records and employee records. Table 10 reports the results of these tests. We find that firm valuations decrease for both types of data breaches (specifications (1) – (2)), although the effect is smaller in magnitude and less statistically significant for breaches that impact employee records. Specifications (3) – (6) show that ROE and P/E ratios appear to respond similarly in magnitudes for both types of data breaches. Finally, specifications (7) and (8) show that firms respond to both types of breaches by increasing CSR, although this response is statistically larger following breaches that impact employee records. These results suggest that the responses documented previously are robust across various types of data breaches. The fact that markets appear to respond somewhat differently to different types of data breaches also casts doubt on the idea that breaches contain value-relevant information about managerial competence or corporate governance, since it is not clear why a breach involving customer records would reveal more about managerial incompetence or firm governance than a breach involving employee records.

Interestingly, the heterogeneity observed in Table 10 suggests that the same negative shock may affect perceptions about a firm's reputation differently across different groups of stakeholders. For example, it is natural to think that a data breach involving employee records may reduce a company's reputation among current and prospective employees more so than among a company's consumers. Given this assumption, the strength of firms' investment in CSR following different types of data breaches is informative about the extent to which CSR investments (versus other types of investments) actually improve a firm's reputation. In particular, the results in Table 10 suggest that the marginal improvement in a company's reputation following a dollar of additional CSR investment may be larger among employees rather than customers. Hence, a company wishing to rehabilitate its reputation may be more likely to invest in CSR (versus other types of investments) if it believes its

reputation has suffered more among current or prospective employees than among current or prospective consumers.

4.5 Robustness

One possible concern with our results is that firms that are affected by a data breach may be significantly different than unaffected firms. If this is true, our results could simply be capturing latent differences across firms rather than the effects of negative corporate reputation shocks. To address this concern, we re-estimate our main results (M/B, ROE, P/E, and CSR) using a nearest-neighbor matching technique that pairs affected firms with control firms that share similar observable characteristics. Specifically, we use nearest neighbor matching along size, market-to-book, and CSR to identify nine matched control firms for each firm that experienced a data breach.[16] Table 11 presents the characteristics of treatment and matched firms. While we improve our match fitness in several characteristics such as profitability, leverage, and governance, treatment firms are still larger and engage in more CSR. Panel B demonstrates that CSR does not significantly differ between treatment and control groups when we compare treated firms to controls firms simultaneously along all dimensions.

We rerun our main regressions using a stacked difference-in-difference specification with fixed effects for each data breach. Table 12 reports these results. $Post\,1 \times Treated$ and $Post\,4 \times Treated$ capture the difference-in-difference estimates for the firms that suffered a data breach relative to the matched firms. We find similar results in this sample of firms despite the sample size being substantially reduced. Market-to-book decreases by more than 2% in the two years following the breach disclosure, although the result is not

[16] We employ a nearest-neighbor approach rather than a propensity score matching approach because data breaches are very rare events, making propensity score matching extremely noisy (indeed, the fit of our prediction model in Table 2 is quite low.

statistically significant due to the more limiting match procedure. Profitability decreases by 3.8 percentage points over the next four years, while P/E ratios decline by -1.7 (though this result is not statistically significant). Furthermore, the increases in CSR among treatment firms is strong and robust, on the order of .2 standard deviations. The results of these matched tests ensure that our main results are not driven primarily by sample selection issues.

One final potential concern with our results is that the valuation metrics reported in Table 12 are generally not statistically significant, though their economic significance remains large. Hence, it is natural to ask whether managers would really undertake costly CSR investments if long-run firm valuations do not statistically decline following data breaches. However, a manager would likely be judged relative to other firms in his/her industry (as in our main results) rather than to a specific peer (as with our matched results). Also, the possibility of a spillover is larger for the matched firm – that is, we cannot rule out that there was actually a statistically significant effect on value, but that the matched firm's value also declined. Collectively, these points help to alleviate our concerns regarding the lack of statistical significance for the valuation metrics reported in Table 12.

5 Conclusions

We exploit unexpected corporate data breaches to study how firms respond to negative reputation events. By focusing on data breaches we are able to study a negative shock to a firm's reputation that is generally not correlated with underlying product quality or direct costs to the firm. We find that the disclosure of unexpected data breaches are followed by negative stock returns in the following month. However, this stock market reaction is weakened for socially responsible firms, consistent with CSR acting as insurance against

negative reputation shocks.

The effects of these reputation shocks are not limited to short-term price responses but persist in the longer term. We find that the disclosure of data breaches negative impact firm value for at least a few years. Market-to-book, profitability, and price-to-earnings ratios decrease in the years following the data breach. At the same time, we find that firms invest in rebuilding their corporate reputation. Firms significantly increase their investment in CSR by an average of 0.4-0.5 standard deviations in the years following an unexpected breach.

Our paper represents the first empirical study to directly link CSR to firm's investment in corporate reputation. Socially responsible investment is often thought to act as insurance against negative reputation shocks. However, prior research has not shown that the need to improve corporate reputations causes increases in CSR. We present the first evidence that firms actively invest in CSR as the result of a negative reputation shock.

References

Acquisti, Alessandro, Allan Friedman, and Rahul Telang (2006), "Is There a Cost to Privacy Breaches? An Event Study." In *Proceedings of the Twenty-Seventh International Conference on Information Systems*.

Albuquerque, Rui, Art Durnev, and Yrjo Koskinen (2013), "Corporate Social Responsibility and Firm Risk: Theory and Empirical Evidence." Working Paper, Boston College.

Armour, John, Colin Mayer, and Andrea Polo (2017), "Regulatory Sanctions and Reputational Damage in Financial Markets." *Journal of Financial and Quantitative Analysis*, 52, 1429–1448.

Barrage, Lint, Eric Chyn, and Justine Hastings (2014), "Advertising, Reputation, and Environmental Stewardship: Evidence from the BP Oil Spill." NBER Working Paper No. 19838.

Barrot, Jean-Noël and Julien Sauvagnat (2016), "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks." *Quarterly Journal of Economics*, 131, 1543–1592.

Bebchuk, Lucian A., Alma Cohen, and Allen Ferrell (2009), "What Matters in Corporate Governance?" *Review of Financial Studies*, 22, 783–827.

Bénabou, Roland and Jean Tirole (2010), "Individual and Corporate Social Responsibility." *Economica*, 77, 1–19.

Bouslah, Kais, Lawrence Kryzanowski, and Bouchra M'Zali (2013), "The Impact of the Dimensions of Social Performance on Firm Risk." *Journal of Banking & Finance*, 37, 1258 – 1273.

Campbell, Katherine, Lawrence A Gordon, Martin P Loeb, and Lei Zhou (2003), "The Economic Cost of Publicly Announced Information Security Breaches: Empirical Evidence from the Stock Market." *Journal of Computer Security*, 11, 431–448.

Chakravarthy, Jivas, Ed DeHaan, and Shivaram Rajgopal (2014), "Reputation repair after a serious restatement." *The Accounting Review*, 89, 1329–1363.

Cheng, Ing-Haw, Harrison Hong, and Kelly Shue (2013), "Do managers do good with other people's money?" Technical report, National Bureau of Economic Research.

Dyck, IJ, Karl Lins, Lukas Roth, and Hannes Wagner (Forthcoming), "Do Institutional Investors Drive Corporate Social Responsibility? International Evidence." *Journal of Financial Economics*.

Edmans, Alex (2011), "Does the Stock Market Fully Value Intangibles? Employee Satisfaction and Equity Prices." *Journal of Financial Economics*, 101, 621–640.

Elfenbein, Daniel W, Ray Fisman, and Brian McManus (2012), "Charity as a Substitute for Reputation: Evidence from an Online Marketplace." *The Review of Economic Studies*, 79, 1441–1468.

Flammer, Caroline (2015), "Does Corporate Social Responsibility Lead to Superior Financial Performance? A Regression Discontinuity Approach." *Management Science*, 61, 2549–2568.

Giuli, Alberta Di and Leonard Kostovetsky (2014), "Are Red or Blue Companies More Likely to go Green? Politics and Corporate Social Responsibility." *Journal of Financial Economics*, 111, 158 – 180.

Godfrey, Paul C., Craig B. Merrill, and Jared M. Hansen (2009), "The Relationship Between Corporate Social Responsibility and Shareholder Value: An Empirical Test of the Risk Management Hypothesis." *Strategic Management Journal*, 30, 425–445.

Gompers, Paul, Joy Ishii, and Andrew Metrick (2003), "Corporate Governance and Equity Prices." *Quarterly Journal of Economics*, 118, 107–155.

Hallegatte, Stephane and Adrien Vogt-Schlib (2016), "Are Losses from Natural Disasters More Than Just Asset Losses?" World Bank Policy Research Working Paper 7885.

Heal, Geoffrey (2005), "Corporate Social Responsibility: An Economic and Financial Framework." *Geneva Papers*, 30, 387–409.

Hong, Harrison and Marcin Kacperczyk (2009), "The Price of Sin: The Effects of Social Norms on Markets." *Journal of Financial Economics*, 93, 15 – 36.

Hong, Harrison and Leonard Kostovetsky (2012), "Red and Blue Investing: Values and Finance." *Journal of Financial Economics*, 103, 1–19.

Hong, Harrison and Inessa Liskovich (2014), "Crime, Punishment and the Halo Effect of Corporate Social Responsibility." Working paper, Princeton University.

Jäger, Simon (2016), "How Substitutable Are Workers? Evidence from Worker Deaths." Working Paper, MIT.

Jarrell, Gregg and Sam Peltzman (1985), "The Impact of Product Recalls on the Wealth of Sellers." *Journal of Political Economy*, 93, 512–536.

Kamiya, Shinichi, Jun-Koo Kang, Jungmin Kim, Andreas Milidonis, and Rene M. Stulz (2018), "What is the Impact of Successful Cyberattacks on Target Firms?" NBER Working Paper No. 24409.

Karpoff, Jonathan M., D. Scott Lee, and Gerald S. Martin (2008), "The Cost to Firms of Cooking the Books." *Journal of Financial and Quantitative Analysis*, 43, 581–611.

Karpoff, Jonathan M, John R Lott Jr, and Eric W Wehrly (2005), "The Reputational Penalties for Environmental Violations: Empirical Evidence." *Journal of Law and Economics*, 48, 653–675.

Kitzmueller, Markus and Jay Shimshack (2012), "Economic Perspectives on Corporate Social Responsibility." *Journal of Economic Literature*, 50, 51–84.

Kreps, David M. (1990), "Corporate Culture and Economic Theory." In *Perspectives on Positive Political Economy*, 90–142, Cambridge University Press, Cambridge.

Lins, Karl V., Henri Servaes, and Ane Tamayo (2017), "Social Capital, Trust, and Firm Performance: The Value of Corporate Social Responsibility during the Financial Crisis." *The Journal of Finance*, 72, 1785–1824.

Liu, Yan and Venkatesh Shankar (2015), "The Dynamic Impact of Product-Harm Crises on Brand Preference and Advertising Effectiveness: An Empirical Analysis of the Automobile Industry." *Management Science*, 61, 2514–2535.

Makridis, Christos Andreas and Benjamin Dean (2017), "The Economic Effects of Cyber Security Failures on Firms: Evidence from Publicly Reported Data Breaches." Working Paper, Carnegie Mellon University.

Margolis, Joshua D., Hillary A. Elfeinbein, and James P. Walsh (2009), "Does it Pay to Be Good...And Does it Matter? A Meta-Analysis of the Relationship between Corporate Social and Financial Performance." Working Paper, Harvard University.

Murphy, Deborah L and Ronald E Shrieves (2009), "Determinants of the Stock Price Reaction to Allegations of Corporate Misconduct: Earnings, Risk, and Firm Size Effects." *Journal of Financial and Quantitative Analysis*, 43.3, 851–612.

Servaes, Henri and Ane Tamayo (2013), "The Impact of Corporate Social Responsibility on Firm Value: The Role of Customer Awareness." *Management Science*, 59, 1045–1061.

Spanos, Georgios and Lefteris Angelis (2016), "The Impact of Information Security Events to the Stock Market: A Systematic Literature Review." *Computers & Security*, 58, 216–229.

Tadelis, Steven (1999), "What's in a Name? Reputation as a Tradeable Asset." *American Economic Review*, 89, 548–563.

Vanhamme, Joëlle and Bas Grobben (2009), "'Too Good to be True!' The Effectiveness of CSR History in Countering Negative Publicity." *Journal of Business Ethics*, 85, 273–283.

Figure 1: **Frequency of Breaches**

(a) **Frequency of Breaches**

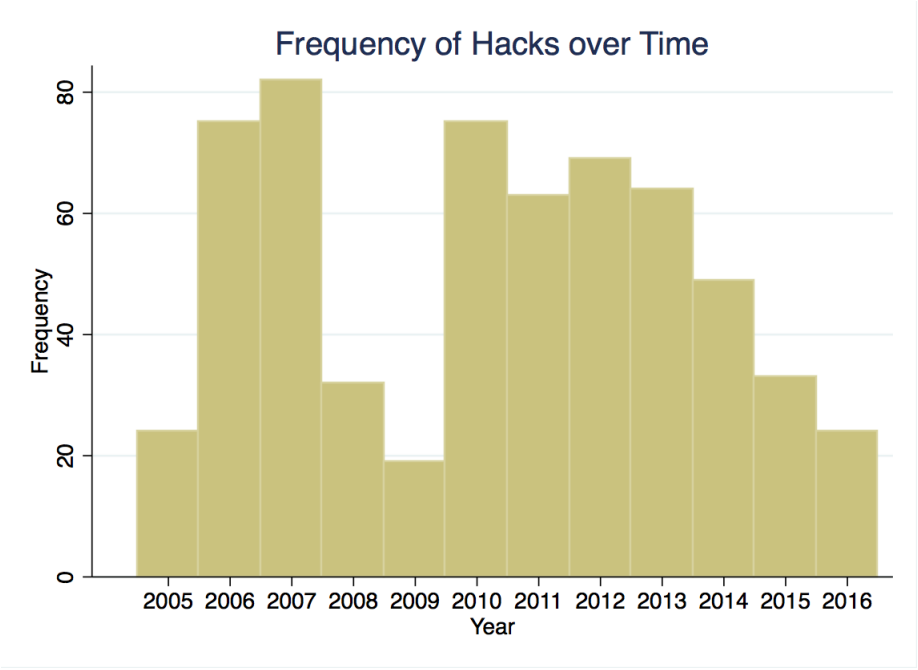

(b) **Industries Affected by Breaches**

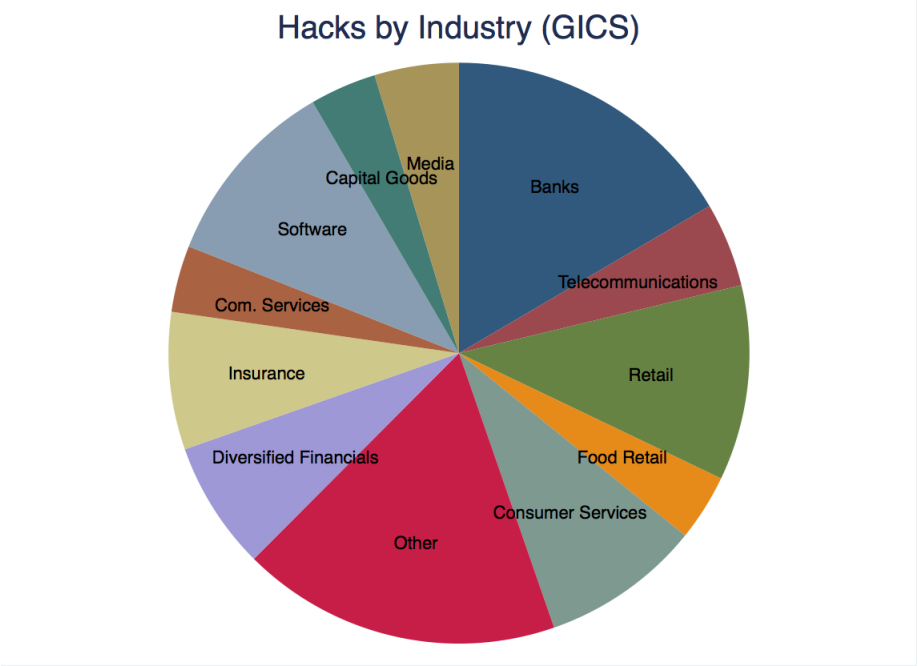

These figures show the distribution of data breaches across time and industry. **Panel (a)** presents the number of breaches per year and **Panel (b)** presents the proportion of breaches across four-digit Global Industry Classification Standard (GICS) industries.

Figure 2: **Number of Records Affected by Breaches**

(a) **Number of Records**

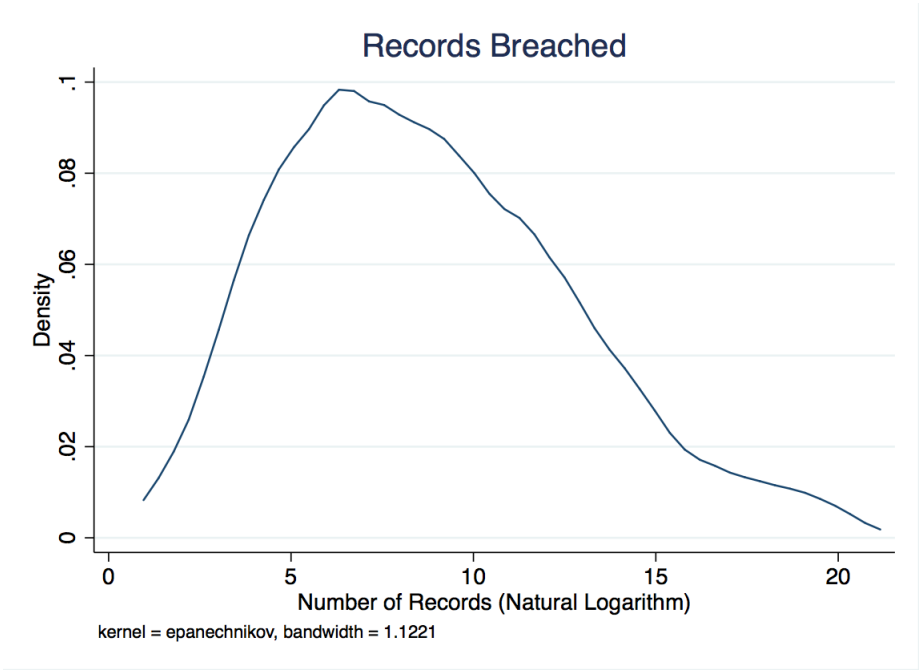

(b) **Number of Records over Time**

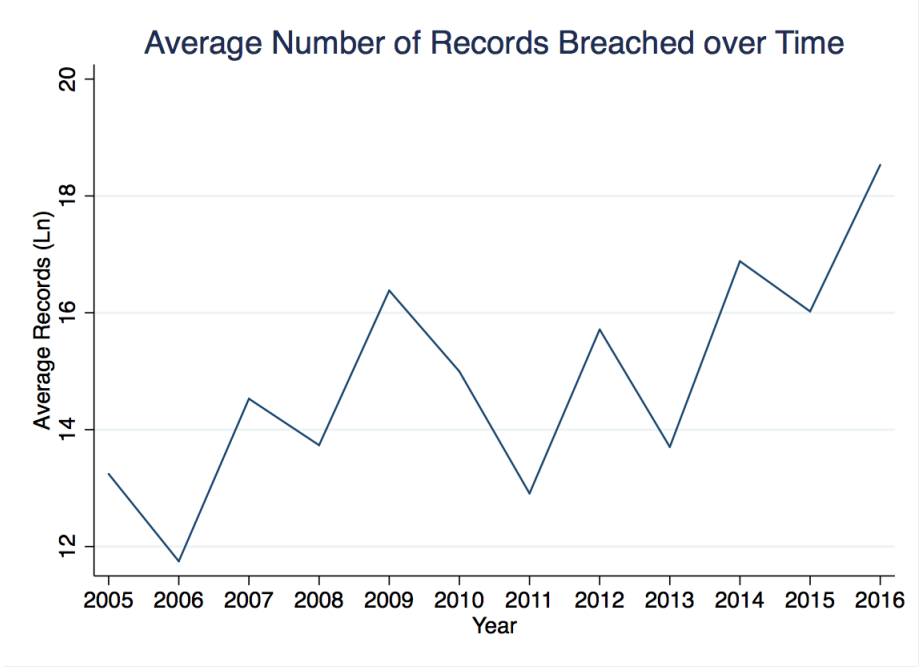

These figures show the distribution of the number of records affected by security breaches. **Panel (a)**presents the a kernal density plot of the natural log of the number of records breached. **Panel (b)** presents the natural log of the average number of records breached across different years.

Figure 3: **Types of Breaches**

(a) **Number of Records**

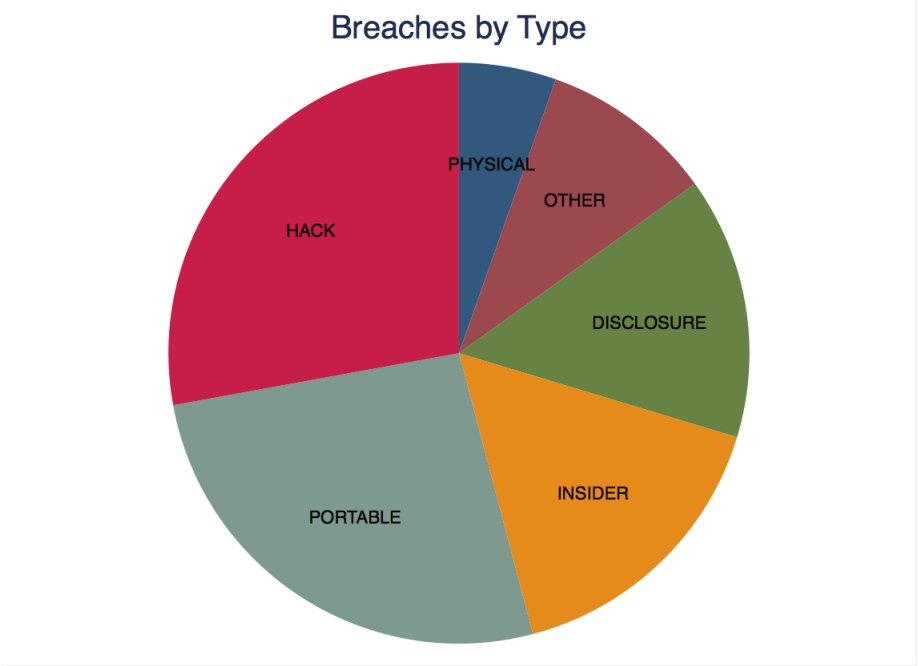

(b) **Number of Records over Time**

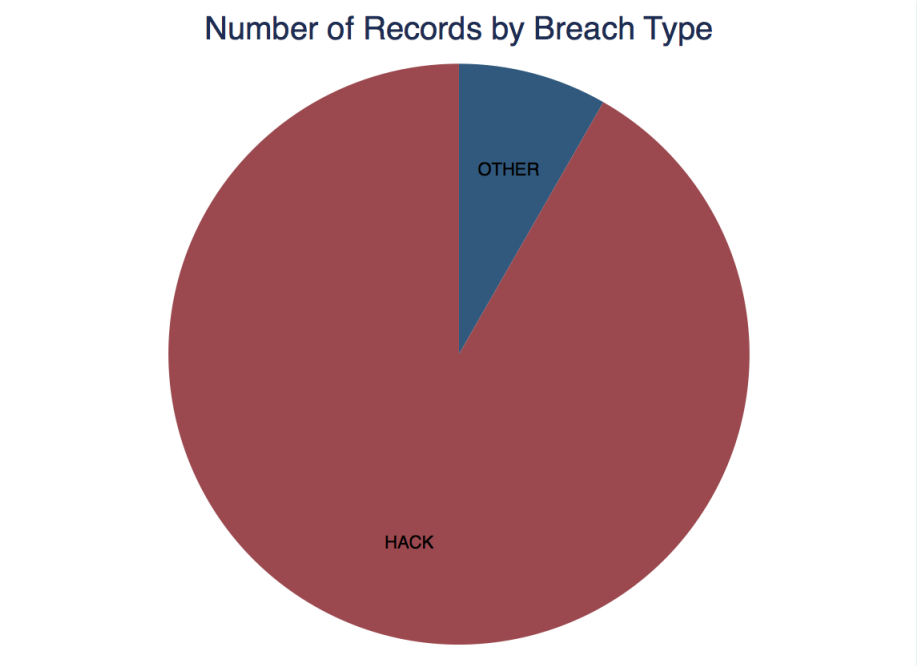

These figures show the distribution of the types of security breaches. **Panel (a)** presents the frequency of different types of breaches. **Panel (b)** presents the proportion of records compromised by different types of breaches.

Figure 4: **Cumulative Abnormal Returns around Breach Announcements**

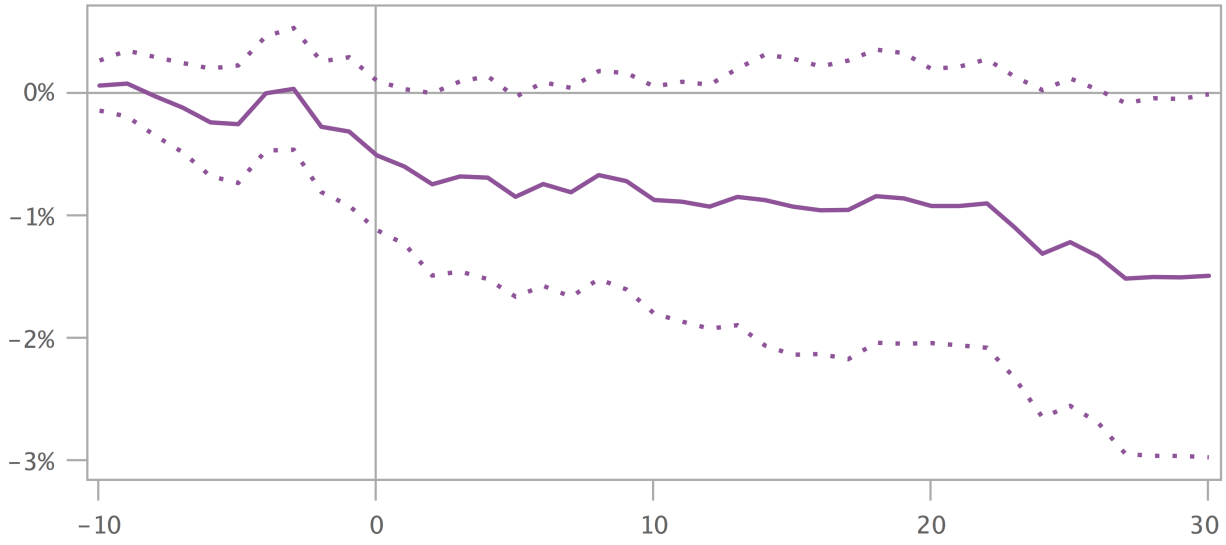

This figure plots the cumulative abnormal returns (CARs) starting ten days before the disclosure of a data breach for which the number of affected records is known. CARs are computed using the Fama-French three factor model

Table 1: Summary Statistics

Panel A: Data Breaches

	Mean	Median	StDev
Total Records	6,173,208	3,482	40,340,917
Employee Records	.331	0	.471
Customer Records	.655	1	.476
Internal Documents	.0139	0	.117
Subsidiary	.0906	0	.288
Multiple Firms	.0523	0	.223
Observations	287		

Panel B: Firms

	Mean	Median	StDev
Norm CSR	.818	.615	1.88
Assets	14,025	21,295	8,297
Leverage	.319	.261	.256
ROA	.0334	.0299	.0904
ROE	.0992	.115	.517
PE	17.1	14.7	19.2
MtB	2.74	2.11	2.39
Q	1.76	1.4	1.15
Sales/Turnover	7,404	10,546	4,229
Nonrecurring	.374	0	.486
Observations	147		

Notes: Panel A covers all data breaches that have been matched to public firms and have a non-missing value for total number of records compromised. The remaining variables are non mutually exclusive indicators for whether the affected records included employee data, customer data, or internal documents. The last two variables are indicators for whether the compromised entity was a subsidiary and whether the data breach affected multiple firms. Panel B includes all public firms that have been matched to data breaches in which at least 1,000 records were compromised. Firm characteristics are measured in the year of the data breach disclosure. The CSR score is normalized such that within the full COMPUSTAT sample the mean is 0 and the standard deviation is 1. All other variables have been winsorized at the 5% level within the full COMPUSTAT sample.

Table 2: Determinants of Data Breaches

	(1) Large Breach	(2) Large Breach	(3) Large Breach
Norm CSR		.00213	.00273
		(.00151)	(.00181)
Entrenchment Index			.0000211
			(.00221)
Constructed Sub G-Index			-.000751
			(.00141)
ln(Assets)	-.00253***	-.0226**	-.0567***
	(.000551)	(.00891)	(.0196)
ln(Assets)^2	.000377***	.00187***	.0035***
	(.0000771)	(.000636)	(.00119)
Leverage	-.00134	-.0088	-.00363
	(.00102)	(.00588)	(.0108)
ln(MtB)	-.000702***	-.00312***	-.00668***
	(.000181)	(.00103)	(.0022)
ROA	.000303	.00254	.00721
	(.000256)	(.00239)	(.00702)
ln(Sales)	-.000268	.000533	.00686*
	(.000175)	(.00104)	(.00384)
Yr x GIC FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Observations	72947	22023	13104
R^2	0.128	0.161	0.192

Notes: The outcome variable is an indicator for whether the firm will suffer a data breach in which at least 1,000 records are compromised. All firm characteristics are measured as of the year prior to the data breach disclosure and are winsorized at the 5% level. Standard errors are clustered at the firm level and reported in parentheses.
* p<.10 ** p<.05 ***p<.01

Table 3: Stock Market Reactions to Data Breaches

	Panel A: Known # of Records Affected			
	(1)	(2)	(3)	(4)
	CAR [-1,3]	CAR [-1,5]	CAR [-1,10]	CAR [-1,30]
Constant	-.0046*	-.00637**	-.007*	-.015**
	(.00276)	(.00297)	(.00371)	(.00673)
Observations	299	299	299	299
R^2	0.000	0.000	0.000	0.000
	Panel B: $\geq 1,000$ Records Affected			
	(1)	(2)	(3)	(4)
	CAR [-1,3]	CAR [-1,5]	CAR [-1,10]	CAR [-1,30]
Constant	-.00764**	-.00897**	-.009*	-.0192**
	(.00387)	(.00415)	(.00487)	(.00842)
Observations	196	196	196	196
R^2	0.000	0.000	0.000	0.000

Notes: CARs are measured relative to the Fama-French three factor model. Expected returns are estimated using a 100-day window and there is a 50-day gap between the estimation window and the date of data breach disclosure. The CAR windows are measured relative to the date of disclosure. Standard errors are reported in parentheses.
* p<.10 ** p<.05 ***p<.01

Table 4: Heterogeneity in Stock Market Reactions

	(1) CAR [-1,30]	(2) CAR [-1,30]	(3) CAR [-1,30]	(4) CAR [-1,30]	(5) CAR [-1,30]	(6) CAR [-1,30]
Constant	-.0192**	-.0339*	.0955	-.0895	.023	-.393
	(.00842)	(.0183)	(.0842)	(.237)	(.0967)	(.253)
High CSR t-1		.039	.0543**	.0571**	.0458*	.0357
		(.0237)	(.0254)	(.026)	(.0254)	(.0259)
ln(Assets) t-1				.0411		.0829*
				(.046)		(.0481)
$\ln(\text{Assets})^2$ t-1				-.000952		-.00288
				(.00229)		(.00239)
Market Leverage t-1				-.156**		-.121
				(.0609)		(.0748)
G-Index t-1					.0179	.00569
					(.011)	(.0115)
E-Index t-1					-.0193	.00194
					(.0141)	(.0156)
ln(Total Records)			-.0105**	-.0128***	-.00909**	-.00767*
			(.00445)	(.00452)	(.00444)	(.00458)
Some Data Missing			.0474	.0363	.0607	.0433
			(.0402)	(.0406)	(.0402)	(.0417)
Employee Records			-.0392	-.102	-.0374	-.0989
			(.0741)	(.0741)	(.0693)	(.0702)
Customer Records			-.042	-.0815	-.0429	-.1
			(.073)	(.0719)	(.0685)	(.0686)
More Xs	No	No	Yes	Yes	Yes	Yes
Observations	196	114	110	109	101	100
R^2	0.000	0.024	0.132	0.223	0.128	0.223

Notes: Data breaches are included if the number of affected records is known and it is at least 1,000. CARs are measured relative to the Fama-French three factor model. Expected returns are estimated using a 100-day window and there is a 50-day gap between the estimation window and the date of data breach disclosure. The CAR windows are measured relative to the date of disclosure. Firm characteristics as measured as of the year prior to data breach disclosure. "High CSR" is defined as firm-year observations that have CSR scores above average for that year-by-industry cell. "Additional Characteristics" include indicators for whether internal documents were affected, whether multiple firms were affected, whether the affected entity is a subsidiary, and whether the breach utilized credit card skimming devices. Standard errors are reported in parentheses.
* p<.10 ** p<.05 ***p<.01

Table 5: Value Reactions to Data Breaches

Panel A: Years 0-1

	(1) ln(MtB)	(2) ln(MtB)	(3) ROE	(4) ROE	(5) PE	(6) PE
Years 0-1	-.245***	-.116***	-.0637**	-.0322	-3.38**	-3.13**
	(.0564)	(.0401)	(.028)	(.0251)	(1.34)	(1.32)
Treated	.215***		.00289		2.19**	
	(.0509)		(.0173)		(1.02)	
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Yr x GIC FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	No	Yes	No	Yes	No	Yes
Observations	74580	73141	84127	82900	84119	82891
R^2	0.303	0.698	0.075	0.330	0.134	0.352

Panel B: Years 0-4

	(1) ln(MtB)	(2) ln(MtB)	(3) ROE	(4) ROE	(5) PE	(6) PE
Years 0-4	-.179***	-.0653	-.0659***	-.0353*	-3.5***	-2.62**
	(.054)	(.0425)	(.0204)	(.0184)	(1.17)	(1.19)
Treated	.233***		.0138		2.77***	
	(.0481)		(.0169)		(1.07)	
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Yr x GIC FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	No	Yes	No	Yes	No	Yes
Observations	74580	73141	84127	82900	84119	82891
R^2	0.303	0.698	0.075	0.330	0.134	0.352

Notes: "Years 0-1 Post" is an indicator for whether a firm has disclosed a data breach in the current or previous year. "Years 0-4 Post" is an indicator for whether a firm has disclosed a data breach within the past five years. Data breaches are included if the number of affected records is known and it is at least 1,000. Firms are only included if there has ever been a data breach in their six-digit GIC industry. Controls include ln(Assets), ln(Assets)^2, and market leverage. Standard errors are clustered at the firm level and reported in parentheses.
* p<.10 ** p<.05 ***p<.01

Table 6: Other Reactions to Data Breaches

Panel A: Years 0-1

	(1) Sales/Assets	(2) Sales/Assets	(3) EBITDA/Equity	(4) EBITDA/Equity	(5) Nonrecurring	(6) Nonrecurring
Years 0-1	.0154	.0186	-.0579**	-.0282	.0676**	.0689**
	(.022)	(.0116)	(.0283)	(.0232)	(.034)	(.0309)
Treated	.0276		-.00925		.0519***	
	(.0375)		(.0256)		(.0178)	
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Year x GIC FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	No	Yes	No	Yes	No	Yes
Observations	84132	82904	81527	80308	84308	83083
R^2	0.437	0.845	0.120	0.422	0.134	0.333

Panel B: Years 0-4

	(1) Sales/Assets	(2) Sales/Assets	(3) EBITDA/Equity	(4) EBITDA/Equity	(5) Nonrecurring	(6) Nonrecurring
Years 0-4	.0156	-.00193	-.0281	-.00513	.0869***	.0864***
	(.0246)	(.0145)	(.0315)	(.0261)	(.0305)	(.0275)
Treated	.0251		-.00954		.035**	
	(.0378)		(.0262)		(.0169)	
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Year x GIC FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	No	Yes	No	Yes	No	Yes
Observations	84132	82904	81527	80308	84308	83083
R^2	0.437	0.845	0.120	0.422	0.134	0.334

Notes: "Years 0-1 Post" is an indicator for whether a firm has disclosed a data breach in the current or previous year. "Years 0-4 Post" is an indicator for whether a firm has disclosed a data breach within the past five years. "Sales/Assets" and "EBITDA/Equity" are winsorized at the 5% level within the full COMPUSTAT sample. "Nonrecurring" is an indicator for whether the firm has a non-zero value for a non-recurring entry for "Nonreccuring Disc Operations" or " Nonrecurring Income Taxes After-tax". Data breaches are included if the number of affected records is known and it is at least 1,000. Firms are only included if there has ever been a data breach in their six-digit GIC industry. Controls include ln(Assets), $\ln(\text{Assets})^2$, and market leverage. Standard errors are clustered at the firm level and reported in parentheses.
* p<.10 ** p<.05 ***p<.01

Table 7: CSR Reaction to Data Breaches

	(1) Norm CSR	(2) Norm CSR	(3) Norm CSR	(4) Norm CSR	(5) Norm CSR	(6) Norm CSR	(7) Norm CSR	(8) Norm CSR	(9) Norm CSR	(10) Norm CSR
Years 0-1 Post	.318**	.235*	.22	.181	.143					
	(.15)	(.138)	(.136)	(.137)	(.118)					
Years 0-4 Post						.515***	.475***	.45***	.424***	.386***
						(.149)	(.146)	(.144)	(.145)	(.129)
Treated	.488***	.191	.13	.14		.364***	.0693	.0141	.0256	
	(.114)	(.117)	(.11)	(.112)		(.117)	(.12)	(.114)	(.116)	
Controls	No	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	No	No	Yes	Yes	Yes	No	No
GIC FE	No	No	Yes	No	No	No	No	Yes	No	No
Yr x GIC FE	No	No	No	Yes	Yes	No	No	No	Yes	Yes
Firm FE	No	No	No	No	Yes	No	No	No	No	Yes
Observations	23278	23154	23154	23138	22739	23278	23154	23154	23138	22739
R^2	0.041	0.080	0.120	0.168	0.605	0.044	0.083	0.123	0.170	0.607

Notes: "Years 0-1 Post" is an indicator for whether a firm has disclosed a data breach in the current or previous year. "Years 0-4 Post" is an indicator for whether a firm has disclosed a data breach within the past five years. The CSR score is normalized such that within the full COMPUSTAT sample the mean is 0 and the standard deviation is 1. Data breaches are included if the number of affected records is known and it is at least 1,000. Firms are only included if there has ever been a data breach in their six-digit GIC industry. Controls include ln(Assets), ln(Assets)^2, and market leverage. Standard errors are clustered at the firm level and reported in parentheses.
* p<.10 ** p<.05 ***p<.01

Table 8: Heterogeneity Across Components of CSR

	(1) Environment	(2) Employee	(3) Community	(4) Diversity	(5) Product
Years 0-4	.601***	-.155	-.111	.335***	-.0327
	(.11)	(.107)	(.0988)	(.12)	(.125)
Controls	Yes	Yes	Yes	Yes	Yes
Year x GIC FE	Yes	Yes	Yes	Yes	Yes
Firm	Yes	Yes	Yes	Yes	Yes
Observations	22739	22739	22739	22739	22739
R^2	0.596	0.433	0.550	0.662	0.581

Notes: "Years 0-4 Post" is an indicator for whether a firm has disclosed a data breach within the past five years. Each column uses as its outcome variable a different component of the overall CSR score. Each component sums strengths and subtracts concerns within a certain area of social responsibility. All measures are normalized such that within the full COMPUSTAT sample their means are 0 and their standard deviations are 1. Data breaches are included if the number of affected records is known and it is at least 1,000. Firms are only included if there has ever been a data breach in their six-digit GIC industry. Controls include ln(Assets), ln(Assets)^2, and market leverage. Standard errors are clustered at the firm level and reported in parentheses.
* p<.10 ** p<.05 ***p<.01

Table 9: Robustness to Governance & Managerial Reaction

	(1) Norm CSR	(2) Norm CSR	(3) Norm CSR	(4) Norm CSR	(5) CEO Leaves	(6) CEO Leaves	(7) CEO Leaves	(8) CEO Leaves
Years 0-1 Post	.197	.131			-.0858***	-.0926***		
	(.15)	(.128)			(.0281)	(.0277)		
Years 0-4 Post			.44***	.353**			-.121***	-.127***
			(.16)	(.144)			(.0315)	(.0314)
Treated	.116		.000781		.0684**		.0945***	
	(.121)		(.126)		(.0275)		(.0335)	
E-Index	.00171	-.0229	.00127	-.0218				
	(.0257)	(.0276)	(.0256)	(.0275)				
G-Index	.0121	.00336	.0124	.0035				
	(.0167)	(.0183)	(.0167)	(.0182)				
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Yr x GIC FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	13573	13321	13573	13321	21402	21242	21402	21242
R^2	0.217	0.633	0.219	0.635	0.192	0.575	0.193	0.576

Notes: "Years 0-1 Post" is an indicator for whether a firm has disclosed a data breach in the current or previous year. "Years 0-4 Post" is an indicator for whether a firm has disclosed a data breach within the past five years. The CSR score is normalized such that within the full COMPUSTAT sample the mean is 0 and the standard deviation is 1. CEO Leaves is an indicator for whether the CEO exits the firm in any given year, taken from Execucomp. Data breaches are included if the number of affected records is known and it is at least 1,000. Firms are only included if there has ever been a data breach in their six-digit GIC industry. Controls include ln(Assets), ln(Assets)^2, and market leverage. Standard errors are clustered at the firm level and reported in parentheses.
* p<.10 ** p<.05 ***p<.01

Table 10: Heterogeneity in Reactions

	(1) ln(MtB)	(2) ln(MtB)	(3) ROE	(4) ROE	(5) PE	(6) PE	(7) Norm CSR	(8) Norm CSR
Years 0-1 Post x Customer Records	-.155***		-.0445*		-3.04*			
	(.0474)		(.0254)		(1.72)			
Years 0-1 Post x Employee Records		-.0856		-.0464		-2.58*		
		(.0593)		(.0446)		(1.45)		
Years 0-4 Post x Customer Records							.241*	
							(.14)	
Years 0-4 Post x Employee Records								.534**
								(.215)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year x GIC FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	73141	73141	82900	82900	82891	82891	22739	22739
R^2	0.698	0.698	0.330	0.330	0.352	0.352	0.605	0.607

Notes: "Years 0-1 Post" is an indicator for whether a firm has disclosed a data breach in the current or previous year. "Years 0-4 Post" is an indicator for whether a firm has disclosed a data breach within the past five years. "Employee Records" indicates that employee records were involved in the breach. "Customer Records" indicates that customer records were involved. These two designations are not mutually exclusive. The CSR score is normalized such that within the full COMPUSTAT sample the mean is 0 and the standard deviation is 1. Data breaches are included if the number of affected records is known and it is at least 1,000. Firms are only included if there has ever been a data breach in their six-digit GIC industry. Controls include ln(Assets), ln(Assets)^2, and market leverage. Standard errors are clustered at the firm level and reported in parentheses.
* p<.10 ** p<.05 ***p<.01

Table 11: Nearest Neighbor Matching

Panel A: Univariate Comparison

	(1) Control	(2) Treatment	(3) (1) vs. (2), p-value
Norm CSR	0.206	0.591	0.001
E-Index	2.977	2.608	0.011
G-Index	6.556	6.412	0.453
ln(Assets)	8.177	9.763	0.000
Leverage	0.251	0.270	0.401
ROA	0.033	0.054	0.117
ln(MtB)	0.869	0.938	0.295
ln(Sales)	7.418	8.936	0.000
N	1007	112	

Panel B: Joint Comparison

	Treated	
Norm CSR	.0116	(.178)
E-Index	-.0234**	(.0152)
G-Index	.00868	(.252)
ln(Assets)	-.107**	(.0483)
ln(Assets)^2	.00988***	(.000375)
Leverage	-.0639	(.385)
ln(MtB)	.0232	(.156)
ROA	-.198**	(.0448)
ln(Sales)	.0412*	(.0736)
Event FE	Yes	
Observations	5566	
R^2	0.197	

Notes: All firm-years with data breaches have been matched to firms in the same year and GIC industry. Control firms are chosen as the nine observations that are closest in three characteristics: ln(Assets), CSR, and market-to-book. Distance is calculated as the sum of absolute normalized differences. The firm data includes all public firms that have been matched to concurrent data breaches in which at least 1,000 records were compromised. KLD is normalized such that within the full Compustat sample the mean is 0 and the standard deviation is 1. All other variables have been winsorized at the 5% level within the full Compustat sample. In Panel B, standard errors are clustered at the event level and p-values are reported in parentheses.
* p<.10 ** p<.05 ***p<.01

Table 12: Matched Results

	(1) ln(MtB)	(2) ln(MtB)	(3) ROE	(4) ROE	(5) PE	(6) PE	(7) Norm CSR	(8) Norm CSR
Post 1 x Treated	-.0356		-.046**		-1.65		.174**	
	(.0503)		(.0231)		(1.65)		(.0832)	
Post 4 x Treated		-.00193		-.0393*		-1.42		.238***
		(.0446)		(.0206)		(1.38)		(.0753)
Event FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6192	7995	6375	8246	6265	8111	5687	7332
R^2	0.349	0.326	0.093	0.066	0.075	0.068	0.383	0.333

Notes: All firm-years with data breaches have been matched to firms in the same year and GIC industry. Control firms are chosen as the nine observations that are closest in three characteristics: ln(Assets), normalized CSR, and market-to-book. Distance is calculated as the sum of absolute normalized differences. The firm data includes all public firms that have been matched to concurrent data breaches in which at least 1,000 records were compromised. KLD is normalized such that within the full Compustat sample the mean is 0 and the standard deviation is 1. All other variables have been winsorized at the 5% level within the full Compustat sample. Standard errors are clustered at the event level and reported in parentheses.
* p<.10 ** p<.05 ***p<.01