FIRM POLICY IN A MULTIPARTITE ECONOMY: SOURCES OF STRATEGIC INVESTMENT, RISK DISCOVERY, AND INPUT VALUATION

A Dissertation

by

JAMES CHANDLER NORDLUND

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	Shane Johnson
Committee Members,	Audra Boone
	Christa Bouwman
	Adam Kolasinski
	Senyo Tse
Head of Department,	Sorin Sorescu

August 2017

Major Subject: Finance

ABSTRACT

This dissertation studies how the interconnectedness embedded in a firm's economic environment shapes a variety of corporate policies, including (1) strategic investment, (2) discovery of sources of risk and its dissemination to shareholders, and (3) valuation of labor inputs utilized in production. I frame the firm's economy as a multipartite graph, wherein corporate directors, executive officers, companies, and product spaces represent different classes of nodes. Section 1 discusses competitive interaction of company nodes induced by shared connections to product space nodes. I examine conditionally random shocks to a firm's financial flexibility and measure the strategic investment responses by rival firms sharing a common product space via a causal indirect treatment effect estimator novel to the finance literature. Section 2 then looks to information flow across companies through shared connections to director nodes. I study the permeation of conditionally random cybersecurity events across a director's current board appointments to investigate the role of director human capital in monitoring, identifying, and disclosing sources of firm risk. Section 3 focuses within a company instead of looking across companies, and is coauthored with Shane Johnson, Adam Kolasinski, and Steve Boivie. We investigate the interaction of executive officer nodes within a firm by analyzing the effect of a chief executive officer's narcissism on his or her valuation of employee human capital.

DEDICATION

To my family.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supported by a dissertation committee consisting of Professors Shane Johnson (advisor), Christa Bouwman, and Adam Kolasinski of the Department of Finance, Professor Senyo Tse of the Department of Accounting, and Professor Audra Boone of the Department of Finance at Texas Christian University.

Section 4 is joint work with Professors Shane Johnson, Adam Kolasinski, and Steve Boivie (Department of Management).

All other work conducted for the dissertation was completed by the student independently.

Funding Sources

Graduate study was supported by fellowships from Texas A&M University.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE Page

LIST OF TABLES

1. INTRODUCTION

This dissertation will study how the interconnectedness embedded in a firm's economic environment shapes a variety of corporate polices, including: (1) strategic investment, (2) discovery of sources of risk and its dissemination to shareholders, and (3) valuation of labor inputs utilized in production. In the language of network theory, the economic environment in which a firm does business is said to be multipartite. This means that there are multiple independent sets of vertices in the graph characterizing the connections in the economy. For example, in figure 1.1, we can visualize the product market competition between firm's **A** and **B** by showing that they both operate in industry **P**. This is represented in the graph by showing that both nodes **A** and **B** (both in the set of firm vertices) share edges with node **P** (in the set of industry vertices). Similarly, if director **D** sits on the board at both firm **A** and firm **B**, both nodes **A** and **B** will share edges with node **D** (in the set of director vertices).

Figure 1.1: An Example Multipartite Network
This figure presents an example multipartite network with director, firm, and industry nodes. A product market rivalry between firms A and B is shown via a shared connection to product space (industry) P. Similarly, two firms that share a director would have a common director node; an example of this is director D, who works at both firm and firm B.

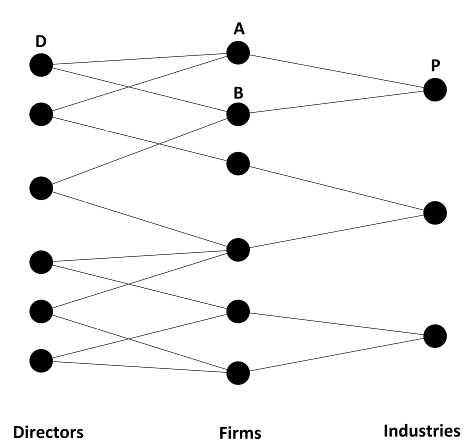

Figure 1.2: Intersection of Boardex and Compustat Databases
This figure presents the entire network of connections that exist within the intersection of the Boardex and Compustat databases. Directors in the database are mapped to the one or more firms at which they are appointed, and each of these firms in turn is assigned to a particular industry space based on the firm's NAICS code.

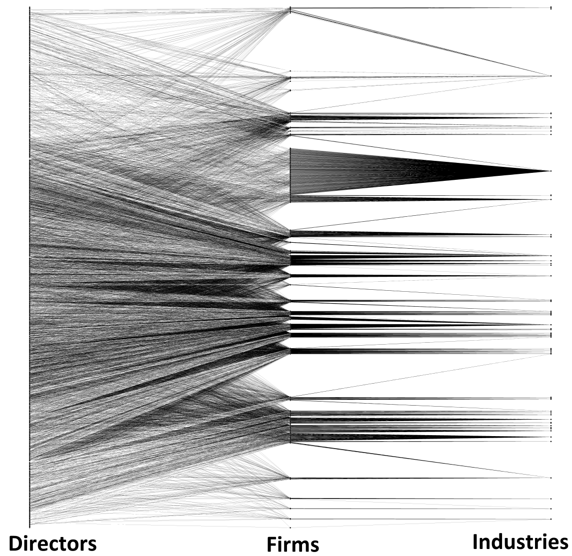

The interconnectedness of firms within an economy is vast. Figure 1.2 shows the shared director and industry connections for firms within the intersection of the Boardex and Compustat databases over calendar year 2012. Here, industries are measured at the 3-digit NAICS level. The data include over 26,000 directors and 3,700 firms. It is clear from the graph that there are a substantial number of director and industry connections for companies in the United States. Economic theory suggests that these inter-firm links lead to behavior different from simplified settings in which firms operate in isolation without, for example, product market rivals and shared human capital. The first two sections of this dissertation will analyze how the network portrayed in figure 1.2 can be used to better understand how certain corporate policies (investment, risk disclosure) are determined. The third section will utilize connections within firms characterized by social interactions of a firm's managerial team.

2. STRATEGIC RESPONSES TO RIVAL FINANCIAL FLEXIBILITY

2.1 Introduction

Out of a number of factors that can affect a firm's capital structure, survey evidence tells us CFOs are more likely to be concerned about financial flexibility than they are anything else.[1] This means that flexibility, the ability of a firm to access or restructure its financing, is of first-order importance to managers. However, academics have only recently begun studying how that flexibility matters.[2] Existing work considers the effect of financial flexibility for an isolated firm; in this paper I investigate the implications of financial flexibility in the richer economic context in which firms face competition.

I propose two alternate channels through which changes to financial flexibility might affect peer firms. First, I consider a "Risk Hypothesis" in which shocks to financial flexibility at one firm reveal information about the risk of inflexibility to managers at competing firms; observing such a shock might lead managers of rival firms to become more cautious about their future flexibility. Second, I posit a "Predation Hypothesis" that recognizes inflexibility could mark a point of weakness for a firm that would permit rivals to strategically prey on the firm. My results show that the latter effect dominates: firms raise debt issuance, investment, and sales growth when peers experience periods of inflexibility. This competitive cost of inflexibility constitutes a previously undocumented benefit of financial flexibility.

I use a double randomization identification strategy novel to the finance literature to study how lost financial flexibility following covenant violations affects rivals who do not violate covenants. Prior research argues such violations are conditionally random events.[3]

[1] Graham and Harvey (2001) survey 392 CFOs and find that, from a list of 14 potential determinants of debt policy, more CFOs classify financial flexibility as important than they do any other factor.

[2] See Denis (2011) for a survey of the literature.

[3] See, for example, Billett, Esmer, and Yu (2016), Chava and Roberts (2008), Demerjian and Owens

The double-randomization framework is necessitated by the fact that the causal estimators familiar to the finance literature (regression discontinuity, difference-in-differences) require the Stable Unit Treatment Value Assumption (SUTVA). This assumption rules out the possibility of indirect treatment effects to non-violators by assuming zero spillover effects.

The hypothesis that a change in flexibility at firm A might induce a causal effect on firm B's outcomes is a direct violation of the SUTVA that requires the treatment effect at firm A to not change the expected outcome for firm B. This is a fundamental assumption in Rubin's (1974) potential outcome framework; the power of the SUTVA is that we can use firm B to predict what firm A might look like in a counterfactual world. This thought experiment is invalidated though if firm B is indirectly treated via spillover effects from firm A. The SUTVA lies at the heart of many matching methods (propensity score matching, differences in differences, regression discontinuity), which makes these methods unsuitable when spillovers are possible.[4] Additionally, I show that if one were to try to directly extend these SUTVA based estimators to test for spillovers (effects at non-violating firms due to rival firm covenant violation), the results are opposite in sign to what a well-identified double randomization estimator would produce.

Following recent developments in the statistics literature, I invoke an observational approximation to a double-randomized experiment to circumvent SUTVA and uncover how treatment affects rival outcomes. These methods are developed in Hong and Raudenbush (2006) and extended in Ferracci, Jolivet, and Berg (2014). To build intuition, consider the goal of evaluating a new vaccination. It is well understood that if part of a community receives a flu shot, the rest of the community benefits via the externality of inoculation

(2014), Ersahin, Irani, and Le (2015), Ferreira, Ferreira, and Mariano (2015), Nini, Smith, and Sufi (2012), Roberts and Sufi (2009), and Zhang (2016) for a partial sampling of papers using this economic context.

[4] I show in the next section that using a SUTVA based estimator when spillovers exist produces an estimated average treatment effect (ATE) that is the sum of the direct and indirect effects. In many cases, asking whether the total (direct plus indirect) effect matters is an economically important question.

(fewer people are communicating the disease).[5] Thus, to test a vaccine, one would want to randomly treat different communities at different intensities. Looking at the infection rates of un-vaccinated individuals across communities with, say, 25%, 50%, and 75% inoculation levels allows researchers to uncover variation in infection rates for un-vaccinated persons. This gets towards the causal indirect effect of vaccination on the un-vaccinated population.

Intuitively, the non-experimental counterpart to double-randomization is to perform double-matching. First, within an industry-group (where I argue spillover effects might occur), I estimate matched expected outcomes between firms in order to define average outcomes as a function of treatment status (covenant violation). Estimation is done group-by-group, so that predicted outcomes are found conditional on being in a group with a particular level of treatment intensity (proportion of firms in the group violating covenants). Second, I take the average outcome of a treated (untreated) firm in a group and match it across groups to wash out observable differences between industries in order to define the average outcome of a treated (untreated) firm as a function of the group's treatment intensity. This generates a causal estimate of indirect treatment effects by showing how average outcomes of untreated firms vary in response to the proportion of peers in the industry that violate covenants.

The selection of the spillover group is chosen in a way to minimize the potential for effects traversing outside of the defined group. Moreover, I additionally show that potential sources of misclassification would lead to outcomes opposite to what I find. As with any matching estimator, I assume that unobservables do not drive difference between covenant violators and non-violators; I follow a well established literature on the effects of covenant violation to control for a large amount of firm-level information. This ensures that covenant violators and non-violators are, at least on a broad dimension of observable

[5] http://www.medscape.com/viewarticle/826481

characteristics, similar in financial health.

From the finance literature, I draw two hypotheses for how financial inflexibility via covenant violation might impact rival firms, leading to indirect treatment effects for non-violators. First, in what I term the "Risk Hypothesis," I argue firms could learn about the risk of entry into covenant violation by observing its ex post effects at firms in violation. This is consistent with the idea in Leary and Roberts (2014) that managers at competing firms learn from each other about optimal capital structure decisions and fits the argument of Roberts and Sufi (2009) that "knowing ex ante that debt contracts impose significant restrictions on corporate behavior and violation of those restrictions impose significant consequences, managers may decide to rely less than they otherwise would on debt financing." Additionally, seeing competitors hamstrung by covenant violation makes the risk of violation especially salient to rival managers. Other economic contexts have shown salient risks to skew firm policy making away from optimal decisions (Dessaint and Matray 2015). This channel would suggest that corporate policies at firms whose peers violate covenants should be more conservative than the policies of firms whose rivals do not violate covenants.

Second, in what I term the "Predation Hypothesis," I argue managers might exploit peer firm covenant violations as an opportunity to steal market share from the violators. This is similar in spirit to theories in which firms without the need to finance externally prey on levered peers through aggressive pricing (Poitevin 1989; Bolton and Scharfstein 1990). It is worth emphasizing, however, that my identification strategy compares violators and non-violators who are ex ante characteristically similar – thus, my results are not driven by firms that are endogenously rich with cash in the pre-treatment period. This hypothesis asserts that corporate policies at firms whose peers violate covenants should be more aggressive than the policies of firms whose rivals do not violate covenants.

My results show that expected changes to net debt issuance, investment, and sales by both treated and untreated firms are functions of treatment intensity. This indicates a rejection of SUTVA and provides causal evidence of strategic reaction to peer-firm changes in financial flexibility that are consistent with the Predation Hypothesis and in contradiction to the Risk Hypothesis. Due to the double-matching estimation, the non-violating firms I identify as strategically responding to peer-firm covenant violation are characteristically similar to the violators. This means that although they were equally likely to violate a covenant themselves, these non-violating firms choose to respond aggressively to peer-firm inflexibility.

The analysis presented here contributes to the literature on financial flexibility. Several recent papers recognize the importance of flexibility in financing (Denis 2011; Denis and McKeon 2012; Gamba and Triantis 2012; Rapp, Schmid, and Urban 2014), with DeAngelo and DeAngelo (2007) arguing "financial flexibility is the critical missing link for an empirically viable [capital structure] theory." The literature on financial flexibility does not consider the decisions of competitor firms.[6] My paper demonstrates a competitive cost to financial inflexibility that is a novel addition to this literature.

My paper also adds to the literature that analyzes the ramifications of debt covenant violation. Chava and Roberts (2008) propose three within firm-lender pair drivers of lost investment stemming from covenant violation (dead-weight losses, creditor intervention, and tightened credit constraints). My paper identifies a fourth, competitive cost to covenant violation. The spillovers in debt issuance, investment, and sales growth that I find give richer economic context to the earlier work of Chava and Roberts (2008), Nini, Smith, and Sufi (2012), and Roberts and Sufi (2009) by showing what competitors do when their rivals announce a new covenant violation. None of these papers disentangle direct effects from potential indirect effects. Consistent with the causal opportunistic responses that I find, a

[6] An exception is Hege and Hennessy (2010), who consider leverage incentives for an incumbent with potential entrants. Choice of both debt level and financial covenants are effected by the incentive to deter entry, since these variables affect the exit outcomes of the entrant.

recent working paper by Billett, Esmer, and Yu (2016) reports that advertising expenses at non-violators tend to be higher when their competitors violate covenants.

Finally, I contribute to very recent efforts in the finance literature to understand how one might estimate effects to indirectly treated units. The strategy here is to utilize a double-matching approach analogous to double-randomization trials commonly exploited in other academic fields. Depending on the economic environment and the nature of how treatment arises, other strategies may be appropriate. For example, both Boehmer, Jones, and Zhang (2015) and Cerqueti, Fiordelisi, and Rau (2015) modify difference-in-difference designs to fit their needs. Many quasi-natural experiments exist within the finance and accounting literature (Atanasov and Black 2016). It is unlikely that many, if any, corporate decisions are made in isolation, so it is useful to have tools that can break down observed causal responses into direct and indirect changes in outcomes because observation of these two separate effects gives richer insight into the specific economic forces that drive differences in outcomes.

Section 2 introduces the novel identification methods employed in this study and section 3 describes the data used. Section 4 documents and discusses the empirical findings of the paper. Section 5 concludes.

2.2 Empirical Methods

This section introduces the identification strategy used for detecting spillovers in treatment effects. The method, developed by Hong and Raudenbush (2006) and extended by Ferracci, Jolivet, and Berg (2014), is essentially a matching estimator within a matching estimator via a two-step procedure.[7] First, on an industry-by-industry basis, it compares outcomes between treated and untreated firms based on their likelihood of treatment.[8] The

[7] I sometimes refer to this estimator as a *hierarchical matching* estimator because matching occurs at both the individual and the group level, although this terminology is not standard in the literature.

[8] For notational convenience, I sometimes refer to firms that do not violate covenants as untreated. This assumes the null hypothesis of no spillover effects.

first step produces two expected outcomes – an expected response for treated firms, and an expected response for untreated firms – for each industry-quarter pair in the data.

Second, aggregating to the industry-quarter level, I model predicted outcomes (again, conditional on treatment) as functions of the intensity of treatment for an industry-quarter. The intensity of treatment is simply the rate of treatment within the group. Because more covenant violations might be expected in some industry-quarter pairs, the second step compares predicted outcomes by matching over a generalized propensity score that allows for the intensity of treatment to fall anywhere within the unit interval.

2.2.1 Peer Effects in Potential Outcome Models

Given a binary treatment, $\tau_i = 1$ if firm i receives the treatment and zero otherwise, and N firms in the population, the $1 \times N$ vector $\boldsymbol{\tau}$ describes the treatment state of the population. The fundamental problem of causal inference is to identify the differences between outcomes under state $\boldsymbol{\tau}$ and outcomes under a different state of treatment, $\boldsymbol{\tau}'$.

Potential outcomes for firm i are modeled as a function of the treatment, so that $y_i(\boldsymbol{\tau})$ describes what the outcomes for firm i would look like under $\boldsymbol{\tau}$. Letting $\boldsymbol{\tau}_{-i}$ denote a $1 \times (N-1)$ vector of treatment assignments with τ_i removed, potential outcomes can be represented by $y_i(\tau_i, \boldsymbol{\tau}_{-i})$. Note that this setup is general in that it allows the outcome of i to be a function of the firm's peers via their treatment: $\boldsymbol{\tau}_{-i}$.

The Stable Unit Treatment Value Assumption presumes a special case of above in which $y_i(\tau_i, \boldsymbol{\tau}_{-i}) = y_i(\tau_i)$.[9] Even with SUTVA in place, causal inference is made difficult by the fact that if one can observe $y_i(\tau_i = 1)$ then the outcome $y_i(\tau_i = 0)$ is unobserved for firm i. Each firm has one observed outcome and thus one unobserved outcome. Differences in differences, regression discontinuity, and propensity score / nearest neighbor matching each invoke SUTVA. In these models, the one unobserved outcome of a treated

[9] To emphasize the cost of this assumption, Manski (2013) re-labels this the "individualistic treatment response" assumption.

firm is approximated by the observed outcome of an untreated firm. In the covenant violation literature, what this means is that a SUTVA estimator will infer what the counterfactual (no covenant violation) outcome for a covenant violator is by using the information of non-violators. The implicit assumption is that these non-violators are not changing their behavior based on the fact that their rival is violating a covenant.

The invocation of SUTVA in a study where spillovers may occur need not invalidate the estimates found in that study. However, one should be aware that in some circumstances the average treatment effect (ATE) will present an inaccurate view of the world.[10] As I show by way of example, the sum of the direct effect of the treatment and the indirect (spillover) effect is the average treatment effect reported by a SUTVA estimator (this is formally presented in Hudgens and Halloran (2008)). This total effect describes the *relative* difference in outcomes between treated and untreated firms. In many instances, whether or not a variable x generates causal effects on differences in outcomes between treated and untreated firms is itself an economically important question. However, the external validity of such an estimate would be limited by the fact that in other contexts the spillover channels might be different, and without explicitly addressing these spillovers it would be difficult to argue how they might change. Additionally, knowledge of indirect and direct effects permits testing of richer economic theories.

Without SUTVA, the function y_i maps a $1 \times N$ vector $\boldsymbol{\tau}$ into $\mathbb{R}$. Because each element τ_j of $\boldsymbol{\tau}$ takes one of two values, there are now $2^N - 1$ unobserved potential outcomes for each firm. This is a problem that Gitelman (2005) refers to as "the 'fundamental problem of causal inference' in the extreme." How can one avoid having to estimate $2^N - 1$ potential outcomes? The key insight of Hong and Raudenbush (2006) is to reduce the dimensionality of the problem by modeling peer-influences via a scalar function so

[10]In situations where the direct and indirect effects are of the same sign, but the latter is larger in magnitude, one could "infer that a treatment is beneficial when in fact it is universally harmful" (Sobel 2006).

that

$$y_i(\tau_i, \boldsymbol{\tau}_{-i}) = y_i(\tau_i, v(\boldsymbol{\tau})). \tag{2.1}$$

Suppose, for instance, that the function $v(\cdot)$ maps into one of two outcomes: $\{H, L\}$. The four potential outcomes are then $y(1, H)$, $y(0, H)$, $y(1, L)$, and $y(0, L)$. The next subsection develops intuition for what form one would expect $v(\cdot)$ to take.

2.2.2 Accounting for Peer Effects

This subsection introduces the functional specification for $v(\boldsymbol{\tau})$. It is important to note that, along with a slightly more detailed econometric approach, spillover effects are more complicated to estimate than direct effects because they require more input by the econometrician as to how and in what direction spillovers should permeate across observations. Different contexts will support different specifications for $v(\boldsymbol{\tau})$. To emphasize the point:

Remark. *Several estimators for spillovers exist, with each identification strategy making different economic assumptions about the nature of how spillovers transmit from one individual to the next.*

For example, Cerulli (2015) specifies a function $v(\boldsymbol{\tau})$ that assumes one-way externalities: the treatment affects outcomes for non-treated units but the non-treated units do not influence outcomes for the treated units. This makes sense in the context of, say, information transfers, but is less likely to be an appropriate assumption when both treated and untreated firms are expected to make decisions that consider their peers' choices.

2.2.2.1 A Cournot Example

To develop intuition for the identification strategy detailed below, consider a simple Cournot oligopoly with N firms. In the base case, these firms are homogeneous in marginal cost so that the linear inverse demand function is given by

$$p = a - bQ = a - b\sum_{i=1}^{N} q_i \tag{2.2}$$

and the profit function for firm i is given by

$$\pi_i = \big[a - b(q_i + Q_{-i})\big]q_i - c_i q_i \tag{2.3}$$

with $Q_{-i} := \sum_{j \neq i} q_i$. In the case of homogeneous cost, $c_i = c\ \forall i$ it is clear the optimal output for any firm is $q_i^* = q^*\ \forall i$.

Suppose an alternative world in which some proportion ρ of firms exogenously are gifted a technology that puts their marginal cost below the cost of their $(1-\rho)N$ peers. Thus $(1-\rho)N$ firms produce at cost c_H and the remaining ρN firms produce at cost c_L. At Nash equilibrium, high-cost and low-cost firms produce, respectively:

$$\begin{aligned} q_H^* &= \frac{a - c_H + \rho N(c_L - c_H)}{b(N+1)} \\ q_L^* &= \frac{a - c_L + (1-\rho)N(c_H - c_L)}{b(N+1)} \end{aligned} \tag{2.4}$$

The treatment effect of the technology compares the change in outcomes between an economy defined by (2.2),(2.3), and (2.4) and an economy defined by (2.2), (2.3), and ($q_i^* = [a - c_H]/[bN + b]\ \forall i$). While economically simple, this one-shot model can build good intuition for how peer effects change firm decisions at equilibrium. In fact, because the cost parameter is the only parameter characterizing a particular firm, one can think of these firms as identical outside of the randomly assigned, treated variable. For matching-based estimators, empiricists are trained to look for un-treated firms that appear otherwise similar to the treated firm of interest.

What happens to outcomes when a random subset of firms have lower costs? Figure 2.1 plots equilibrium firm production (q^*) under three different scenarios. Perhaps not surprisingly, firms with a lower marginal cost than in the counterfactual case of $c_i = c_H \; \forall i$ produce relatively more output. Important to note is that firms with marginal cost structure c_H under both the realized and the counterfactual world do not have the same outcomes in both scenarios *because the outcomes depend on their peers*. Peer-firm dependence reveals itself in the system (2.4) wherein q_H^* is a function of c_L and q_L^* is a function of c_H.

Graphically, one can see both the direct treatment effect and the indirect treatment effect as the two vertical distances labeled on the figure (open circles on the figure denote counterfactual values and closed circles denote production determined by (2.4)). The points above the counterfactual values represent firms that experience reduced marginal costs, and the points below the counterfactual values represent firms who have the same marginal costs as in the counterfactual case. SUTVA-estimators would report the treatment effect as the sum of the direct and indirect effects. It is perfectly correct to take the total vertical distance between outcomes for treated and untreated firms to report that a lower marginal cost positively affects equilibrium production for treated firms. Importantly, however, such an estimate is incapable of telling whether any of the effect is driven by changing outcomes for peers.

2.2.2.2 *Treatment Intensity*

Observe that variation in the number of treated firms in Figure 2.1 affects both the direct treatment effect and the indirect treatment effect. This suggests that the scalar function $v(\cdot)$ introduced in (2.1) should capture variation in the proportion of treatment within the economy. Indeed, the mapping utilized in Hong and Raudenbush (2006) is $v(\boldsymbol{\tau}) = 1$ if the proportion of treatment is high and $v(\boldsymbol{\tau}) = 0$ if the proportion of treatment is low. This produces four possible outcomes: $y_i(1,1)$, $y_i(1,0)$, $y_i(0,1)$, and $y_i(0,0)$; which is

Figure 2.1: Equilibrium Choices Under Different Marginal Cost Schemes

This figure shows equilibrium levels of output by firms competing in a Cournot oligopoly, as described in section 2.2.2.1. The filled-in circles depict outcomes when a subset ($n = 4, 6, 8$) of firms have marginal costs that are lower than the counterfactual world of equal marginal costs across all firms (the open circles). The example is motivated by a hypothetical experiment in which firms are randomly placed along a street, with their address on the street numbering 1 to 12. A social planner grants a cost-reducing technology to all firms with an address above some threshold. Conceptually, this operates akin to a regression discontinuity (RD) framework, with the firm's address acting as the forcing variable. There are two takeaways from the figure. First, firms to the left of the cutoff (those not receiving treatment) have outcomes that differ from the counterfactual of equal costs due to spillovers. Second, the size of spillover effects are positively correlated with the proportion of firms that directly receive treatment. This observation helps suggest a functional form for $v(\cdot)$ in section 2.2.

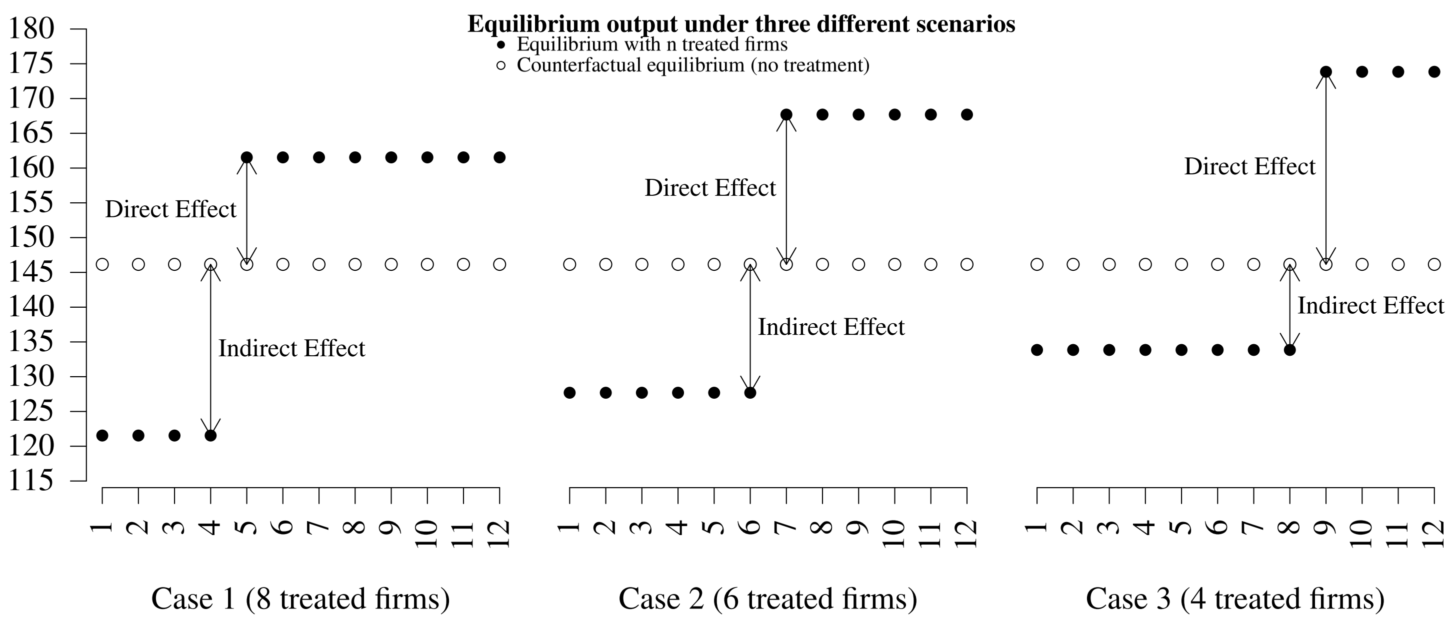

substantially fewer than 2^N in all but the most trivial case.

Ferracci, Jolivet, and Berg (2014) extend the methodology in Hong and Raudenbush (2006) by instead mapping $v(\cdot)$ into $[0,1]$ rather than $\{0,1\}$. They propose taking the proportion of treated individuals, so that $v(\boldsymbol{\tau}) = \bar{\boldsymbol{\tau}}$. In both of these definitions, the implicit assumption is that treatment of any individual j contributes equally to the peer effect via $v(\boldsymbol{\tau})$. When applied to firms, this is less likely to be true. Rather, a weighted average seems more appropriate so that treatment of a large competitor is allowed a different indirect effect from treatment of a very small competitor. Thus in my context I assume

$$v(\boldsymbol{\tau}) = \frac{\sum_{i=1}^{N} \omega_i \tau_i}{\sum_{i=1}^{N} \omega_i} \tag{2.5}$$

with ω_i equal to the pre-treatment level of firm assets.

2.2.3 Identification of Treatment Effects

The unit of observation in this study is firm i in industry j at quarter t. Let $C_i = c$ denote the fact that firm i competes with the set of firms in c. Rather than impose the non-interference assumption on potential outcomes so that $y_i(\boldsymbol{\tau}_t) = y_i(\tau_{i,t})$, I impose the much weaker assumption of *group-level non-interference* and assume that peer effects, if they exist, are isolated within an industry. Thus potential outcomes follow

$$y_i(\boldsymbol{\tau}_t) = y_i(\tau_{i,t}, v(\boldsymbol{\tau}_t)) = y_i(\tau_{i,t}, \rho_{j,t}) \tag{2.6}$$

with

$$\rho_{j,t} = \frac{\sum_{\{k|C_k=C_i\}} \omega_{k,t}\tau_{k,t}}{\sum_{\{k|C_k=C_i\}} \omega_{k,t}} \tag{2.7}$$

so that the weighted average level of treatment within an industry j at time t is allowed to affect outcomes for firms in that industry. In my context, what this means is that I allow firms to change their behavior based not only on whether or not they violate a covenant but also on the proportion of covenant violation within the industry.

There are two dimensions of unobserved outcomes in this model. First, we do not observe $y_i(1 - \tau_{i,t}, \rho_{j,t})$. This means that we do not observe what the outcomes for a covenant violator would have been had it not violated a covenant, nor do we observe what the outcomes for a non-violator would have been had it violated a covenant. Second, we do not observe $y_i(\tau_{i,t}, \rho')$ for $\rho' \neq \rho_{j,t}$. This means that we do not observe what the outcomes for a firm would have been if a different proportion of firms in the industry had violated covenants. Estimation of these two counterfactual outcomes are achieved in two steps. The identification strategy, discussed below, follows that of Ferracci, Jolivet, and Berg (2014).

2.2.3.1 Potential Outcomes Within a Peer Group

Conditional on operating in a given industry j at time t, take $\rho_{j,t}$ to be fixed to some value ρ. Across this set of firms, the only comparison left to make is to measure differences in outcomes $E[y_i(\tau_{i,t} = 1, \rho)|C_i = c]$ and $E[y_i(\tau_{i,t} = 0, \rho)|C_i = c]$.

I assume that there are confounders x_i such that treatment assignments are conditionally random:

$$y_i(\tau, \rho) \perp \tau_i \quad | \quad x_i, C_i = c, \qquad \forall \tau, \rho, c, i. \tag{2.8}$$

The selection of covariates, x_i, to achieve condition (2.8) are described in the next section. This assumption requires that within an industry, for any rate of covenant violation ρ, the assignment of firms to covenant violation status is independent of firm outcomes conditional on x. Along with conditional independence, I require a common support condition, which states that there is no value of x for which the probability of receiving treatment is either 0 or 1. Unlike condition (2.8) which is untestable, common support is easily imposed by simply estimating potential outcomes over the set of observations for which predicted probabilities of treatment are sufficiently within $(0, 1)$. The combination of these two assumptions allows one, for a given group c, to compare treated and untreated outcomes via traditional matching estimators. This is achieved via an augmented inverse-probability-weighted estimator. The final product of step one are the estimated outcomes:

$$\hat{E}[y_i(\tau, \rho)|C_i = c, \rho = \rho_{j,t}] \tag{2.9}$$

for $\tau \in \{0, 1\}$.

Thus, for an industry group, c, with a particular value for ρ, expected outcomes for treated and untreated firms are obtained. The set of firms ultimately used in computing expected outcomes are similar along all observable characteristics (covariate balance). Verifying balance group-by-group, and then computation of an aggregate statistic to test for balance, is left to the appendix. Step two below models how expected responses change over ρ.

2.2.3.2 Potential Outcomes Across Groups

As mentioned earlier, outcomes are not observed for $\rho' \neq \rho_{j,t}$. To estimate these effects, I assume a set of confounders, w at the industry level such that treatment intensity is conditionally independent of expected outcomes in a given industry:

$$E[y_i(\tau, \rho)|C_i = c] \perp \rho_j \quad | \quad w_j, \quad \forall c, \tau, \rho. \tag{2.10}$$

The vector w_j is chosen as the group-level mean of x_i following Hong and Raudenbush (2006). This vector includes both firm-level predictors of covenant violation as well as industry-level measures, such as the competitiveness of the group. The assumption in (2.10) is that the rate of covenant violation in an industry is independent of outcomes for the average firm conditional on a rich et of controls.

Since ρ is a continuous variable over the unit interval, the common support assumption imposed on this step follows Flores et al. (2012). As in step 1, common support is met by restricting estimation over the set of observations for which the condition is most likely to hold. Then, using a generalization of the propensity score (Hirano and Imbens 2004) for continuous treatment levels, one can match expected outcomes in (2.9) over levels of ρ. This yields the final output[11]

$$\hat{E}[y_i(\tau, \rho)|C_i = c]. \tag{2.11}$$

If either $\hat{E}[y_i(0, \rho)|C_i = c]$ or $\hat{E}[y_i(1, \rho)|C_i = c]$ vary over ρ, indirect treatment effects exist. Variation in the latter expected outcome indicates that peer firm responses change based on the level of industry treatment, while variation in the former shows that treated firms change their reaction to treatment based on how many of their rivals find themselves in the same situation. Variation in either case suggests the traditionally utilized individual-level non-interference condition, $y_i(\boldsymbol{\tau}) = y_i(\tau_i)$, is too restrictive.

2.2.4 Two Sources of Endogeneity

There are two separate but equally important endogeneity concerns that necessitate an empirical model that extends beyond ordinary least squares. First, treated firms may be characteristically different from untreated firms. This is a well-understood problem in the empirical finance literature, and I show in section 2.3.2 that treated firms are indeed different than their untreated peers across a number of dimensions. The concern is that the difference in observed outcomes between the two groups may be different not because of the treatment variable, but rather because there are other distinctions between treated and untreated firms that generate variation in responses. Theoretically, these characteristic dif-

[11] Standard errors for the estimated function $\hat{E}(y_i|\tau_i, \rho)$ are calculated via a hierarchical bootstrap procedure. This re-samples clusters – industry-quarters – of data with replacement. Within each drawn cluster, a random subsample of the firms are kept. This is the most efficient way of bootstrapping multilevel data (Ren et al. 2010). The nature of the bootstrap procedure accounts for the possibility of cross-correlated shocks between τ_i and τ_j for $i, j \in c$.

ferences are washed out by the matching performed in section 2.2.3.1. Appendix B shows how to compute a test for covariate balance in which the null hypothesis is that the model sufficiently balances observable characteristics in each of the intra-group regressions, with the alternative hypothesis that at least one variable in at least one of the groups is not balanced between the treated and untreated firms.

The second endogeneity concern in a study of peer effects is that the treatment at the group level – treatment intensity, ρ – is not randomly assigned. Consider two levels of ρ: $\{H, L\}$, with $H > L$. A higher level of treatment intensity H could be indicative of worse industry performance. In this case, differences in outcomes for untreated firms, $E(y|0, H) - E(y|0, L)$ might be driven by observable, characteristic differences between one industry group and another. The matching performed in section 2.2.3.2 controls for differences in the likelihood of a certain treatment intensity to ensure that differences in outcomes over ρ are due do causal effects of treatment, rather than other differences in industry groups.

2.2.5 Selection of a Peer Group

An important trade-off exists between the estimation of (2.9) and (2.11). When peer groups are narrowly defined (for example, a 4-digit NAICS industry classification), the firms defined in a set of competitors, c, are those that most likely compete directly amongst each other. If peer effects exist, it is natural to want to look within this narrow set. Restriction of c to narrow groups, however, ignores the fact that firms can change the scope of their product space. For example, Apple transitioned from a computer company (NAICS 3341) to one that also sold smartphones (NAICS 3342) by moving across 4-digit NAICS groupings. A broader classification of competitive groups reduces the likelihood of omitted competitors. This is important because the existence of spillovers outside of the group c violate the assumption of group-level non-interference.

An additional issue with narrow classifications for c is that, as the set of peer firms shrinks, so too does the number of observations for which we can estimate (2.9). This can lead to sample size issues, especially when incorporating a rich set of confounders x into the model. On the other hand, broad definitions for peer groups will reduce the number of group-quarter observations in the data necessary to estimate (2.11). Thus, after selecting an economic basis for assigning peer groups, it is important to utilize a classification scheme that both satisfies the group-level non-interference assumption as well as leaves ample data to estimate both intra-group and inter-group effects.

When the data allow (i.e. where there are sufficient group-quarter observations), the broadest classification of competitive groups seems the most conservative choice. Doing so minimizes the likelihood that spillover effects carry outside of the specified group c. A broad group definition will possibly include a number of firms that are not in competition with the treated units. This biases estimated indirect effects towards zero, making it harder to reject the null hypothesis of no indirect treatment effect. In this paper, I classify groups based on 2-digit NAICS codes to keep c as broad as possible.[12]

Recent work by Zhang (2016) suggests covenant violation by a firm changes its relationship with its suppliers. Thus, a final consideration in the classification of peer groups is to acknowledge that customers and suppliers may also be affected by a firm's covenant violation. Because a truly complete network of inter-firm relationships would likely have few if any sub-sets of self-contained groupings of firms, a rigorously defined network of all possible upstream and downstream trade is empirically prohibitive. There would be far too few groups with which one could perform inter-group matching to tease out the

[12]One may wonder whether a revealed, product-based categorization of industry grouping as in Hoberg and Phillips (2015) would be an appropriate alternative to NAICS (e.g., FIC codes). The difficulty for FIC codes stems from data loss; firms without a valid product description or without a detailed enough description in their 10-Ks cannot be incorporated in this measure (Hoberg and Phillips 2010). This complicates measurement of ρ and other group-level variables because some potential competitors might not be assigned an FIC code. To the best of my knowledge, the statistical properties of double-randomization under various forms of missing data remains an unanswered, albeit potentially interesting, question.

effects of ρ. Instead, I use the fact that, according to the Bureau of Economic Analysis' Input/Output Make table, 93% of trade between firms stays within a 2-digit NAICS grouping. This implies that the bulk of customer/supplier effects are contained within the groupings chosen for analysis.

Does the potential for customer/supplier effects matter to the results? Certainly, and researchers employing double-matching methods must think carefully about how outcomes might change because of it. In my context, a firm that encounters a higher expected cost of restructuring its debt due to covenant violation invests relatively less than peers that do not (Chava and Roberts 2008). One would expect customers and suppliers to be hit in the same direction. For example, if a mining company slows its rate of growth, a downstream steel manufacturer would likewise slow its growth. Therefore a group that includes customers and suppliers would have expected outcomes between treated and untreated firms that are less far apart, given that some untreated firms (the upstream and downstream firms) would look more like a directly treated company, at least with respect to certain outcomes like investment. My results show that expected outcomes for treated and untreated firms diverge as treatment intensity rises, which implies that these effects are more conservatively estimated than would have been the case if customer/supplier interactions did not exist.

2.3 Data

2.3.1 Covenant Violations

Data on covenant violations are observed at a quarterly level, and were compiled for the analysis in Nini, Smith, and Sufi (2012) by the authors of that paper.[13] The sample analyzes all 10-K and 10-Q filings between 1996 and 2009 for U.S. firms outside of the financial industry and matches them to Compustat records.

Per Nini, Smith, and Sufi (2012), covenant violation waivers frequently expire and

[13]Thanks to Amir Sufi for making the data available on his website.

must be extended. Thus one can observe high serial correlation in violation for a given firm. To isolate the effects driven by truly new changes to a firm's covenant status, I follow the literature and focus on *new violations*. These are covenant violations reported by firms that have not reported any such violation over the last four quarters. This allows me to ensure that shocks identified are first instances in which debtors gain increased control rights.

The emphasis on new violations creates a class of "never-takers" in my sample. Those firms that violated a covenant in period t cannot be classified as treated in periods $t+1$ to $t+4$. Thus, when predicting investment outcomes at the firm level, I am careful to only match newly treated firms to those who were *eligible* for treatment but did not *happen* to violate a covenant.[14]

2.3.2 Firm Fundamentals

All other data come from quarterly Compustat records, with each of the variables used are in the appendix. The outcomes of interest in this paper are net debt issuance, investment, and sales growth. Net debt issuance is defined, as in Roberts and Sufi (2009), as a change in total balance sheet debt. This change is normalized by lagged assets. Investment is given by capital expenditures scaled by prior period PP&E, as in Chava and Roberts (2008). These two variables are measured as changes in the outcome from the quarter following the covenant violation relative to the prior quarter. Sales growth is the sales in the period after the violation less the sales in the quarter of the violation, divided by sales from the quarter of violation.

I use capital expenditures as a measure of investment for consistency with the earlier literature. This allows me to break out indirect and direct treatment effects to help better describe some of the previous results known to occur following debt covenant violation.

[14]For similar reasons, I exclude zero-leverage firms when predicting outcomes.

Moreover, the elasticity of capital expenditure restrictions with respect to covenant violation is higher than that of other loan terms such as interest rates or collateral usage (Nini, Smith, and Sufi 2009), suggesting capital expenditures are a key concern for creditors.[15]

Following the literature on covenant violation, I build a rich set of control variables to predict the likelihood of covenant violation. All variables are measured one period prior to the quarter of treatment. These follow the same predictors of covenant violation used in Nini, Smith, and Sufi (2012). This includes variables on which covenant terms are often written: operating cash flow to average assets, leverage ratio, interest expense to average assets, net worth scaled by total assets, and the current ratio (Roberts and Sufi 2009). I also control for cash and cash equivalents, normalized by total assets, since internal financing is the natural substitute to external financing and can deter predation (Chi and Su 2016). Return on assets likewise captures a firm's ability to raise cash via utilization of its productive assets. I also control for the size of the firm (log of total assets), the proportion of fixed assets (PP&E/total assets) at the firm, and the growth opportunities of the firm (market to book). Market controls w are industry averages of all x variables. Table 2.1 summarizes these variables.

A difference-in-means test between treated and untreated firms shows that these controls can be expected to have explanatory power in determining treatment level. Table 2.2 shows that treated firms have, on average, less cash as a percentage of assets, more debt, and lower growth prospects (as measured by market to book).

2.4 Do Covenant Violations Affect Non-Violating Firms?

Based on the identification outlined in section 2.2, spillovers in treatment should create variation in the expected outcomes over group-level treatment intensity. This section tests

[15] An additional reason to favor capital expenditure policy changes is that many firms report no R&D, and recent evidence suggests omitted R&D is not truly zero R&D but may be strategically omitted. (Koh and Reeb 2015).

Table 2.1: Descriptive Statistics

This table summarizes the variables used in this study. Net debt issuance is the change in balance sheet debt scaled by lagged total assets (Roberts and Sufi 2009). Investment is capital expenditures scaled by lagged PP&E (Chava and Roberts 2008). Sales growth is the change in sales scaled by lagged sales. Debt issuance, investment, and sales changes are measured over the quarter following covenant violation. Control variables are observed in the quarter before an announced covenant violation, and follow the definitions of Nini, Smith, and Sufi (2012). See Appendix A for details.

	Mean	Median	Std. Deviation
Net debt issuance (bps)	55.427	0.000	487.858
Investment (%)	10.797	4.748	358.437
Sales growth (%)	3.926	2.327	20.521
Operating income / average assets (%)	0.368	2.282	8.750
Leverage ratio	0.262	0.191	0.326
Interest expense / average assets (%)	0.608	0.276	1.217
Net worth / assets	0.416	0.500	0.559
Current ratio	2.814	1.929	2.924
Market / book	2.451	1.504	3.040
Cash / asset	0.190	0.088	0.225
ROA (%)	-2.966	0.602	13.143
log(assets)	5.150	5.124	2.236
PP&E / assets	0.262	0.185	0.231
HHI	0.060	0.044	0.053
Observations	147835		

Table 2.2: Differences in Observables Between Violators and Non-violators

This table shows the mean and median of the observables by treatment status. The third column shows the p-value for a t-test of equality in means between the treated and untreated groups. The fourth column reports the p-value for a test with the null hypothesis that there is equality in means on a group-by-group basis. This is done by aggregating group-by-group t-tests of equality via Fisher's method.

	Means				Medians		
	Untreated	Treated	p	Fishers p	Untreated	Treated	p
Operating income / average assets (%)	.358	.944	.001	0	2.291	1.817	0
Leverage ratio	.262	.296	0	0	.19	.259	0
Interest expense / average assets (%)	.607	.665	.021	.098	.273	.456	0
Net worth / assets	.416	.422	.652	.886	.501	.452	0
Current ratio	2.824	2.202	0	0	1.933	1.735	0
Market / book	2.462	1.743	0	0	1.509	1.225	0
Cash / asset	.191	.113	0	0	.089	.047	0
ROA (%)	-2.977	-2.253	.008	.081	.612	.12	0
log(assets)	5.152	5.011	.002	0	5.129	4.899	0
PP&E / assets	.262	.277	.002	0	.185	.209	0
HHI	.06	.062	.122		.044	.044	.004

whether covenant violation shocks to creditor control rights at competitor firms affects outcomes of peers by examining expected outcomes as a function of the proportion of treated firms in the industry group. Unlike the firm-level measure for treatment, which has been used in a number of prior studies, group-level treatment via hierarchical matching is new to this literature. Before utilizing treatment intensity to derive the main results, I first review a few characteristics of the group-level measure.

2.4.1 Analyzing Treatment Intensity

This subsection discusses some of the properties of the group-level treatment measure, ρ. In particular, I detail which subsection of the unit interval has sufficient data for peer effects estimation, the nature of shocks that drive higher treatment intensity, and whether the measure avoids picking up characteristic differences between industry groups.

2.4.1.1 The Range of the Measure

Variation in the measure of treatment intensity given by (2.7) achieves identification in treatment effect spillovers via the identifying assumption in (2.10). Prior papers, for example Nini, Smith, and Sufi (2012), show that covenant violations are not altogether rare events. Figure 2.2 shows the time series of covenant violations at both a quarterly and annual (existence of a new violation at any point in the year) level.

Roughly 6% of firms will report a new covenant violation in any given year (recall that a new violation is defined to be the first indicated violation of a debt covenant over the last four quarters). New covenant violations are noticeably more likely to appear in the fourth quarter of the fiscal year. Sharp spikes in the quarterly data reflect this property. This feature motivates the inclusion of fiscal quarter fixed effects in both x and w (intra-group and inter-group confounders).[16]

[16] Because w is measured at the industry level, I use quarter of year fixed effects in the inter-market matching stage (step 2) to approximately capture fiscal quarter effects at a group level. This is because observations are traced in calendar time, and thus different firms may be at different points in their fiscal year.

Figure 2.2: Time Series of Covenant Violations

This figure shows the time series of covenant violations at an annual and quarterly level. At the annual frequency, a firm is marked as violating a covenant if it experiences a new covenant violation (first report of covenant violation following at least four quarters of non-violation) at any point in its fiscal year.

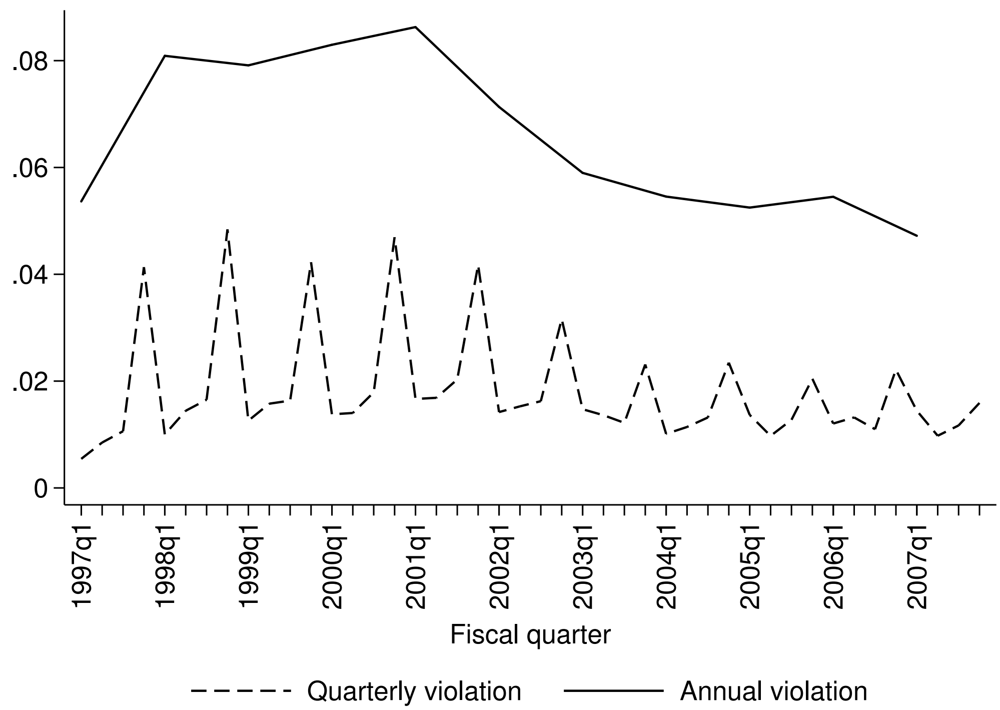

On average, about 1.8% of firms report a new covenant violation in any given quarter. This means that in nearly 2% of firm-quarter observations, firms give up some degree of control to creditors after violating a covenant. The extant literature on debt and competition has focused on the case in which rivals to push a firm into bankruptcy by undercutting the firm and making it unable to repay its debt obligations (Bolton and Scharfstein 1990). In the data, however, distressed exit occurs at a rate of 0.55%. This means creditor control via covenant violation is over three times as likely as the more extreme scenario of distressed exit that usually motivates competitive costs stemming from debt.

A large proportion of industry-quarter pairs experience zero new covenant violation. These groups do not contribute useful information because there is no intra-group variation in treatment. On the other side, while a rare few industry-quarters experience a relatively high degree of treatment intensity, these group-level observations fall far enough into the right tail of the distribution that matching them by generalized propensity score becomes problematic (they are omitted by imposition of the overlap assumption). Thus, the range of group treatment intensity, ρ, over which I can estimate spillover effects is $0\% < \rho < 3.5\%$.

2.4.1.2 Do Shocks Signal More of the Same?

Although conditional independence of treatment (covenant violation) is assumed in (2.8), this does *not* mean that the quasi-random "shock" must be purely idiosyncratic. More than one firm in an industry group might violate a similar debt covenant due to a common industry shock that affects all firms. If these shocks have some persistence, higher intensity of treatment in period t should predict an increased likelihood of treatment in period $t+1$.[17] One reason to think that untreated rivals might mirror their treated peers is that industry-wide events would induce more conservative financing and investment policy across the industry (as in the Risk Hypothesis).

[17]See Hennessy and Strebulaev (2015) for a detailed model of how time-varying treatment probabilities can influence expected outcomes.

Table 2.3: Predicting Covenant Violation

This table reports the logistic estimates predicting covenant violation (at treatment eligible firms) as a function of treatment intensity and firm covariates. Unless otherwise noted, all predictors are measured one period before observing covenant violation. Regressions also include fiscal quarter fixed effects and industry averages of all independent variables. Marginal effects of each variable are reported in place of beta coefficients. The second and fourth columns are similar to the first and third columns, respectively, save for the addition of squared predictors to control for potential non-linear effects.

Treatment intensity	0.034**	0.038**	0.034**	0.038**
	(0.017)	(0.016)	(0.017)	(0.016)
Treatment intensity (second lag)			-0.007	-0.007
			(0.020)	(0.018)
Operating income / average assets	-0.065*	-0.106	-0.065*	-0.106
	(0.034)	(0.104)	(0.034)	(0.104)
Leverage ratio	0.027***	0.034***	0.027***	0.034***
	(0.003)	(0.004)	(0.003)	(0.004)
Interest expense / average assets	0.016	0.237**	0.016	0.237**
	(0.061)	(0.104)	(0.061)	(0.104)
Net worth / assets	0.017***	0.006***	0.017***	0.006***
	(0.002)	(0.002)	(0.002)	(0.002)
Current ratio	-0.003***	-0.004***	-0.003***	-0.004***
	(0.000)	(0.001)	(0.000)	(0.001)
Market / book	-0.004***	-0.004***	-0.004***	-0.004***
	(0.001)	(0.001)	(0.001)	(0.001)
Cash / assets	0.053*	0.033	0.053*	0.033
	(0.029)	(0.096)	(0.029)	(0.096)
ROA	0.008**	0.002	0.008**	0.002
	(0.004)	(0.007)	(0.004)	(0.007)
log(assets)	-0.002***	0.011***	-0.002***	0.011***
	(0.000)	(0.002)	(0.000)	(0.002)
Change in log(at)	0.004	0.008***	0.004	0.008***
	(0.003)	(0.002)	(0.003)	(0.002)
PP&E / assets	-0.006*	0.002	-0.006*	0.002
	(0.004)	(0.006)	(0.004)	(0.006)
Change in PP&E / assets	0.055***	0.049***	0.055***	0.049***
	(0.015)	(0.014)	(0.015)	(0.014)
Squared controls?	No	Yes	No	Yes

Marginal effects reported

Standard errors (clustered on industry) in parentheses

* p <0.10, ** p <0.05, *** p <0.01

Table 2.3 reports the predicted likelihood of a treatment eligible firm (a firm that has not reported a covenant violation in the last four quarters) as a function of prior period treatment intensity, as well as a number of controls. The first column predicts treatment as a function of x and w, while the second includes squared terms of all variables. The marginal effect $\rho_{j,t-1}$ is positive and significant ($p < 0.05$). Changing the treatment intensity from the 25^{th} percentile to the 75^{th} percentile raises the likelihood of covenant violation by 2.6%. This indicates that at least part of the shock that drives firms to violate debt covenants is driven by persistent industry-wide events.[18] To determine whether higher order lags contribute additional information, the last two columns repeat the analysis in the first two, this time adding a second lag in treatment intensity. The effect of the second lagged term is small and indistinguishable from zero.

Because lagged treatment intensity positively predicts the likelihood of covenant violation by firms currently not in violation, I conclude the channel through which the Risk Hypothesis is expected to operate has merit. The optimal response by these firms might be to take on more conservative financing and investment policy. This would decrease the risk of mirroring the fate of their treated peers.

2.4.2 Net Debt Issuance

Do peer effects impact the causal response to treatment? I use the methodology detailed in section 2 to show how debt financing varies over industry treatment intensity. This extends the estimated negative relative treatment effect in Roberts and Sufi (2009) to vary over $\rho_{j,t}$. Figure 2.3 shows the estimated direct and indirect treatment effects. The counterfactual world of no covenant violation is the case $\rho = 0$. For comparison purposes I take the outcomes at the minimal observed treatment intensity, $\rho = \rho_{\text{min}}$, to approximate zero

[18] For this reason, I include lagged industry-level covenant violation in the vector of group-level covariates, w, when predicting firm outcomes. The results are similar when employing a conditional fixed-effects logistic model to control for industry fixed effects.

spillover effects (because effects at $\rho = 0$ are unmeasurable due to overlap conditions); outcomes for $\rho > \rho_{\text{min}}$ include a spillover component. Note that outcomes at $\rho = \rho_{\text{min}}$ (or indeed even $\rho = 0$) are not mean zero because they are evaluated at $x = \bar{x}$; the x vector is suppressed from the expectation expression for notational convenience.

The difference in outcomes between treated and untreated firms imply negative total treatment effects; treated firms reduce net debt issuance substantially whereas untreated firms raise net debt issuance. The increase in net debt issuance by untreated firms grows as the more firms in the industry violate debt covenants, which indicates positive indirect treatment effects. With few firms treated, net debt issuance at untreated firms is roughly -45 basis points, whereas an industry with nearly 3.5% of firms treated yields an expected change of roughly -22 basis points. This response contradicts the Risk Hypothesis that non-violators will become more conservative upon seeing rivals violate covenants. The positive spillover effect is suggestive evidence of the Predation Hypothesis, although the investment and sales growth analysis presented later provides a stronger argument for this hypothesis.

The main interest in this paper is the slope of the expected outcome, not the level.[19] The former gets more specifically to the treatment effect from spillovers since it shows the expected change in outcomes for an untreated firm as a function of the intensity of treatment within the peer group. For maximum flexibility the second step of the matching estimator is done nonparametrically with a kernel function (Flores et al. 2012), however a crude linear approximation to the observed outcome suggests net debt issuances by untreated firms rises by 6.8 basis points for every 1% increase in treatment intensity. This makes a strong case for spillovers across treated to untreated firms (and evidence that SUTVA-based estimators ignore the full richness of treatment effects).

[19]That is, I am interested in $\partial E(y|\tau_i = 0, \rho)/\partial \rho$ because the mean is itself irrelevant. For similar reasons, SUTVA-based estimators do not emphasize $E(y|\tau_i = 1)$ but rather are concerned with the relative outcome: $E(y|\tau_i = 1) - E(y|\tau_i = 0)$. This is similar to $\partial E(y)/\partial \tau$ for discrete τ.

Figure 2.3: Estimated Change in Net Debt Issuance Following Covenant Violation

This figure shows the estimated causal change in net debt issuance as a function of industry treatment intensity (proportion of new covenant violation in the industry). Both violators, $\tau = 1$, and non-violators, $\tau = 0$, are allowed to have outcomes that are effected by peer treatment, ρ. Predicted changes in net debt issuance are determined according to the double-matching algorithm presented in section 2.2. Confidence intervals at the 10% level are shown as dashed lines, where standard errors are computed via a hierarchical bootstrap. Change in debt issuance is measured as the first difference of the change in balance sheet debt, scaled by lagged assets. Test statistics for group-by-group covariate balance are reported below the figure (see appendix for details). The counterfactual world with no covenant violation is approximated by $y(0, \rho_{\text{min}})$, spillovers are identified by observing $\partial y(0, \rho)/\partial\rho \neq 0$. The results show that non-violators increase net debt issuance as peer treatment increases.

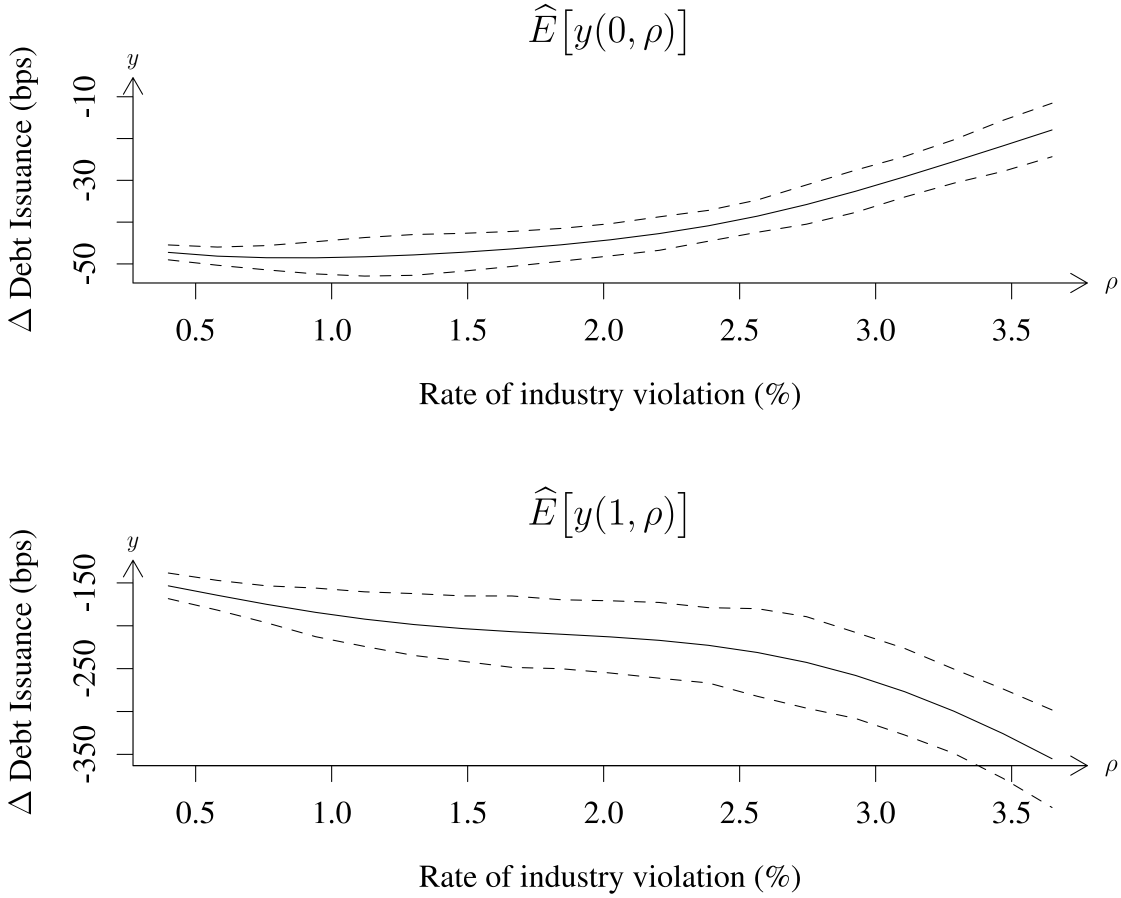

First stage covariate balance (null hypothesis is covariate balance):

- Fisher's Method: $-2 * \sum_{j=1}^{m} \log(p_j) = 0.0011$ (p-value $= 1$)
- Koziol & Perlman: $\sum_{j=1}^{m} J_j = 38.6$ (p-value $= 1$)

In contrast to untreated firms, the expected change in net debt issuance by treated firms decreases as more of their peers find themselves in similar circumstances. When few firms are treated, covenant violators reduce their debt issuance by an average of about 145 basis points. In contrast, when many firms are in a state of covenant violation, treated firms reduce net debt issuance by almost 400 basis points. This makes sense given that more severe shocks are likely to hit a greater number of firms.

To what extent do spillover effects matter? Moving the level of group treatment intensity from 1% to 3% increases the expected level of net debt issuance by untreated firms by 10.9 basis points. This is a 19.6% increase from the mean net debt issuance of 55.427 basis points. Over the same interval, net debt issuance by treated firms falls by 40.4 basis points. Thus, the relative difference in expected outcomes between treated and untreated firms grows by 51.3 basis points, with 21% of that gap driven by indirect treatment effects (spillovers) at untreated firms.

A natural question is whether all firms are equally likely to tap external financing. In particular, the likelihood of a firm approaching creditors for financing should be related to how much cash the firm can raise via operations. Therefore, I partition the expected level of net debt issuance by whether the firm is above or below its industry mean level of operating income (scaled by average assets).[20] Figure 2.4 presents the estimated spillover effects for non-violators conditional on whether the firm has a high amount of operating income.

The results show a clear difference in borrowing between firms that have a high vs. low amount of operating income. The predicted change in net debt issuance for untreated firms with lower income demonstrates a clear indirect treatment effect of rival covenant

[20] Step one of the matching estimator calculates $\hat{\mathrm{E}}[y_{i,j}(\tau_i,\rho)|C_i = c, \rho = \rho_j, x_{i,j}]$, although for convenience the vector x is usually omitted in the text. Partitioning expected outcomes by dichotomizing the k^{th} element of x yields $\mathrm{E}\{\hat{\mathrm{E}}[y_{i,j}(\tau_i,\rho)|C_i = c, \rho = \rho_j, x_{i,j}]|x^k_{i,j} > \bar{x}^k_j\} = \hat{\mathrm{E}}[y_{i,j}(\tau_i,\rho)|C_i = c, \rho = \rho_j, x_{i,j}, x^k_{i,j} > \bar{x}^k_j]$ by the law of iterated expectations. I then estimate the last term in the previous equation as a function of ρ via generalized propensity score matching.

Figure 2.4: Heterogeneities in Effects on Debt Issuance (Non-Violators)

This figure extends the result in figure 2.3 to allow for heterogeneous effects of non-violators, $\tau = 0$. I partition predicted levels of net debt issuance by observations with above and below average operating income in the prior period. The results show that the increase in net debt issuance by non-violators is driven primarily by firms who do not bring in as much cash through operations.

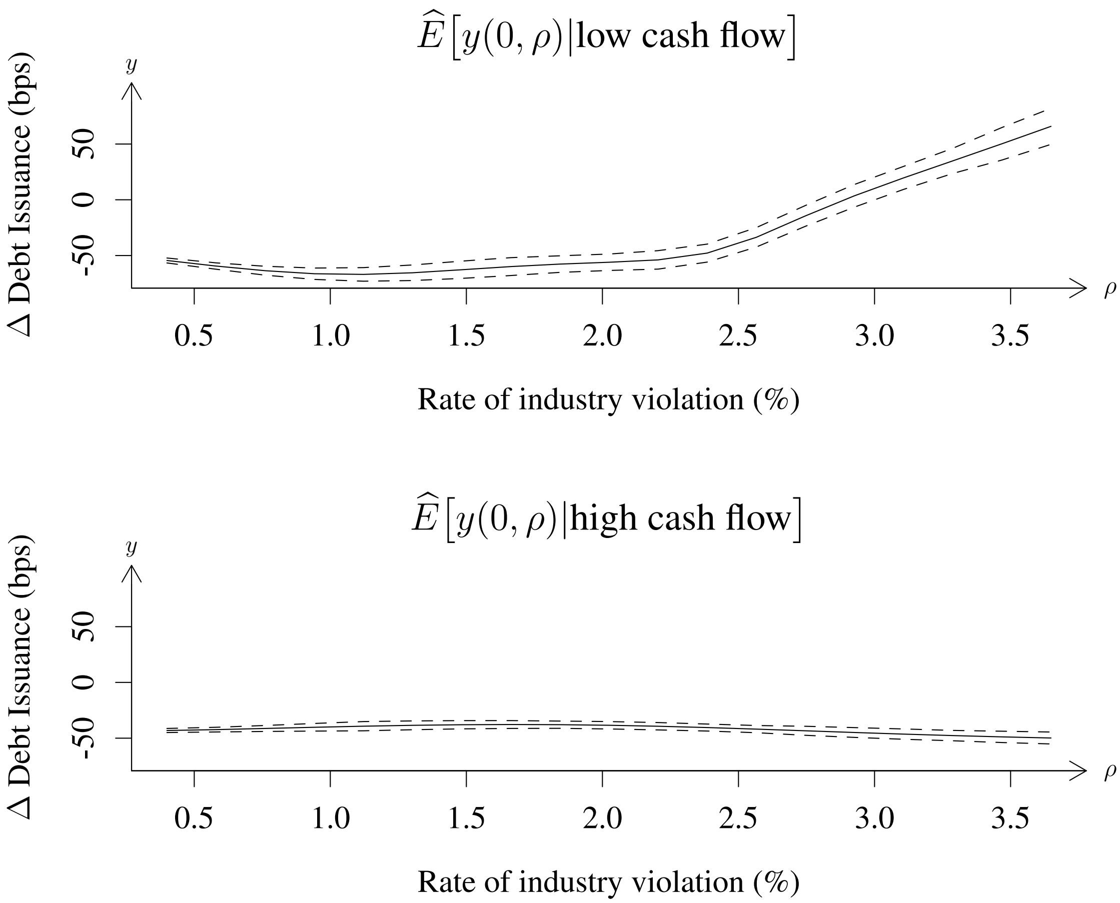

violation. As the proportion of competitors violating covenants rises from 1% to 3%, spillover effects at untreated, low-income firms rise by 200%.

In contrast the untreated, high-income firms show almost no sensitivity in debt issuance to peer-firm treatment. This makes sense, given that these firms raise more cash via operations. A similar effect is found for treated firms in figure 2.5: low income covenant violators lose a greater degree of access to debt markets than do higher income firms.

In sum, the borrowing patterns of untreated firms suggests that "untreated" is indeed a misnomer. These non-violating firms, although characteristically similar to the firms in my sample that violate debt covenants, seek additional debt financing in the period following peer firm treatment. This effect is stronger for firms that are more likely to need external financing to fund aggressive investment policies (as predicted by the Predation Hypothesis).

2.4.3 Strategic Investment

Do competitors increase investment when their peers violate debt covenants? It is clear from figure 2.3 that treated firms lose a substantial degree of financial flexibility following covenant violation. I now turn to whether competitors react to this expected loss in cash flow from financing. Figure 2.6 shows the corresponding change in investment to treated and untreated firms. When few firms violate covenants, the treated (violation) firms reduce investment by an average of roughly 50 basis points. At higher treatment intensity, when the expected decline in net debt issuance is even greater, this reduction balloons to about 550 basis points.

At low levels of treatment intensity, untreated firms exhibit an average increase in capital expenditures by roughly 40 basis points. At the highest levels of observed treatment intensity, this average outcome doubles in magnitude. Untreated competitors raise investment by 80 basis points when more of their peers lose at least part of their access to debt

Figure 2.5: Heterogeneities in Effects on Debt Issuance (Violators)

This figure extends the result in figure 2.3 to allow for heterogeneous effects of violators, $\tau = 1$. I partition predicted levels of net debt issuance by observations with above and below average operating income in the prior period. The results show that the decrease in net debt issuance by violators is driven primarily by firms who do not bring in as much cash through operations since banks may view their repayment risk to be higher.

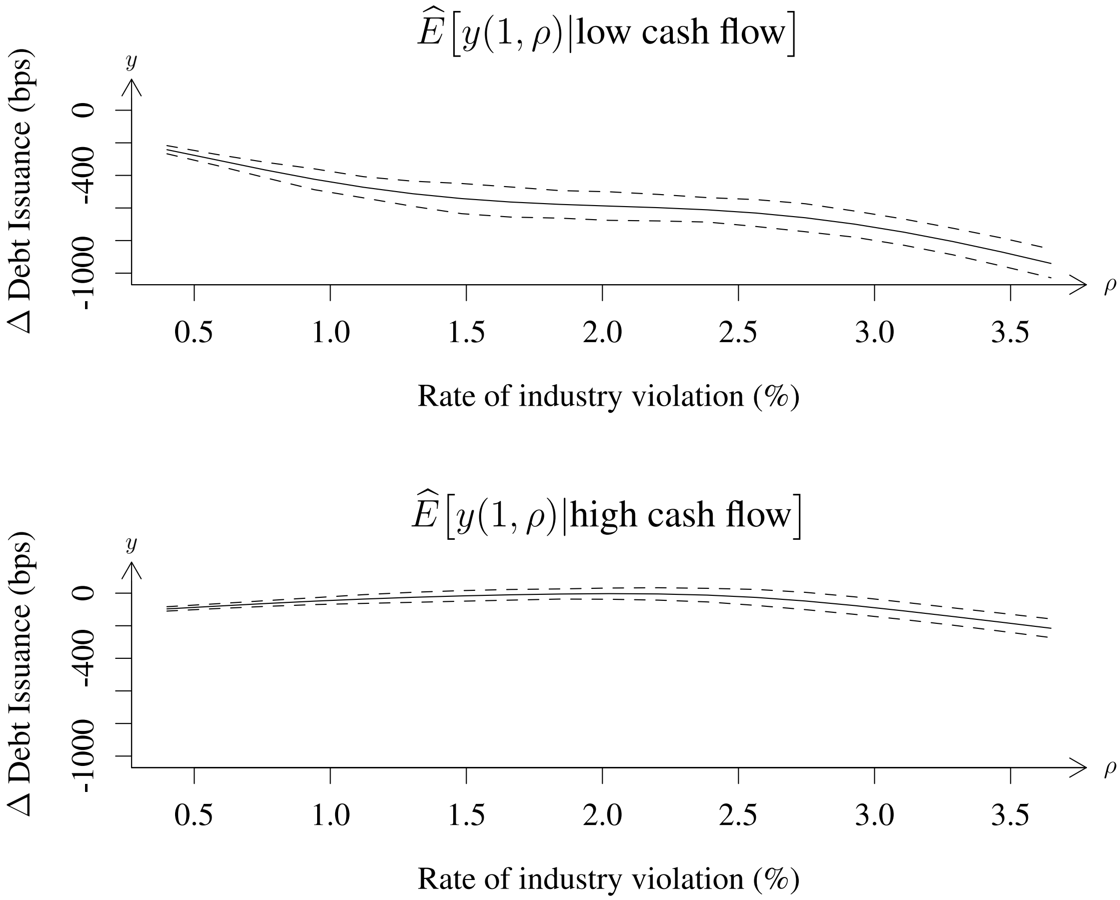

Figure 2.6: Estimated Change in Investment Following Covenant Violations
This figure shows the causal responses – changes in expected outcomes – of treated and untreated firms based on the degree of group treatment intensity. Confidence intervals at the 10% level are shown as dashed lines, where standard errors are computed via a hierarchical bootstrap. Change in investment is measured as the first difference of capital expenditures, scaled by lagged PP&E. Test statistics for group-by-group covariate balance are reported below the figure (see appendix for details).

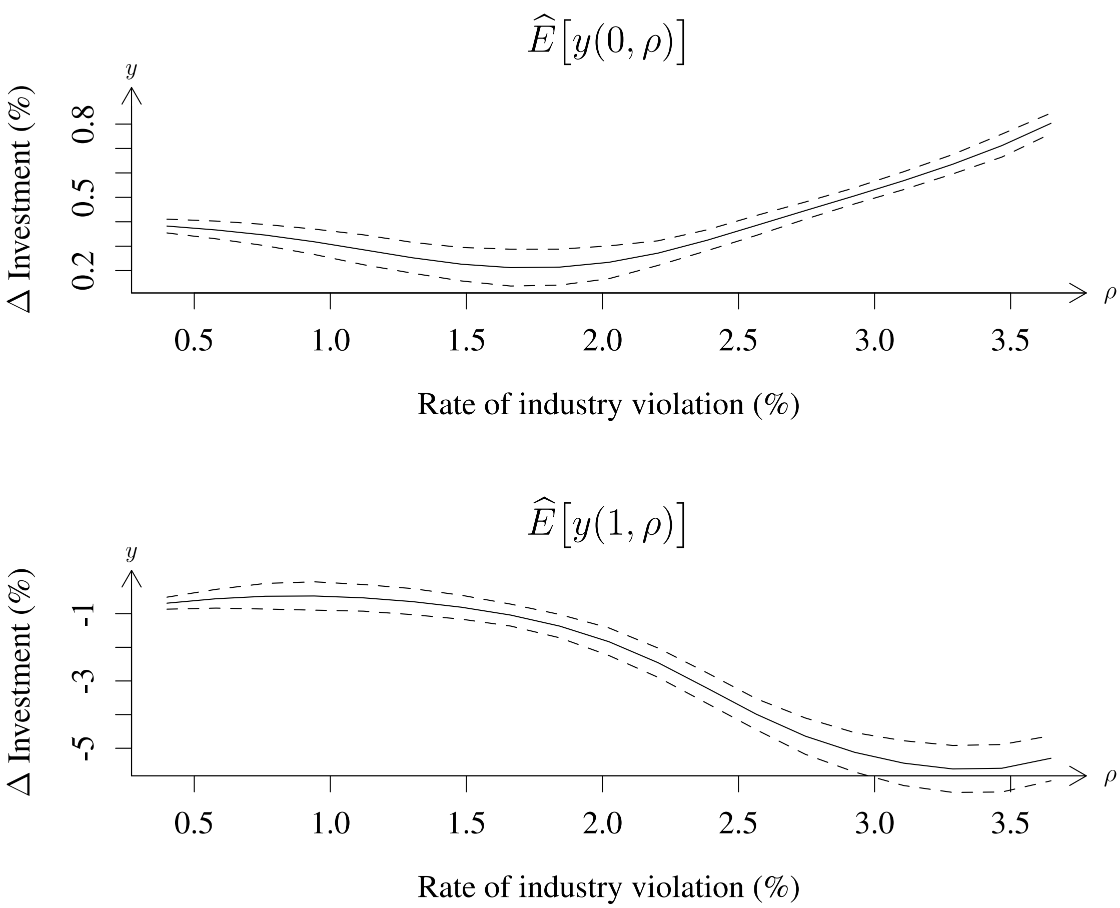

First stage covariate balance (null hypothesis is covariate balance):

- Fisher's Method: $-2 * \sum_{j=1}^{m} \log(p_j) = 0.0013$ (p-value $= 1$)
- Koziol & Perlman: $\sum_{j=1}^{m} J_j = 39.3$ (p-value $= 1$)

financing. The upward-sloping indirect treatment effect is consistent with opportunistic investment motives. Untreated competitors know their peers will lose financial flexibility following a covenant violation (and even more so when treatment intensity is high).

At the mean level of treatment intensity, I find roughly an 80 basis point gap in investment between treated and untreated firms (consistent with the point estimates of Chava and Roberts (2008), whose calculation of an average treatment effect (ATE) includes both a direct and an indirect effect). However, I show that part of this gap is driven by untreated firms increasing their investment above the counterfactual level (the level of investment in a situation where no competitor violates a covenant). For larger shocks that push a larger proportion of firms into violating covenants, the relative gap in investment between treated and untreated firms grows substantially, in part because of more aggressive investment on the part of untreated firms.

As with net debt issuance, spillovers constitute a nontrivial amount to the relative difference between treated and untreated firms. Moving the level of group treatment intensity from 1% to 3% raises the expected level of investment by untreated firms by 23.2 basis points. This is a 2.1% increase over the mean level of investment in the data. In contrast, over the same interval, expected investment by firms that violate covenants falls by 487.3 basis points. Thus, approximately 4.5% of the increase in the difference in expected investment between treated and untreated firms follows from spillover effects to firms that do not violate covenants.

The change in industry investment patterns following a shock to financial flexibility highlights a previously overlooked cost of inflexibility. Whereas previous papers have emphasized the role creditors can play, outside of default, in shaping corporate investment decision-making, none recognize the spillover effect that this has to competitor firms. Moreover, unlike in the theoretical incumbent-entrant model of Hege and Hennessy (2010), wherein an incumbent hoping to deter entry will reduce financial flexibility and consequently lower its innovative output (via debt overhang) and competitor innovation (via entry deterrence), the exogenous shock to realized inflexibility here is shown to increase investment at competitor firms.

2.4.4 Sales

Do firms benefit when rivals violate covenants? The observed increase in investments and net debt issuance by untreated firms suggests their intent is to capture market share from their rivals. I test whether this strategy is successful by considering the expected level of sales growth following violation. Figure 2.7 plots sales growth for treated and untreated firms as a function of industry treatment intensity.

Here, we see that outcomes for untreated firms improve, and outcome for treated rivals deteriorate, as group treatment intensity rises. Consistent with the more aggressive financing and investment decisions shown by untreated firms, these firms exhibit stronger sales growth when more of their peers violate covenants. Again, the spillover constitutes a significant portion of the difference in outcomes between treated and untreated firms. Moving the level of group treatment intensity from 1% to 3% raises sales growth at untreated firms by 1.4 percentage points. This constitutes a 35.7% increase over the sample mean of 3.926% quarterly growth. Over the same interval, sales growth drops 13.3 percentage points at firms that violate covenants. Spillovers therefore make up 9.5% of the increasing gap in expected sales growth between treated and untreated firms.

2.4.5 A More Naive Approach

One natural question that arises from the discussion of SUTVA and use of a hierarchical matching estimator is what cost the econometrician pays for use of an estimator not explicitly meant to incorporate spillover effects. In the context of covenant violation, numerous authors have recognized the endogenous nature of violation: some firms are more prone to violate covenants than others, given their fundamentals. This motivates match-

Figure 2.7: Estimated Change in Sales Following New Covenant Violations

This figure shows the causal responses – changes in expected outcomes – of treated and untreated firms based on the degree of group treatment intensity. Confidence intervals at the 10% level are shown as dashed lines, where standard errors are computed via a hierarchical bootstrap. Sales growth is the percentage increase in sales. Test statistics for group-by-group covariate balance are reported below the figure (see appendix for details).

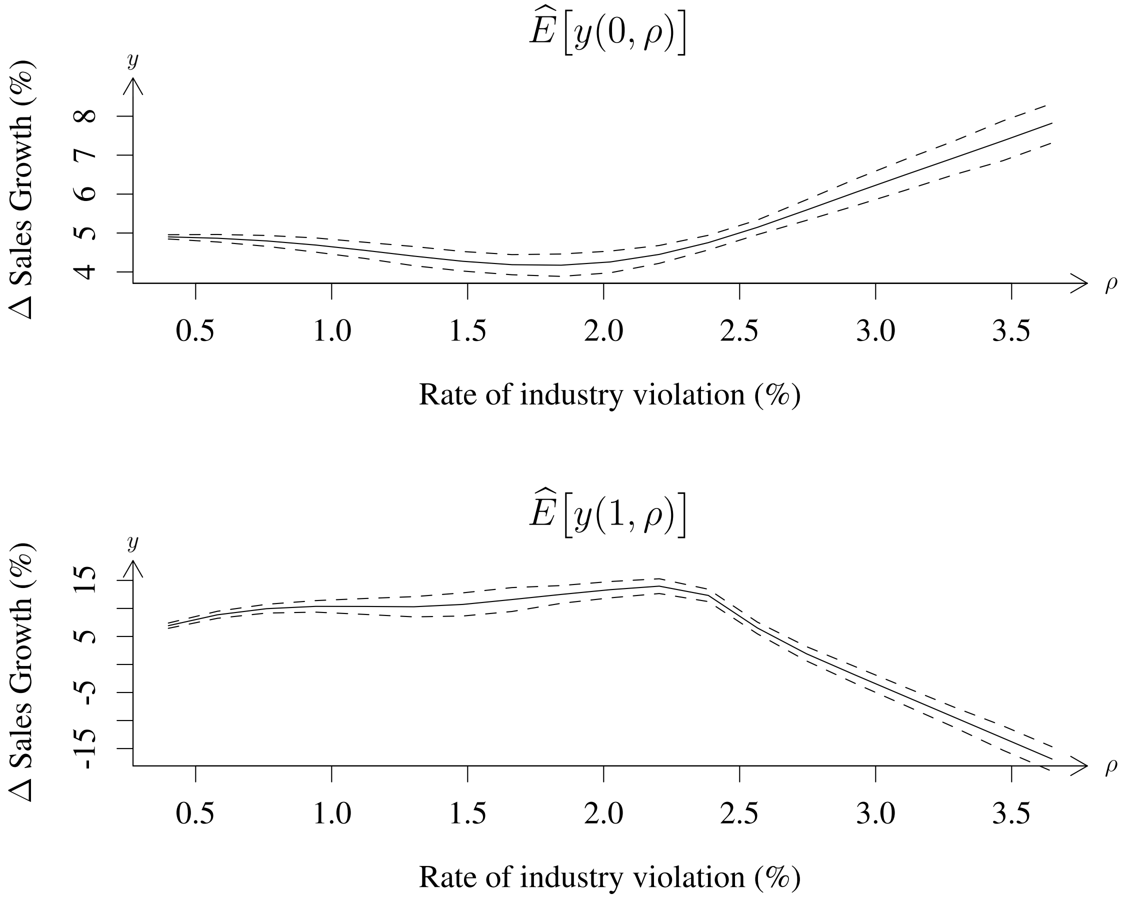

First stage covariate balance (null hypothesis is covariate balance):

- Fisher's Method: $-2 * \sum_{j=1}^{m} \log(p_j) = 0.0019$ (p-value $= 1$)
- Koziol & Perlman: $\sum_{j=1}^{m} J_j = 42.2$ (p-value $= 1$)

ing estimators that contrast outcomes between violators and non-violators by including confounders that predict covenant violation. As discussed earlier, however, the proportion of firms in the industry is also not entirely random. One should expect worse industry performance to correlate with greater rates of covenant violation. Suppose that one were to attempt to estimate the effect of spillovers by simply including ρ, industry treatment intensity, as a control variable in a linear regression framework that one typically finds in the empirical study on covenant violation. That is, add ρ to a regression that would include an indicator for treatment status, firm-level covariates x, and industry-level covariates w. This functional specification follows that of Chava and Roberts (2008), Roberts and Sufi (2009), and Nini, Smith, and Sufi (2012).[21]

Table 2.4 shows the estimation results. Based on the results from the hierarchical matching estimator, the effect of ρ should be positive in sign for net debt issuance, investment, and sales growth. In stark contrast to this claim, the point estimate of ρ is negative in all specifications and statistically distinguishable from zero in some. This means inclusion of ρ into a regression without addressing the endogeneity that some industries will have a higher rate of violation due to observable industry characteristics can lead a researcher to make the wrong conclusion. For example, table 2.4 shows that a naive regression will report that non-violators decrease investment when a higher proportion of firms in the industry violate covenants. This could follow from the fact that a greater rate of covenant violation in the industry is associated with poor industry prospects. The collective evidence of this table clearly shows that inappropriately specified regression frameworks will not consistently luck out and draw the econometrician to the correct conclusion.

[21]This approach is referred to as a "quasi-discontinuity" by Nini, Smith, and Sufi (2012); while not strictly a regression discontinuity, Roberts and Sufi (2009) report that this method and a strict regression discontinuity used on a subsample of firms with at least one of two accounting covenants yield similar estimates on the effect of covenant violation.

Table 2.4: Testing Spillovers Without Addressing Endogenous Rates of Violation

This table presents results for an econometric approach that does not explicitly address the endogeneity of both violation and violation intensity (ρ), unlike hierarchical matching estimates presented in figures 2.3-2.7. Firm level controls, industry-level controls, and squares of these controls are included, though not reported. Industry, time, and fiscal quarter fixed effects are also included. Columns 1 and 2 test for changes in net debt issuance, columns 3 and 4 consider changes to investment, and columns 5 and 6 look at sales growth. Odd numbered columns use data on all firms, whereas even numbered columns show results for the subsample of non-violators. Nonlinear effects for industry treatment intensity (ρ) are included to match the observed nonlinear effects in figures 2.3-2.7.

	(1)	(2)	(3)	(4)	(5)	(6)
Covenant violation	-0.0069**		-0.0047*		-0.0034	
	(0.021)		(0.082)		(0.603)	
Proportion of peers in violation (ρ)	-0.1210	-0.1422	-0.2119*	-0.2398*	-0.2967	-0.3348
	(0.305)	(0.235)	(0.098)	(0.063)	(0.629)	(0.586)
Covenant violation $\times$ ρ	0.1919		0.0555		0.3077	
	(0.201)		(0.709)		(0.321)	
ρ^2	3.0596	3.8173	5.0078	5.5749*	4.4272	5.0935
	(0.301)	(0.206)	(0.126)	(0.090)	(0.771)	(0.738)

Additional controls, industry, time, and fiscal quarter fixed effects not reported.

Marginal effects reported

Standard errors (clustered on industry-quarter) in parentheses

* p <0.10, ** p <0.05, *** p <0.01

2.5 Conclusion

To what extent does financial flexibility (the ability to raise and restructure financing) matter to a firm? Past work has emphasized the direct effects of flexibility: firms endogenously defer external financing to future periods in which they anticipate the arrival of costly investment projects. This paper extends flexibility to competitive effects. Strategic reactions by rivals could mitigate the costs of inflexibility if sources of shocks to flexibility are partially correlated across firms, and thus competitors become more conservative when rivals lose flexibility (Risk Hypothesis). Alternatively, strategic reactions by rivals could exacerbate the costs of inflexibility if these competitors view inflexibility as an opportunity to gain market share (Predation Hypothesis). I show that a change to the flexibility of peer firm financing causally raises competitor net debt issuance, investment, and sales growth. Collectively, these results imply that financial inflexibility makes a firm appear weaker to rivals, consistent with the Predation Hypothesis.

An important aspect of this paper is the use of a spillover effects estimation procedure that can uncover externalities in treatment. SUTVA-based estimators (regression discontinuity, difference-in-differences, propensity score matching) are completely sufficient to establish relative effects. However, a substantially richer understanding of equilibrium firm incentives can come from determining whether externalities to untreated peers are negative or positive. In the present context, strategic responses account for 5 to 20 percent of the variation in the difference in outcomes between firms that do violate covenants and those that do not. Few techniques for measuring spillovers have been implemented thus far, although the nature of competition between firms suggests exogenous shocks to peer firms would warrant response by rivals.

3. DIRECTOR EXPERIENCE AND CYBERSECURITY EVENTS

3.1 Introduction

> Close to half of the directors we surveyed said "we know [cyber security] is an issue, but we don't even know what kinds of questions we should be asking."
>
> – Jean-Marc Levy, interview for NYSE Governance Services' *Boardroom View*

A number of large-scale, high-profile security breaches (Target Corp., Home Depot, Community Health Systems, Yahoo, etc.) has pushed cybersecurity risk to the forefront of public attention. The need for cybersecurity is quite clear; an analysis of S&P 500 firms found that nearly 90% of the aggregate value of these firms came from intellectual property and other intangibles (Stathis 2015). The repercussions of a cybersecurity breach can be quite large: Verizon lowered its bid price for Yahoo by $350 million after Yahoo announced that its users' information had been stolen (Knutson 2017). On average, a firm's stock price drops 0.37% on news of a data breach. A rules-based, "check-the-boxes" style of cybersecurity management appears to be at best a partial solution to the problem. Target Corp., for example, was certified compliant with the Payment Card Industry's (PCI) security standards just weeks before it fell victim to a massive malware attack that stole customer payment information (Bjorhus 2014). The rising need for stronger cybersecurity monitoring has led the issue to become the number one topic in corporate board rooms (FTI 2014). Today, 89% of directors regularly discuss cybersecurity at board meetings (NACD 2017).

Due to the importance of the topic, some have asked the SEC to provide increased guidance regarding how firms should disclose cybersecurity risk in their annual filings (Langevin and Himes 2015). The SEC, in turn, has called on boards of directors to become more proactive in managing cybersecurity risk (Aguilar 2014). There is even a legislative effort to require boards to disclose the presence of a "cybersecurity expert," much like how SOX required the disclosure of a "financial expert" (Reed, Collins, and Warner 2017). Many directors openly admit an inadequate understanding of cyber risk. One survey found that 39% of directors felt they would perform better at cyber risk oversight if they had a better understanding of how to do so (Scally 2014).

Despite the rising discussion of cybersecurity in the popular press, the increasing role cybersecurity plays in boardroom discussions, and the ongoing policy debate about how to best structure cybersecurity monitoring and risk disclosure, little academic work exists on the topic. I study two important questions in this paper. First, how does cybersecurity experience affect a director's risk monitoring? In particular: are firms overseen by experienced directors less likely to suffer a data breach, and do firms that are overseen by experienced directors disclose cybersecurity risk differently than those that are not? Second, are there reputational penalties for experiencing a data breach?

I use a database of publicly reported cybersecurity events (security breaches) at U.S. companies over 2005 to 2016.[1] These data are provided by Privacy Rights Clearinghouse, which aggregates reports of cybersecurity events across the country. To my knowledge, this paper is the first in the finance literature to use this dataset. The data are comprised of public announcements of security breaches across the 47 states that have laws requiring the disclosure of breaches that involve personally identifiable information.[2] The data also include notifications triggered by the Gramm-Leach-Bliley Act and the Health Insurance Portability and Accountability Act of 1996, which are federal regulations covering, respectively, financial institutions and health plans/providers.[3]

I treat these cybersecurity events as conditionally random shocks to a director's expe-

[1] The dataset is available for download at www.privacyrights.com

[2] New Mexico became the 48^{th} state with breach disclosure regulation in April 2017. The two remaining states without disclosure rules are Alabama and South Dakota.

[3] A thorough discussion of state and federal regulations on breach notification laws is available at www.steptoe.com/assets/htmldocuments/SteptoeDataBreachNotificationChart.pdf.

rience with cybersecurity. This variation in director exposure to cybersecurity is valuable in testing the effect of director experience on firm cybersecurity risk. In general, it is difficult to test whether director expertise causally affects a firm's cybersecurity outcomes because of the endogenous matching between directors and firms. Firms at higher risk of data breaches may be more likely to hire a director with cybersecurity experience. Thus, to sidestep this matching issue, I look to announced data breaches and study effects on firms that already employ a director when she is exposed to a cybersecurity event. This removes the endogenous matching between director expertise and firm risk by generating a shock to a director's cybersecurity human capital.

To separate firm-level effects of a data breach from director-level monitoring changes, I focus my analysis of director risk monitoring on the *other* firms at which the director is employed when she experiences a cybersecurity event. I refer to these firms as *cyber-affiliated* firms, since they employ directors affiliated with a data breach at another company. One important point to recognize is that although this identification strategy removes the endogenous matching between a director with cybersecurity experience and a firm, it does not remove the possibility some firms are more likely to be cyber-affiliated than others. Indeed, several observable characteristics of firms are different between those that are and are not cyber-affiliated. To address this issue, I propensity score match cyber-affiliated firms to not-cyber-affiliated firms and show that these observable differences do not exist in the matched set.

Within the matched set of firms for which breach affiliation is assumed to be a conditionally random event, I find that firms that employ a breach affiliated director are 15% more likely to disclose cybersecurity risk as a risk factor in the firm's 10-K filing. The increased likelihood of cybersecurity risk disclosure rises once a director experiences their first data breach somewhere else, and does not change following additional cybersecurity experiences. However, although the incidence of cybersecurity disclosure rises for firms that are connected (via a director) to a data breach, the quality does not. I proxy for the quality of a risk disclosure by considering whether the addition of a cybersecurity disclosure happens via inclusion of a new risk factor to the 10-K filing or via an insertion of cyber-related text into an existing risk factor in the document. The latter case becomes more likely when firms employ a breach-affiliated director.

I next show that the propensity of a firm to suffer a data breach is lower when the firm happens to employ a director that experiences a cybersecurity event at one of their other firms. The marginal probability of a data breach falls 87% when one of a firm's directors experiences a data breach at one of their other appointments. As with the incidence of disclosure, effects on the realized probability of suffering a cybersecurity event are isolated to the first instance in which a director experiences a data breach. Subsequent breaches have zero economic and statistical effect.

Finally, I consider the reputational costs of cybersecurity events by analyzing the likelihood of director turnover following a data breach. Directors who experience more than one cybersecurity event are substantially (59%) more likely to lose their positions at hacked firms. First time directors are more likely to hold on to their positions. There are no effects at a hacked director's other appointments (where a cybersecurity event did not occur). It is well known that directors face turnover risk following monitoring failures (Ertimur, Ferri, and Maber 2012; Fich and Shivdasani 2007; Srinivasan 2005). However, the extent to which the labor market learns about director quality is still unknown. This paper shows that reputational penalties for poor cyber monitoring depend, in part, on the director's history of cybersecurity events. Whereas first-time hacked directors experience no increased turnover risk (and simultaneously, demonstrate higher monitoring going forward), directors that are affiliated with multiple cybersecurity events are punished in the labor market.

3.2 Data

This section introduces the data used in the paper. Firm-level accounting data come from Compustat. Director-level information comes from Boardex. The Compustat sample runs through fiscal year 2015, but the Boardex data ends June 2015. I link Boardex company identifiers to the CRSP/Compustat universe following Engelberg, Gao, and Parsons (2012) – this match procedure utilizes Levenshtein matching, ticker symbols, and International Security Identification Numbers (ISIN) to connect identifiers. Data on cybersecurity events and on risk factor disclosure are unique to my paper and detailed below.

3.2.1 Cybersecurity Events

This paper uses all security breaches reported between 2005 and 2016 and recorded by Privacy Rights Clearinghouse (PRC). This dataset aggregates all reports of cybersecurity breaches published according to a state's breach disclosure law. These laws require individuals in that state to be notified if their personal information is lost by a company entrusted with it. The first state to enact such a law was California in 2002. Following a wave of state-level laws in the mid-2000s, all states except two (Alabama and South Dakota) have disclosure laws in place. The dataset also includes breaches reported in accordance with the Gramm-Leach-Bliley Act or the Health Insurance Portability and Accountability Act of 1996, which are federal regulations covering, respectively, financial institutions and health plans/providers.

Although breach notification laws vary by state, many of the fundamental aspects of the disclosure requirements are similar. Personal information loss that triggers disclosure includes social security numbers, drivers license numbers, or financial information, although some states also extend this to include biometric information, a passport number, date of birth, etc.[4] The disclosure timeline also varies by state. For example, California law

[4]https://www.bna.com/complicated-compliance-state-data-breach-notification-laws/

requires disclosure of a breach be made "in the most expedient time possible and without unreasonable delay" (California Civil Code 1798.29) while in Ohio the disclosure must be made "in the most expedient time possible but not later than forty-five days following its discovery" (Ohio Code 1349.19).

The PRC data categorizes reported data losses by the nature of the event that caused the loss. All cybersecurity events are classified as either: HACK (electronic entry by an outside party), CARD (electronic payment fraud, e.g. capturing credit card terminal data), INSD (data loss due to an insider taking advantage of a system), PHYS (lost, stolen, or discarded paper documents), PORT (lost, stolen, or discarded portable electronic devices), STAT (lost, stolen, or discarded electronic devices not designed to be moved), DISC (unintended/accidental disclosure), or UNKOWN (all other events). I use all intentional forms of attack (HACK, CARD, INSD) in my analysis because accidental events – for example, a PORT event resulting from a misplaced laptop – are less likely to be subject to director level oversight.

The institution name identified by PRC is matched into firm names from CRSP using the Levenshtein algorithm to compare string values. Perfect matches are kept, and the remainder of the data are reviewed by hand and attached to a `permco` identifier where appropriate. Most of the observations in PRC do not link to CRSP because they are small private firms, non-profits (including education), or privately owned health care providers.

One limitation of data breach information is that the severity of the breach is not always clear. A breach in the PRC data will include an estimated number of records lost (i.e., number of customer's who lost personal information) where possible, but this information is not always known. In the event PRC learns new information about the number of records that were lost (e.g. when the company updates their initial report after discovering a broader breach then originally believed), the data in PRC are revised to reflect the new numbers.

3.2.2 Risk Factor Disclosure

Data on firm disclosure of risk factors is text mined from the DirectEdgar database, which pre-processes 10-K filings and separates out Item 1A into an HTML document that I parse with Python's BeautifulSoup html pareser. Item 1A of a firm's 10-K filing discloses certain risk factors that a firm faces. I search for terms relating to cybersecurity risk (e.g. "hackers"). The full list of search terms is provided in the appendix.

I study both the incidence of cyber disclosure and the way in which cybersecurity disclosures are added to a firm's 1A filing. In the former case, a simple text search of whether the firm uses any of the aforementioned cybersecurity risk words is sufficient. In the latter, I am interested in the quality of this disclosure. This is measured in two ways. First, I look at whether a firm begins disclosing cyber security risk by adding a entirely new risk factor to the firm's 1A filing or by adding cybersecurity risk to an existing disclosure. Second, I look at cybersecurity disclosures that include cyber-relevant terminology in the header of the risk factor and compare them to disclosures where cyber terminology is exclusively used in the body of the risk factor. The former features cybersecurity risk more prominently whereas the latter is more likely to be a case in which the firm includes multiple sources of rare event risk (disastrous weather, terrorism, et cetera) in a single factor. Figure 3.1 presents examples of each case.

Figure 3.1: Two Example Disclosures of Cybersecurity Risk

Panel A: Example case of a firm adding a new risk factor that discloses cyber risk

An interruption or breach in security of our information systems may result in financial losses, loss of customers, or damage to our reputation.

We rely heavily on communications and information systems to conduct our business. In addition, we rely on third parties to provide key components of our infrastructure, including loan, deposit and general ledger processing, internet connections, and network access. These types of information and related systems are critical to the operation of our business and essential to our ability to perform day-to-day operations, and, in some cases, are critical to the operations of certain of our customers. The risk of a security breach or disruption, particularly through cyber attack or cyber intrusion, including by computer hackers, has increased as the number, intensity and sophistication of attempted attacks and intrusions from around the world have increased. As a financial institution, we face a heightened risk of a security breach or disruption from threats to gain unauthorized access to our and our customers' data and financial information, whether through cyber attack, cyber intrusion over the internet, malware, computer viruses, attachments to e-mails, spoofing, phishing, or spyware.

Our customers have been, and will continue to be, targeted by parties using fraudulent emails and other communications to misappropriate passwords, credit card numbers, or other personal information or to introduce viruses or other malware through "trojan horse" programs to our customers' computers. These communications appear to be legitimate messages sent by the Bank, but direct recipients to fake websites operated by the sender of the e-mail or request that the recipient send a password or other confidential information via e-mail or download a program. Despite our efforts to mitigate these tactics through product improvements and customer education, such attempted frauds remain a serious problem that may cause customer and/or Bank losses, damage to our brand, and increase in our costs.

Although we make significant efforts to maintain the security and integrity of our information systems and we have implemented various measures to manage the risk of a security breach or disruption, there can be no assurance that our security efforts and measures will be effective or that attempted security breaches or disruptions would not be successful or damaging. Even the most well protected information, networks, systems and facilities remain potentially vulnerable because attempted security breaches, particularly cyber attacks and intrusions, or disruptions will occur in the future, and because the techniques used in such attempts are constantly evolving and generally are not recognized until launched against a target, and in some cases are designed not to be detected and, in fact, may not be detected. Accordingly, we may be unable to anticipate these techniques or to implement adequate security barriers or other preventative measures, and thus it is virtually impossible for us to entirely mitigate this risk. A security breach or other significant disruption could: 1) Disrupt the proper functioning of our networks and systems and therefore our operations and/or those of certain of our customers; 2) Result in the unauthorized access to, and destruction, loss, theft, misappropriation or release of confidential, sensitive or otherwise valuable information of ours or our customers, including account numbers and other financial information; 3) Result in a violation of applicable privacy and other laws, subjecting the Bank to additional regulatory scrutiny and expose the Bank to civil litigation and possible financial liability; 4) Require significant management attention and resources to remedy the damages that result; or 5) Harm our reputation or cause a decrease in the number of customers that choose to do business with us. The occurrence of any such failures, disruptions or security breaches could have a negative impact on our results of operations, financial condition, and cash flows.

10-K Filing by Bank of Hawaii Corporation for fiscal year end December 31, 2011

Figure 3.1 (Continued)

Panel B: Example case of a firm updating an existing risk factor to include cyber risk

Operations risks may adversely affect our business and financial results.

The operation of our electric generation, and electric and gas transmission and distribution systems involves many risks, including breakdown or failure of expensive and sophisticated equipment, processes and personnel performance; operating limitations that may be imposed by equipment conditions, environmental or other regulatory requirements; fuel supply or fuel transportation reductions or interruptions; transmission scheduling constraints; and catastrophic events such as fires, explosions, severe weather or other similar occurrences. In addition, our information technology systems and network infrastructure may be vulnerable to internal or external cyber attack, unauthorized access, computer viruses or other attempts to harm our systems or misuse our confidential information.

We have implemented training and preventive maintenance programs and have security systems and related protective infrastructure in place, but there is no assurance that these programs will prevent or minimize future breakdowns, outages or failures of our generation facilities or related business processes. In those cases, we would need to either produce replacement power from our other facilities or purchase power from other suppliers at potentially volatile and higher cost in order to meet our sales obligations, or implement emergency back-up business system processing procedures.

10-K Filing by The Empire District Electric Company for fiscal year end December 31, 2011

Operations risks may adversely affect our business and financial results.

The operation of our electric generation, and electric and gas transmission and distribution systems involves many risks, including breakdown or failure of expensive and sophisticated equipment, processes and personnel performance; operating limitations that may be imposed by equipment conditions, environmental or other regulatory requirements; fuel supply or fuel transportation reductions or interruptions; transmission scheduling constraints; and catastrophic events such as fires, explosions, severe weather or other similar occurrences.

We have implemented training, preventive maintenance and other programs, but there is no assurance that these programs will prevent or minimize future breakdowns, outages or failures of our generation facilities. In those cases, we would need to either produce replacement power from our other facilities or purchase power from other suppliers at potentially volatile and higher cost in order to meet our sales obligations.

These and other operating events may reduce our revenues, increase costs, or both, and may materially affect our results of operations, financial position and cash flows.

10-K Filing by The Empire District Electric Company for fiscal year end December 31, 2010

In panel A, Bank of Hawaii Corporation adds a risk factor discussing cybersecurity risk to its 10-K filing for the period ending 12/31/2011. Its previous filing, for the period ending 12/31/2010 did not include this factor. Thus, I classify Bank of Hawaii Corporation as including a new cybersecurity disclosure. The risk factor disclosure by Bank of Hawaii Corporation is three paragraphs long, and it discusses both the potential cybersecurity threats faced by the firm as well as possible consequences of a security breach. In panel B, The Empire District Electric Company updates an existing risk factor to include a discussion of cybersecurity risk. I classify this firm as updating on old disclosure to include cybersecurity risk. The added disclosure is one sentence long and recognizes that a cybersecurity threat exists.

To identify the different risk factors in the html document, I analyze the structure of formatting rules used in the document. The headers to each risk factor are assigned special formatting as visual cues to the reader. For example, some documents use bold font, some use italics, some use a different color, etc. I look over the distribution of formatting rules applied to the entire document and assume that the most common formatting structure are the body paragraphs to a risk factor and that the second most common formatting rule is the one identifying risk factor headers. Before reviewing the distribution of formatting rules, I exclude table and list paragraphs (e.g. those including a `<tr>` or `<li>` tag) from the distribution since some factors can include long bullet lists of subitems. I place all table text into the preceeding body paragraph so that it is included when searching for cybersecurity words.

After finding the risk factor headers for a firm in years t and $t-1$, I vectorize the headers using tf-idf and compute the pairwise cosine similarity between each risk factor header in t and each risk factor header in $t-1$. This yields a matrix of cosine scores. To determine whether a factor in t is novel, or whether it is an updated version of a factor in $t-1$, I create a mapping between the factors in t and $t-1$ that maximizes the total pairwise cosine similarity between the two years. Table 3.1 provides an example matrix of cosine similarity scores.

Table 3.1: Example Mapping of Risk Factors Between $t-1$ and t

This table provides an example mapping between risk factors for a firm at time t and time $t-1$. Each cell reports the cosine similarity between rf_t^i and rf_{t-1}^j. In this example, rf_{t-1}^b maps to rf_t^C, rf_{t-1}^a maps to rf_t^B, and rf_t^A is classified as a new risk factor. To establish a mapping between $t-1$ and t I convert each element, $a_{i,j}$ in the matrix (a cosine similarity) by $|1-a_{i,j}|$ and apply the Hungarian Method to the resultant matrix.

	rf_t^A	rf_t^B	rf_t^C
rf_{t-1}^a	0.032	$\mathbf{0.95}$	0.09
rf_{t-1}^b	0.12	0.21	$\mathbf{0.99}$

This matrix of cosine similarity scores records how closely each risk factor header at time $t-1$ matches a risk factor header at time t. I determine the best mapping between the time $t-1$ factors and the time t factors by transforming each element of the matrix, $a_{i,j}$ by $|1-a_{i,j}|$ and applying the Hungarian Method to the resultant matrix. This method maps each risk factor in $t-1$ to a factor in t such that no two factors in $t-1$ point to the same factor at t, and the total cosine similarity of each pair of factors is maximized. In the reported analysis, an assigned pair is considered to be a successful match (the risk factor at t is the same as at $t-1$) if the cosine similarity between the two is at least 0.95 (similarity is measured on a scale of 0 to 1). Results are robust to alternative cutoffs (e.g. 1, 0.9). I allow for scores less than 1, a perfect match, in order to allow for the possibility that risk factor headers could re-phrase or modify text from one year to the next. However, as scores move further below 1, there is an increased likelihood that the assigned $(t-1,t)$ pair of risk factors discuss different concepts.

3.2.3 Summary Statistics

In the analysis that follows, I create two indicator variables that turn on when a director experiences a data breach at one of her firms. First, to allow for differential effects for the first time a director is exposed to a data breach, I create an indicator variable "First experience with cyber" that turns on the first time a director experiences a data breach (it does not turn back on for any additional breaches). Second, I create an indicator variable "Additional experience with cyber" that turns on whenever a director experiences a cybersecurity event after her first breach. A director's first experience may be different than additional experiences because more information is revealed to the director at the first event. In my data, 696 directors are affiliated with a data breach, and 12% of these experience more than one data breach in their careers. Additionally, because about half of corporate boards have 3-year staggered election cycles (Fos et al.), the treatment dummy variables "First experience with cyber" and "Additional experience with cyber" switch on for two years at a time in all of my analysis. This aligns the treatment duration with the average director's post-breach turnover risk concern that I discuss in detail in the next section.

In all specifications, I use a number of controls to capture the potential cybersecurity risk that a firm faces. I include four firm size and performance measures: ROA, log of total assets, log of the number of employees, log of sales. These controls capture observable differences in firms that may lead to a firm being targeted by hackers (e.g. those with more sales may have more customer information to steal, and those with more employees may have more employee information to steal). Additionally, I control for observable board characteristics: independence of the board, average director age, board size, board busyness, whether the firm has a tech committee, and whether the firm has a risk committee.[5] I also control for the firm's cyber history: whether the firm has had a data breach in the past,

[5]These committee assignments are not given standard names in the Boardex database and the appendix lists the committee names I use to classify a committee as assigned "tech" or "risk."

whether the firm disclosed cyber risk in its past filings, and (if disclosure was made) the length of the cyber disclosure. Finally, I control for the number of industry competitors (at a 2-digit NAICS level) that experience a data breach.[6]

Table 3.2 reports the summary statistics for all variables used in the paper. Panel A reports director-level observables. The variables of interest here are an indicator for whether the director experiences their first data breach (across all board appointments past and present) and an indicator for whether the director gains additional experience with cybersecurity via a data breach. These director-level variables are used in predicting turnover hazards of directors following cybersecurity events. To control for whether a director has specific experience with technology oversight, I create a classification for "tech experts." This is an indicator variable for whether the director was ever appointed (across all board positions past or present) to a technology committee. Definitions for what committees I consider to be technology committees are reported in the appendix.

Panel B reports firm-level observables. The variables of interest here are an indicator variable for whether any director at the firm experiences their first data breach (at one of their other board appointments) and an indicator variable for whether any director at the firm experiences an additional data breach (at one of their other appointments). These are used in predicting changes in cybersecurity monitoring at the firm level. In my sample, 35% of firms report cybersecurity in their annual disclosure of risk factors.

3.3 Cybersecurity Monitoring

Does cybersecurity experience affect a director's monitoring of cybersecurity risk? The assumption implicit in calls for increased cybersecurity expertise on boards is that directors without such expertise will not be as capable at monitoring this type of risk. One source of resistance to such a requirement is that the labor pool of qualified, "cybersecurity experts"

[6]Finer partitions to more detailed industry categorizations are impossible due to the relative sparseness of data breaches in the data.

Table 3.2: Summary Statistics

This table reports summary statistics of the data used. Panel A reports director-level variables, while panel B reports firm-level variables. All variables are defined in section 2.

Panel A

	Mean	Std. Dev.
First experience with cyber	0.009	0.094
Experience with cyber	0.012	0.111
Age	61.108	9.039
Age over 65	0.331	0.470
Female	0.111	0.315
Num. directorships	1.610	0.987
Independent dir.	0.940	0.237
Previous cyber exp.	0.009	0.092
On tech committee	0.014	0.117
On risk committee	0.019	0.138
On audit committe	0.581	0.493
Tech expert	0.031	0.173
Risk expert	0.034	0.181
Audit expert	0.753	0.431

Panel B

	Mean	Std. Dev.
Firm is hacked	0.007	0.084
Discloses cyber risk	0.358	0.479
First experience with cyber	0.052	0.223
Additional experience with cyber	0.037	0.189
Percent independent	0.828	0.087
ROA	-0.025	0.217
log(total assets)	6.661	2.057
log(number of employees)	1.220	1.197
log(sales + 1)	5.999	2.205
Avg. director age	60.481	4.792
Number of directors	8.456	2.466
Busy board	0.039	0.194
Firm has tech. committee	0.030	0.170
Firm has risk committee	0.032	0.177
Hack(s) in industry	8.420	8.204

may be insufficient to supply the needs of public companies. Therefore, as an alternative, others have attempted to address the issue of IT in the boardroom by proscribing a set of governance questions directors should ask management regarding IT policy (including cybersecurity).[7] It is not clear then, whether having a director classified as a "cybersecurity expert" on the board is necessary to effectively monitor cybersecurity risk.

The empirical challenge in identifying an effect of cybersecurity experience on monitoring cybersecurity risk is the endogenous matching between director skills and firm demands. Firms at greater risk of suffering a cybersecurity event, for example, may be more likely to hire a cybersecurity expert. To identify an effect of cybersecurity experience on director monitoring, I use cybersecurity events reported in the PRC dataset to measure observable changes in a director's first-hand experience with cybersecurity risk. Because changes in cybersecurity monitoring at a breached firm cannot be isolated to director choices, the emphasis of my tests is on changes in cyber monitoring at a breach-affiliated director's *other* firms (if they have any other board appointments). I ensure that these board positions were in place at the time of the reported data breach to avoid the possibility that these directors are endogenously hired for their cybersecurity experience. I then test whether the change in cyber experience resultant from a data breach changes a director's monitoring policy at these other firms at which they work as directors.

One remaining identification challenge to studying cybersecurity monitoring at breach-affiliated firms is that some firms may be more likely to employ a director that is exposed to greater cybersecurity risk at their other appointments. Panel A of table 3.3 shows that many observable firm characteristics differ between firms that are and are not breach-affiliated. For example, breach affiliated firms have larger firms, because having more directors on the board raises a firm's chance that at least one director experiences a data breach. Larger firms are more likely to be breach affiliated, likely because larger firms employ directors

[7]This was the approach used by the Canadian Institute of Chartered Accountants in 2004.

who have positions at other large firms, and large firms have a higher probability of being breached.

I wash out these observable differences between firms by propensity score matching breach-affiliated firms to those that are not breach-affiliated.[8] Descriptive statistics for the matched subsample are reported in panel B of table 3.3. Across all observable dimensions included in this study, matched firms appear identical to breach-affiliated firms. Under the assumption of unconfoundedness, whether or not a firm in this matched sample is breach-affiliated is a conditionally random event.

After matching, I use two outcomes to indirectly infer the amount of cybersecurity monitoring at a firm. First, I study the disclosure of cybersecurity risk in Item 1A of a firm's 10-k filing. Second, I look to the estimated likelihood of a firm suffering a data breach. These two measures capture, respectively, the disclosure of risk and the realization of risk at a firm. Unfortunately, the specific cyber policies of a firm are not publicly reported, and expenditure on IT is not specifically broken out in a firm's reporting of capital expenditures. However, despite these limitations, the two proxies of monitoring I use are quite powerful. Disclosure of cyber risk is itself a topic of much policy debate, and the propensity of a data breach event is exactly the outcome of interest in cyber monitoring. Stronger monitoring should result in greater risk disclosure and a decreased likelihood of experiencing a data breach, and one would therefore expect that breach affiliated firms would exhibit a higher propensity to disclose cybersecurity risk and a lower propensity to be hacked.

3.3.1 Cybersecurity Risk Disclosure

Are directors with cybersecurity experience more likely to recognize cyber as a form of risk? Although SEC guidance on risk factor disclosure recommends these factors be

[8]To ensure all firms are eligible to be classified as breach-affiliated, only firms whose directors hold multiple appointments are considered in the matching estimation.

Table 3.3: Mean Differences Between Indirectly Treated and Untreated Firms

This table reports summary statistics for two subsamples: firms that employ a breach-affiliated director (where the breach occurs at some other board appointment for the director) and firms that do not. Panel A reports statistics for the full sample. Panel B reports statistics for the matched sample where observations are matched according to the propensity score for being breach-affiliated.

Panel A

	Means			Medians		
	Untreated	Treated	*p*	Untreated	Treated	*p*
Percent independent	.825	.866	0	.857	.889	0
ROA	-.03	.035	0	.02	.054	0
log(total assets)	6.517	8.338	0	6.576	8.482	0
log(sales + 1)	5.83	7.961	0	5.961	8.228	0
log(number of employees)	1.112	2.463	0	.731	2.51	0
Number of board connections	11.884	23.266	0	10	22	0
Avg. director age	60.467	60.642	.206	60.714	61	.114
Number of directors	8.313	10.11	0	8	10	0
Hack(s) in industry	8.486	7.653	0	6	6	.167
Discloses cyber risk	.243	.355	0	0	0	0
Firm was previously hacked	.025	.098	0	0	0	0
Busy board	.031	.132	0	0	0	0
Firm has tech. committee	.028	.048	0	0	0	0
Firm has risk committee	.032	.036	.442	0	0	.442

Panel B

	Means			Medians		
	Untreated	Treated	*p*	Untreated	Treated	*p*
Percent independent	.867	.865	.577	.889	.889	.88
ROA	.041	.035	.169	.055	.054	.516
log(total assets)	8.303	8.323	.743	8.336	8.467	.278
log(sales + 1)	7.954	7.947	.907	8.079	8.197	.304
log(number of employees)	2.437	2.451	.764	2.382	2.51	.664
Number of board connections	23.241	23.164	.843	23	22	.691
Avg. director age	60.722	60.642	.589	61	61	.595
Number of directors	10.055	10.093	.647	10	10	.229
Hack(s) in industry	7.647	7.676	.908	6	6	.306
Discloses cyber risk	.355	.355	.995	0	0	.995
Firm was previously hacked	.09	.097	.491	0	0	.491
Busy board	.115	.128	.25	0	0	.25
Firm has tech. committee	.057	.049	.283	0	0	.282
Firm has risk committee	.035	.036	.899	0	0	.899

as firm-specific as possible, this is often not the case (*Concept Release on Business and Financial Disclosure Required by Regulation S-K*). Therefore in this section I will investigate not only the incidence of cybersecurity risk disclosure but also the types of disclosures that are made.

I study the propensity to disclose cyber risk by modeling this disclosure as a function of firm-level observables. The effects of interest are the coefficients on "First experience with cyber" and "Additional experience with cyber" which are dichotomous variables that capture whether any board member is exposed to a data breach *at a different firm.* I also control for whether the firm itself was breached, whether the firm disclosed cyber risk in the past, as well as firm-level performance variables (e.g. sales) that might affect a firms likelihood of being targeted by a data breach and governance variables (e.g. board independence) that capture observable differences in board structures. All regressions include time and industry fixed effects to absorb constant latent time and industry variation in cybersecurity disclosure.

I approach risk disclosure in three ways. First, I test whether director exposure to a data-breach changes a firm's incidence of cybersecurity risk disclosure in the firm's 10-k filing. As discussed in section 3.2.2, the incidence of risk disclosure is determined by performing a text search on Item 1A of the 10-K filing for words related to cybersecurity risk.

Table 3.3.1 reports estimated marginal effects of director cyber experience on the probability of a firm disclosing cyber risk. As a baseline, the typical firm in my matched subsample is 50% likely to report cybersecurity risk. This incidence rises by 7.4 percentage points (p-value < 0.00) after a director at the firm experiences their first data breach at one of their other boards. The 7.4 percentage point increase in disclosure likelihood translates into an 15% increase over the baseline probability of disclosing cybersecurity risk. A firm's propensity to disclose cybersecurity risk does not significantly change when any of

Table 3.4: Probability of Disclosing Cybersecurity Risk

This table estimates the probability of a firm disclosing cybersecurity risk as a function of firm-level observables. All variables are defined in section 2. The variables of interest are “First experience with cyber” and “Additional experience with cyber.” The former indicates that at least one of the firm’s directors is connected to a cybersecurity event at another firm and that this is the director’s first experience with a data breach. The latter indicates that at least one of the firm’s directors is connected to a cybersecurity event and that this is not the director’s first experience with a cyber event. Industry and time fixed effects are included but not reported. Coefficient estimates on additional control variables, listed in Panel B of Table 3.2, are not reported for brevity. Column one uses a logit specification, column two uses a complementary log-log model, and column three uses a linear probability model.

First experience with cyber	0.074***	0.049**	0.041***
	(0.020)	(0.023)	(0.014)
Additional experience with cyber	0.013	-0.012	0.016
	(0.028)	(0.026)	(0.018)
Firm was breached	-0.037	0.039	0.021
	(0.095)	(0.077)	(0.072)
Controls	Yes	Yes	Yes
Pr(Disclose | Experience = 0)	.49	.49	.48
N	2246	2246	2256

Industry and time fixed effects not reported.

Marginal effects on the probability of the outcome variable are reported.

Standard errors (clustered on industry) in parentheses

* p <0.10, ** p <0.05, *** p <0.01

its directors experience additional data breaches at other appointments.

The incidence of cybersecurity risk disclosure, however, does not yield insight into how risk disclosure is made. In particular, as discussed in section 3.2.2, firms can include mention of cybersecurity risk by one of two means. First, firms can add a new risk factor to their 10-K filing. Second, firms can update an existing risk factor to include a discussion of cybersecurity risk.

I construct two proxies to capture the quality of a new cybersecurity risk disclosure. First, I divide cybersecurity disclosures between those that happen by adding a new risk factor to a firm’s filing and those that happen by updating an existing risk factor. This captures the quality of a disclosure under the assumption that a newly added risk factor is more specific to cybersecurity whereas a risk factor that is updated to include a mention of cybersecurity is less specific.[9] Second, I divide cybersecurity disclosures between those that include cybersecurity words in the header of a factor and those that only include cybersecurity terms in the body of the factor. This captures the quality of a cybersecurity disclosure under the assumption that disclosures that feature cybersecurity more prominently in the header of the risk factor text do a better job of highlighting the risk associated with cybersecurity events.

Table 3.5 reports the estimated marginal effects of a director's outside experience with cybersecurity on a firm's disclosure of cybersecurity risk. In Panel A, I proxy for the quality of a new cybersecurity disclosure by dividing new disclosures between those that happen via a new risk factor and those that happen via an addition to an existing risk factor. I assume that new risk factors are more specific descriptions of cybersecurity risk and therefore of higher quality. The marginal change in the probability of switching from no new cybersecurity disclosure to a new cybersecurity disclosure via a new risk factor is a 13.7 percentage point increase (p-value < 0.00) for firms that employ a breach experienced director. The marginal change in the probability of switching from no new cybersecurity disclosure to a new cybersecurity disclosure via an addition to an existing risk factor is a 18 percentage point increase (p-value < 0.00) for firms that employ a breach experienced director. Thus, a new data breach experience for a director significantly changes the likelihood that a connected firm discloses cybersecurity risk, consistent with table 3.3.1. However, the last column of table 3.5 shows that the marginal change in the probability of disclosing cybersecurity risk via updating an existing risk factor is 4.3 percentage points

[9] Strictly speaking, an existing risk factor that adds a cybersecurity component is necessarily less than 100% on the topic of cybersecurity. It is possible however that a newly risk factor that includes a cybersecurity component is still a "kitchen-sink" disclosure that adds other risks besides cybersecurity risk into a single disclosure, and therefore no more specific.

Table 3.5: Probability of Adding Cybersecurity Disclosure

This table estimates the probability of a firm adding a cybersecurity risk disclosure to its 10-K filing by considering categorical outcomes proxying for the quality of a new disclosure. Marginal effects computed from a multinomial logistic regression are reported. Controls, time fixed effects, and industry fixed effects are included but not reported below to conserve space. Control variables are identical to those in table 3.3.1. Standard errors are clustered on the industry and reported in parentheses below the marginal effects.

Panel A

Case 1: No change in cybersecurity disclosure (baseline)

Case 2: At least one *new* factor disclosing cyber is added

Case 3: The only new cyber disclosure addition is through updating an *existing* factor

	Pr(2)-Pr(1)	Pr(3)-Pr(1)	Pr(3)-Pr(2)
First experience with cyber	$.137^{***}$	$.180^{***}$	$.043^{*}$
	(.037)	(.041)	(.022)
Additional experience with cyber	.046	.047	.000
	(.059)	(.076)	(.036)

Panel B

Case 1: No change in cybersecurity disclosure (baseline)

Case 2: Cyber is added to the *header* of a risk factor (more prominent)

Case 3: Cyber is added only to the *body* of a risk factor

	Pr(2)-Pr(1)	Pr(3)-Pr(1)	Pr(3)-Pr(2)
First experience with cyber	$.058^{*}$	$.096^{**}$	.039
	(.033)	(.047)	(.026)
Additional experience with cyber	.019	.061	.042
	(.028)	(.037)	(.027)

higher (p-value $= 0.05$) than the probability of disclosing cybersecurity risk via release of a new risk factor. This implies that although the *incidence* of cybersecurity disclosure is positively affected by director experience with data breaches, the *quality* of disclosure, if anything, falls.

Panel B of table 3.5 uses the prominence of cybersecurity in the risk disclosure as a proxy for the quality of the disclosure. In particular, I assume that use of cybersecurity language in the header of a risk factor constitutes a higher quality disclosure than cases in which cybersecurity is exclusively discussed in the body of the factor. This is a second method to parse out differences between risk factors that exclusively discuss cybersecurity risk versus risk factors that lump in cybersecurity risk with many other rare events (natural disasters, terrorism, et cetera). The marginal probability of a firm disclosing cybersecurity risk in the header of a risk factor increases by 5.8 percentage points (p-value $= 0.08$) when a firm's board includes a breach affiliated director. The marginal probability of disclosing cybersecurity risk exclusively in the body of the risk factor's text increases by 9.6 percentage points (p-value = 0.04) over the baseline case of no new cybersecurity disclosure. While the differences in marginal effects here is not significant at conventional levels (p-value = 0.14), the point estimate is consistent with panel A: breach-affiliated directors increase the incidence of cybersecurity disclosure, but more often through lower quality disclosure.

3.3.2 Future Cybersecurity Events

As a second measure of director cybersecurity monitoring, I use the propensity of the firm to suffer a data breach in the following year. One limitation of cyber risk disclosure is that it could be cheap talk. The board may encourage the firm to disclose cyber risk merely in an attempt to limit their liability to shareholders without taking real steps to ensure cybersecurity events are made less likely. Therefore, in this subsection I study the realized likelihood of suffering a data breach to learn whether directors with cybersecurity experience lower their firms propensity to suffer a data breach. Importantly, I look at shocks to a director's experience with cybersecurity and its effect on firms *at which the director is already employed.* This removes the potential that breach-affiliated directors will endogenously seek out less risky board appointments in the future.

Table 3.3.2 presents the estimated marginal effects on a firm's propensity to suffer a

Table 3.6: Probability of Experiencing a Data Breach

This table estimates the probability of a firm suffering a data breach as a function of firm-level observables. All variables are defined in section 2. The variables of interest are "First experience with cyber" and "Additional experience with cyber." The former indicates that at least one of the firm's directors is connected to a cyber security event at another firm and that this is the director's first experience with a data breach. The latter indicates that at least one of the firm's directors is connected to a cyber security event and that this is not the director's first experience with a cyber event. Industry and time fixed effects are included but not reported. Coefficient estimates on additional control variables, listed in Panel B of Table 3.2, are not reported for brevity. Column one uses a complementary log-log model specification, column two uses logit, and column three uses a linear probability model.

First experience with cyber	-0.004*	-0.005*	-0.016**
	(0.002)	(0.002)	(0.006)
Additional experience with cyber	0.001	0.000	0.008
	(0.001)	(0.001)	(0.011)
Firm was previously hacked	0.004**	0.004**	0.081***
	(0.002)	(0.002)	(0.017)
Controls	Yes	Yes	Yes
Pr(Breach | Experience = 0)	.0048	.0053	.0285
N	2651	2651	2651

Industry and time fixed effects not reported.

Marginal effects on the probability of the outcome variable are reported.

Standard errors (clustered on industry) in parentheses

* p <0.10, ** p <0.05, *** p <0.01

data breach. Note that for completeness, I report results for a linear probability model alongside the logit and complimentary log-log (sparse outcome robust) models. However, the predicted probabilities of suffering a data breach are negative in over 35% of the observations under a linear probability model which suggests that the linear model specification is a poor fit for the data.

The estimated marginal effect of employing a director who experiences their first data breach at another firm is to lower the firm's propensity to suffer a data breach by 0.5 percentage points (p-value 0.04). This lowers the firm's probability of having a cybersecurity event by an economically significant 87%. In contrast, employing a director who experiences additional data breaches at another firm does not change the firm's risk of having a cybersecurity event (p-value 0.98).

The combined results of tables 3.3.1, 3.5, and 3.3.2 imply that directors who are exposed to their first corporate data breach increase their cyber monitoring at their other current firms. These connected firms (those that also employ these directors) are more likely to disclose cyber risk and less likely to experience a data breach. Monitoring gains do not change when directors are exposed to additional cybersecurity events.

3.4 Director Turnover

Given that directors increase cybersecurity monitoring after experiencing a data breach, and that their monitoring does not increase for each additional breach experienced, how do labor market outcomes change in response to breach events? Past work documents reputational costs (increased turnover) for directors following intentional misrepresentations at the firm (Ertimur, Ferri, and Maber 2012; Fich and Shivdasani 2007). Higher turnover risk also exists for audit committee members following restatements (Srinivasan 2005) because restatements indicate a monitoring failure by the directors on this committee. In this section, I investigate turnover outcomes for breach-affiliated directors.

Similar to Srinivasan (2005), I consider the possibility that directors will face differential reputational penalties for working at a firm that suffers a data breach based on the specific role of the director. In my context, I search the Boardex database for all committee names that indicate the director is assigned to a technology role at the board (e.g. "cyber security," "information technology and security," or "I.T. oversight"). The full list of committee names that I classify as a technology role are provided in the appendix. Directors assigned to a technology role on the board may have more direct oversight of cyber risk and therefore stronger reputational penalties following a data breach. The number of firms with technology-related committees is quite small. Thus, as an additional check, I also classify directors as technology experts if they have ever served on a technology-related committee on any of their past or present board appointments. Directors who are technology experts may be perceived to have greater responsibility over cyber policies at a firm even if that firm does not have its own technology committee.

As an additional check for heterogeneous reputational penalties, I follow survey evidence from the National Association of Corporate Directors regarding the allocation of cyber responsibilities at the board (NACD 2017). In this survey, only 5% of directors report that their firm assigns cyber risk oversight to a technology committee. This is consistent with the observed rarity of technology committees in the Boardex data. Per the survey, 11% of firms assign the cyber risk oversight role to the risk committee, 51% assign the role to the audit committee, and 41% assign the responsibility to the full board (the survey allowed directors to select multiple responses). Because I cannot observe which committee at a particular firm is assigned the bulk of cyber monitoring duties, I create a dummy variable "cyber responsibility" in the following way. If the board has a technology committee, I assume the majority of cyber monitoring falls with it. If the board does not have technology committee but has a risk committee, I assume the majority of cyber monitoring falls with it. If the board has neither a technology committee nor a risk committee, I assume the principal responsibility for monitoring cyber risk falls on the audit committee.

Finally, I consider an interaction effect of cybersecurity experience with whether or not a firm had previously disclosed cybersecurity risk. The interaction term here is a test of whether or not directors that previously warned shareholders of the potential for a cybersecurity event are less likely to lose their job following a data breach.

I model the turnover hazard for a director using the Cox proportional hazard model. I control for director-level observables: age, an indicator for whether the director is over 65, the number of directorships the director holds, an indicator for if the director is female, an indicator for if the director is classified as independent, indicators for whether the director is on the technology, risk, or audit committees (where technology and risk are defined as described above), indicators for whether the director is a technology, risk, or audit expert (again, as defined above), and whether the director had any experience with a data breach before being hired at the firm. I also control for firm-level observables: return on assets, log of total assets, the percent of institutional holdings, the percentage of independent directors at the firm, board size, and whether the firm previously disclosed cyber risk in its annual filings.

Table 3.7 presents the estimated hazard ratios. Recall that in a hazard regression, a ratio of 1 indicates no effect. When a firm reports a data breach, and a director at the firm has previously experienced one or more cybersecurity events, the director's hazard of turnover is 1.59 (p-value < 0.00). This hazard ratio implies that the director is 59% more likely to leave the firm relative to the baseline case in which no breach is reported. When a firm reports a data breach and this is the first time that a director at the firm is associated with a cybersecurity event, that director faces a hazard of turnover of 0.67 (p-value $=0.04$). Thus directors who experience their first data breach are 33% less likely to turn over in the following year. This is consistent with staying on the board to review and clean up the aftermath of a cybersecurity event. A director's hazard of turnover at a particular firm does not change if one of the director's other firms reports a data breach.

I find no evidence that the reputational cost of experiencing a data breach varies according to a director's committee appointments. I test for these effects in columns 2-4 of table 3.7. The statistically insignificant effects on various committee memberships, interacted with breach announcements, is consistent with the apparent uncertainty regarding which of a firm's directors are assigned a cyber monitoring role (NACD 2017).

Table 3.7: Hazard of Director Turnover Following a Cybersecurity Event

This table reports hazard ratios for the right hand side variables of interest used in predicting a director's turnover. In columns 2-4, variable Z is defined by the column header. All variables are defined in the text. The table shows increased director turnover following data breaches, and that this turnover risk changes as a function of whether the director is associated with data breaches one or more times.

		Tech_Committee	Tech_Expert	Cyber_Responsibility	Mentions_Cyber
First experience with cyber (firm breached)	0.672**	0.668**	0.635**	0.557**	0.588**
	(0.129)	(0.130)	(0.125)	(0.151)	(0.124)
First experience with cyber (firm breached) × Z		1.183	1.641	1.387	2.004
		(1.519)	(1.469)	(0.531)	(1.024)
Experience with cyber (firm breached)	1.586***	1.613***	1.755***	1.571**	1.760***
	(0.253)	(0.260)	(0.286)	(0.335)	(0.301)
Experience with cyber (firm breached) × Z		0.559	0.306	1.037	0.563
		(0.597)	(0.225)	(0.325)	(0.251)
First experience with cyber (other firm breached)	0.718	0.706	0.721	0.707	0.824
	(0.154)	(0.154)	(0.164)	(0.225)	(0.196)
First experience with cyber (other firm breached) × Z		1.685	0.870	0.988	0.448
		(2.169)	(0.571)	(0.420)	(0.259)
Experience with cyber (other firm breached)	1.169	1.167	1.190	0.944	1.113
	(0.200)	(0.203)	(0.215)	(0.234)	(0.218)
Experience with cyber (other firm breached) × Z		1.229	0.949	1.563	1.163
		(1.314)	(0.511)	(0.523)	(0.449)
Controls	Yes	Yes	Yes	Yes	Yes
N	129419	129419	129419	129419	129419
N directors	36299	36299	36299	36299	36299
N turnovers	11781	11781	11781	11781	11781

Hazard ratios reported

* p <0.10, ** p <0.05, *** p <0.01

3.5 Conclusion

This paper uses a novel database of cybersecurity events (public announcements of data breaches) to study director learning and reputational penalties. Directors who gain first-hand experience from a data breach at one of their firms put this knowledge to work at their other appointments. These connected firms are 15% more likely to disclose cybersecurity risk in their annual filings, and 87% less likely to suffer a data breach. These two results imply that directors with real cybersecurity experience are more successful at monitoring cybersecurity risk. Given reports that 89% of directors regularly discuss cybersecurity at board meetings (NACD 2017), it seems unlikely that these effects are purely driven by a saliency channel.

Interestingly, although the cybersecurity experienced directors lower a firm's likelihood of suffering a data breach and increase a firm's propensity to disclose cybersecurity risk, these directors do not lead firms to more detailed disclosures of cybersecurity risk. Rather, experienced directors are more likely to add a cyber risk disclosure to an existing risk factor (instead of creating a new risk factor dedicated to cybersecurity risk) and less likely to feature cybersecurity prominently in the header of the risk factor's text. One explanation for this could be that experienced directors recognize the limitations of more detailed disclosure in fully anticipating possible threats to the company and therefore chose to disclose the risk via a broader, more general statement.

4. CEO NARCISSISM, HUMAN CAPITAL, AND FIRM VALUE

4.1 Introduction

We examine empirically how CEO narcissism affects firm value through the channel of human capital, in particular through the increased turnover of employees, both at the executive level and in the general labor pool. A large literature in finance, management and accounting (discussed below) examines how CEO psychological characteristics impact firm investment, acquisition policy, and profitability. There is little work, however, on how such characteristics impact firm value through the channel of managing and retaining a firm's employees and their human capital. Narcissism is a consequential and clinically well-defined personality trait that has been thoroughly studied in the psychology literature. It has been found to be highly prevalent among people in powerful positions, including CEOs (see e.g., Judge, LePine, and Rich (2006)), and it has a strong effect on how an individual interacts and works with people (see e.g., Nevicka et al. (2011)). CEO narcissism is thus likely to significantly impact how firms utilize their human capital. Labor costs account for over half of U.S. GDP, so human capital is undoubtedly an important input to firm value.[1] Thus, the research question of how CEO narcissism impacts firm value through the human capital channel is of first order importance, and yet it remains unexplored.

Both theory and empirical findings in the extant literature on narcissism motivate our examination of how CEO narcissism impacts firm value through the human capital channel. One of the defining characteristics of narcissists is a lack of empathy (see e.g., Maccoby (2004)). Hence, narcissism in the CEO could increase firm value by making the CEO less hesitant to terminate executives and workers who are under-performing. Prior

[1] For 2015, GDP and total labor compensation were, respectively, \$17.95 and \$9.68 Trillion, according to the Bureau of Economic Analysis.

research additionally suggests, however, that narcissists also tend to underestimate the value of skills in others, possibly leading narcissists to reduce firm value by underestimating the loss of value in human capital when making termination decisions. In addition, prior research indicates that others find it unpleasant to work with narcissists, so CEO narcissism could harm firm value by causing the voluntary departures of non-CEO executives who possess valuable, firm-specific human capital that is costly to replace.[2]

The observations above motivate the hypothesis that the likelihood of turnover of non-CEO top management team (TMT) members should increase in the presence of higher levels of CEO narcissism. Following Aktas et al. (2016), we use a measure of CEO narcissism previously validated in the psychology literature: the ratio of first person singular pronouns to the sum of first person singular and first person plural pronouns in unscripted CEO speech (e.g., Foster and Campbell (2007), Raskin and Shaw (1988).) Consistent with the hypothesis, we find that CEO narcissism increases the odds of departures of other top management team (TMT) members (+13.6% as narcissism moves from one standard deviation below to one standard deviation above the mean). We further examine how the CEO's narcissism is related to the stock price reaction to news of a TMT member's departure. The effect of narcissism on the announcement return is negative if the departing TMT member has been with the firm for three years or longer. Since longer-serving TMT members are more likely to have accumulated valuable firm-specific human capital, this result implies that the tendency of narcissists to undervalue human capital can reduce firm value when CEO narcissism induces the turnover of long-serving TMT members.

Secondary results support the hypothesis that TMT member turnovers in a firm with a narcissistic CEO are driven by narcissism. Regardless of CEO narcissism, non-CEO executives who are closer to the CEO in pay prominence (adjusted for industry norms)

[2]See e.g., Bushman and Thomaes (2011), Morf and Rhodewalt (2001), Zeigler-Hill, Myers, and Clark (2010), and Hogan, Curphy, and Hogan (1994).

are less likely to depart. But, the positive relation between the likelihood of non-CEO departure and CEO narcissism is significantly amplified for non-CEOs with greater pay prominence. Theory and evidence from the psychology literature suggests that a narcissistic CEO would likely perceive other highly-paid executives as drawing or threatening to draw attention away from the CEO, whose narcissism causes him to crave such attention. Thus, highly-paid TMT members are either forced out by the narcissistic CEO or they leave voluntarily because it is unpleasant to work with him. Although the data do not permit us to reliably distinguish between terminations versus voluntary departures, we note that either cause can lead to the loss of valuable human capital.[3]

We also find a positive relation between CEO narcissism and mass layoffs of non-executive employees. CEO narcissism increases the odds of a mass layoff (+15.3% as narcissism moves from one standard deviation below to one standard deviation above the mean), but how the stock market reacts to the news of narcissist-initiated layoffs depends on the value of workers' human capital. Because we cannot obtain data on tenure for non-executive employees as we do for TMT members, we use the average real hourly wage in the industry as our proxy for the value of human capital for workers. While controlling for industry unit labor costs, we find that the stock market reacts more negatively to a narcissist-initiated mass layoff than to that of a non-narcissist when the value of human capital is high. We infer that the tendency of narcissistic CEOs to undervalue human capital destroys firm value through the channel of increased mass layoffs in firms with high levels of human capital. The results for both the likelihood of mass layoffs and the stock price reaction to mass layoffs of higher-wage employees are parallel to those for the TMT members, which suggests that CEO narcissism has effects on the human capital of

[3]Although still not a definitive test, the relation between non-CEO executive turnover and CEO narcissism is weaker when the non-CEO would forfeit more pay if he were to leave the firm. The fact that the potential pay forfeited upon departure matters is consistent with the proposition that at least a non-trivial proportion of the departures are voluntary.

both executive and non-executive employees.

Our work contributes to two important strands of literature. First, we extend the literature on the effects of CEO psychological traits on firm value. Malmendier and Tate (2015) review the large finance literature on the effects of CEO optimism and overconfidence.[4] Other studies examine the effects of CEOs' life experiences.[5] Though narcissism has been less studied in finance, Aktas et al. (2016) examine its effects on takeover negotiations and Ham, Seybert, and Wang (2017) its effects on investment policy and profitability. Kaplan, Klebanov, and Sorensen (2012) examine the effects of several personality traits related to narcissism. There is also a large literature in management that examines how CEO narcissism relates to firm performance, innovation and other firm actions.[6] We add to this strand of literature by showing a link between CEO narcissism and the turnover of human capital, and by showing that the degree to which this link impacts firm value depends on the value of the human capital of the departing employees. To our knowledge, we are the first to show a link between a CEO psychological trait and the loss of human capital by firms at both the non-CEO executive level and the broad non-executive employee level, as well as the associated effects on firm value. Our results complement the large set of extant results that link CEO psychological traits to other capital (e.g., investment, mergers and acquisitions, etc.). It is worth emphasizing that there may be off-setting benefits to narcissistic CEOs that make their net contributions to firm value positive despite their negative effects on human capital; thus, our results about human capital effects should not be taken to imply that employing narcissistic CEOs is suboptimal.

[4] A sampling of this literature includes Banerjee et al. (2015) and Banerjee, Humphery-Jenner, and Nanda (2014), Billett and Qian (2008), Bouwman (2014), **cjrs11** Deshmukh, Goel, and Howe (2013), Galasso and Simcoe (2011), Goel and Thakor (2008), Graham, Harvey, and Puri (2013), Hirshleifer, Low, and Teoh (2012), Kaplan, Klebanov, and Sorensen (2012), Kolasinski and Li. (2013), **mt05**; Malmendier and Tate (2005) and Malmendier and Tate (2008), Malmendier, Tate, and Yan (2011), Otto (2014), and Roll (1986).

[5] See **bs03** Benmelech and Frydman (2015), Graham and Narasimhan (2004), Malmendier and Nagel (2011), Xuan (2009), Yim (2013).

[6] See **ch11**; Chatterjee and Hambrick (2007), Gerstner et al. (2013), **wpl13** Patel and Cooper (2014), and Zhu and Chen (2015a) and Zhu and Chen (2015b).

Second, we extend the literature on the executive labor market by focusing on the turnover of non-CEO executives and on the value of their human capital. While there is an extensive body of research on the causes and consequences of CEO turnover, there is comparatively much less research on the antecedents and effects of the turnover of non-CEO executives.[7] In the management literature, some early research focused on the rate of turnover of non-CEO executives, but there has been relatively little work on specific predictors of individual top manager turnover and the associated valuation effects.[8] Here, we provide evidence that an aspect of the CEO's personality contributes to non-CEO executive turnover, as well as how it impacts firm value through the human capital channel. Our finding of more negative stock market reactions to departures of non-CEO executives with longer tenure from firms led by narcissistic CEOs also sheds light on the value of the human capital of non-CEO executives.

A potentially important byproduct of our research is the validation of the use of the ExecuComp database as a tool for future non-CEO executive turnover studies. While ExecuComp is the most widely used data source for executive compensation research, it has not been used to study non-CEO executive turnover in part because of concerns about whether it is an appropriate tool for this task. In our paper, we use hand-collected data to validate a subsample of turnovers inferred using an algorithm applied to ExecuComp data, and we thereby demonstrate that non-CEO executive turnover can be studied on a large scale with an existing database.

4.2 Background and Hypothesis Development

Narcissism is a well-defined and well-studied concept in the fields of social psychology and organizational behavior. Research in this area classifies narcissism into one of two

[7] Walsh (1988), Murphy and Zummerman (1999), Parino (1997), Huson, Parrino, and Starks (2001), Goyal and Park (2002).

[8] Wagner, Pfeffer, and O'Reilly C. A. (1984), Walsh (1988), Walsh and Ellwood (1991), and Wiersema and Bantel (1993).

types, grandiose or vulnerable, with both demonstrating egocentric tendencies. Either form of narcissism would suggest a biased valuation of the input of peers or subordinates, as well as a lack of empathy, but an understanding of the two types helps to pin down the mechanisms through which CEO narcissism could operate. It is worth noting that a single narcissist can exhibit both of these two forms of narcissism from time to time.

Grandiose narcissists tend to have a very high self-image or ego. This manifests itself in numerous ways. Such narcissists tend to be charismatic individuals (Maccoby 2004), and this charisma causes others to view them positively, at least initially (Carlson, Vazire, and Tf 2011). They want to have more authority, tend to be more expressive, and view themselves as superior (Emmons 1987; Raskin and Terry 1988). Grandiose narcissists overestimate their intelligence and attractiveness (Gabriel, Critelli, and Ee 1994). Often, they view themselves positively because they are high performers. Grandiose narcissists' sense of superiority leads them to feel entitled to the admiration and praise of others, and they tend to react aggressively if they do not get the recognition and praise they feel they deserve (Rosenthal and Pittinsky 2006). A grandiosely narcissistic CEO, therefore, is less likely to treat the ideas of others as equal to his own. Such a CEO's egocentrism causes him to over-emphasize (in his mind) his own contribution to the firm while discounting the value of others' involvement, while at the same time making it difficult for him to empathize with others (e.g., Maccoby (2004)).

The second type, vulnerable narcissists, need frequent reinforcement of their extremely positive self-image. This need for reinforcement, often described as "fragility," has also been demonstrated in multiple studies. For instance, such narcissists tend to be extremely sensitive to criticism or slights (Maccoby 2004), and they constantly work to enhance their own self-esteem (Campbell 1999). If a vulnerable narcissist becomes threatened or his ego declines, he is more likely to become aggressive (Bushman and Thomaes 2011; Horton and Sedikides 2009). Vulnerable narcissists tend to over-react to relatively small failures and problems (Horton and Sedikides 2009), and their egos are more fragile (Zeigler-Hill, Myers, and Clark 2010). Although such narcissists are fragile, they are also grandiose in their self-concept (Morf and Rhodewalt 2001) and have a strong desire to gain self-enhancement (Wallace and Baumeister 2002). These traits make vulnerable narcissists lack empathy (e.g., Maccoby (2004)), similarly to grandiose narcissists. These traits also make a vulnerably narcissistic CEO more likely to take dissenting opinions personally, rather than treat them as good faith objective disagreements about how to best maximize firm value. As a result, such a CEO is less capable of impartially incorporating feedback from others, thereby de-valuing their input. Finally, since vulnerable narcissists share with the grandiose type the tendency to overestimate their own value, they will also likely under underestimate the value and contributions of others.

As shown in prior research, both types of narcissists' high ego, vulnerable narcissists' fragility, as well as grandiose narcissists' sense of entitlement to the admiration of others, make it difficult for narcissists of either type to work with other competent adults in teams.[9] Thus, we would expect CEO narcissism to create problems within the top management team (TMT) of a firm. An egocentric CEO who puts excessive stock in his own contribution to the firm, as in the case of a grandiose narcissist, or who feels threatened by the skills of his peers, as in the case of a vulnerable narcissist, should be more likely to fire a TMT member, regardless of the fact that the TMT member might have valuable ideas or unique abilities. Moreover, we note that a narcissistic CEO of either sort, whose self-centered personality makes it difficult to empathize with others, is less likely to feel an internal, social cost to firing a person. In addition, the large ego of a narcissist of either sort, the constant validation demanded by a vulnerable narcissist, and the entitlement of a grandiose narcissist make it unpleasant to work for such individuals. The unpleasant work

[9]See, e.g., Carlson, Naumann, and Vazire (2011), Carlson, Vazire, and Tf (2011), Farwell and Wohlwend-Lloyd (1998), John and Robins (1994), Maccoby (2003) and Maccoby (2004).

environment, in turn, could make competent TMT members more likely to leave a firm voluntarily if the CEO is a narcissist.

Given the greater likelihood that narcissistic CEOs undervalue or feel threatened by others, and the greater difficulty of working with them, we formulate our first hypothesis:[10]

Hypothesis 1a: CEO narcissism should be associated with higher TMT Turnover.

Because vulnerable narcissists need positive reinforcement of their self-image, they tend to feel threatened when others around them are recognized or otherwise perceived as valuable relative to the CEO (Farwell and Wohlwend-Lloyd 1998). Within a firm, the pay of an executive is a measure of how much the firm values that executive. Hence it is likely that CEOs who are vulnerable narcissists will feel relatively more threatened by other TMT members who are highly paid relative to them. A grandiose narcissist, on the other hand, will tend to underestimate the value of other highly paid executives and not take their ideas seriously, which will make working with a narcissist relatively more unpleasant for more highly-paid non-CEO executives who presumably have relatively more valuable ideas. We thus formulate a corollary to our first hypothesis:

Hypothesis 1b: The effect of CEO narcissism on TMT turnover will be greater if the TMT member has more prominent (higher) pay relative to the CEO.

Because narcissists of both types lack empathy, narcissistic CEOs will have fewer emotional qualms in terminating a member of the TMT they perceive as performing poorly, which should be value increasing if the member indeed has low value. On the other hand, as narcissists of both types tend to overestimate their own abilities while underestimating that of others (e.g., Gabriel, Critelli, and Ee (1994)), they will tend to underestimate the

[10] As we describe later in the discussion of our data and results, the data do not permit us to classify whether TMT members are involuntarily terminated or depart voluntarily.

loss of value in human capital when making termination decisions. This could lead to suboptimal termination decisions. Underestimating the loss of value in human capital associated with a TMT member departure also means that the narcissistic CEO should be less likely to attempt to retain a TMT member who threatens to leave voluntarily unless the difficult working conditions under the narcissistic CEO improve. As with termination, this could lead to suboptimal decisions to not retain valuable TMT members. We thus formulate the following hypothesis:

Hypothesis 2: The effect of CEO narcissism on the stock price reaction to a TMT member departure will be more negative when the TMT member's human capital value is higher.

The hypotheses above relate to members of a firm's top management team. We next develop hypotheses about a firms' lower-level employees. Except for it being unpleasant to work with a narcissistic CEO, the effects discussed for TMT members should also operate at lower levels of employees, albeit at a broader level. Moreover, it seems plausible that when a CEO is contemplating a mass layoff of a firm's employees due to poor performance or the need to restructure, greater empathy for the firm's employees will, all else equal, reduce the likelihood that the CEO decides to implement the layoff. Thus, a lack of empathy and an underestimation of lower-level employees' value should make narcissistic CEOs of either sort more willing to initiate mass layoffs of non-executive workers, leading us to our next hypothesis:

Hypothesis 3: Narcissistic CEOs are more likely to initiate mass layoffs of non-executive workers, and such decisions should depend less on poor firm performance.

As discussed above in the context of TMT members, however, the tendency for narcissists of either sort to undervalue others will likely lead to their underestimating the potential human capital cost of layoffs. As a result, we predict the following:

Hypothesis 4: The effect of CEO narcissism on the stock price reaction to a layoff announcement will depend on the value of the human capital of the workers being laid off. Specifically, narcissism will have a more negative effect for layoffs of employees with higher value human capital.

4.3 Sample Selection, Data and Construction of Variables

To generate our sample, we begin with the population of firms covered in the Compustat Executive Compensation (ExecuComp) database. We utilize two sources of data to construct our measure of executive narcissism: (i) The Wall Street Transcript (twst.com) interview transcripts; and (ii) conference call transcripts available from Lexis-Nexis. Both have been used in prior studies of CEO narcissism (e.g. Aktas et al. (2016), Chatterjee and Hambrick (2007)). The Wall Street Transcript dataset begins in 1997, while the Lexis-Nexis conference call transcript database starts in 2001; thus, our analysis sample begins in 1997. Given the construction of our primary measure of TMT turnover and its look-ahead verification requirement (discussed below), the analysis sample ends in 2009. We obtain dates of mass non-executive worker layoff announcements from Ravenpack, and data on industry-level unit labor costs and real wages from the Bureau of Labor Statistics. We obtain dates of news of non-CEO executive departures from firms' 8-K filings. Stock price and return data are from CRSP. Other control variables are collected from the BoardEx and Compustat databases. Given variable construction and sample requirements, our final estimation sample uses roughly 1,800 firms and includes over 11,000 non-CEO executives with over 3,500 turnover events and over 2,200 CEOs. We next discuss our variables in more detail.

4.3.1 CEO Narcissism

Prior studies in the finance literature use one of two measures of CEO narcissism: the size of the signature in the annual report (Ham, Seybert, and Wang 2017) or the ratio of first person singular to the sum of first person singular and first person plural pronouns computed from analysis of CEO speech (Aktas et al. 2016), while studies in management use others. We use the pronoun ratio for three primary reasons. First, to our knowledge, pronoun usage is the only unobtrusive narcissism measure proposed in prior research in any discipline that has been validated against direct, clinical measures of narcissism, such as the Narcissism Personality Index (NPI) (Foster and Campbell 2007; Raskin and Shaw 1988). Furthermore, linguistic usage in general has been found in the psychology literature to be an accurate way of measuring personality traits (Pennebaker and King 1999). Second, only the pronoun measure can be constructed using machine-readable data for a sample as large as ours. Third, and perhaps most importantly, we note that only the first person pronoun ratio in a CEO's speech is under the CEO's unilateral control (we only include the unscripted Q&A portion of the conference calls), whereas alternative measures of narcissism can plausibly be determined or influenced by others. The measure based on the CEO's signature size in the annual report, along with other measures used outside of Finance, such as the differential in CEO compensation relative to other executives, the prominence of the CEO's photograph in the annual report,[11] and personal mentions of the CEO in company reports (e.g. Chatterjee and Hambrick (2007)), all potentially reflect the decisions or inputs of others. Our narcissism measure, like all others employed in CEO studies, cannot distinguish between grandiose and vulnerable narcissism. However, our hypotheses make identical predictions about the effects of both types of narcissism, so our inability to distinguish between the types is not a significant drawback in our context.

[11] The photograph size measure cannot be constructed for a large sample from machine-readable data.

We construct the CEO pronoun usage ratio from transcripts of conference calls and unscripted interviews. Neither interview transcripts from The Wall Street Transcripts nor conference call transcripts from Lexis-Nexis have linking variables to the ExecuComp dataset. For both sources, we apply the Levenshtein algorithm to link CEO names in ExecuComp to speaker names in the interview and conference call samples. This algorithm computes the distance between strings by calculating the number of changes required to be made to one string variable to make it equal to another. Below are examples using actual names of CEOs in our data:

"Nigel Travis" and "Nigel Travis" has a match score of 0 (perfect match)

"Steven Cooper" and "Steve Cooper" have a match score of 1

"David King" and "Dave King" have a match score of 2

"William Merritt" and "William Moffitt" have a match score of 3

"David Singer" and "Arvind Sanger" have a match score of 4

As demonstrated by the above, match quality deteriorates substantially as distances grow past one. In the reported results, we include all matches with a score of zero or one. Untabulated results show that the inference from our main results is the same when the sample is expanded to include matches with scores equal to two or three. We drop observations where CEO narcissism is not observed (that is, we do not observe CEO participation in the conference call or in a Wall Street Transcript interview) – this is the primary source of observation loss in our study.

Once we match transcripts with each CEO, we compute the ratio of the number of first person singular to the sum of first person singular and first person plural pronouns uttered by the CEO in each transcript. We are careful to separate the CEO's speech from that of interviewers, analysts and other executives participating in a conference call. For each point in time in the data, the CEO's narcissism is the mean pronoun ratio over all transcripts matched to the CEO up through that time period, weighted by the total number of words in each transcript. Our sample mean raw narcissism score is 0.184, which is similar to the mean pronoun-ratio scores of previous studies (e.g. Chatterjee and Hambrick (2007) and Aktas et al. (2016))., with a standard deviation of 0.08. The distribution of raw narcissism is highly skewed, however, so we transform it by adding a positive constant and taking the natural logarithm. The constant is chosen such that the skewness of the transformed variable is as close to zero as possible, and the transformation is defined for all observations. Finally, to ease interpretation, we standardize the transformed variable by subtracting the sample mean and dividing by the sample standard deviation. Reported statistics for narcissism use this transformed and standardized variable.

4.3.2 Top Management Turnover and Mass Layoffs

In principle, we wish to define turnover as a dummy variable equal to one for observations where a non-CEO top manager leaves the firm, and zero otherwise. We thus need to identify as precisely as possible when a manager leaves a firm. Given the large size of our executive pool, we rely on automated methods using data from the ExecuComp database. Securities and Exchange Commission (SEC) regulations require firms to report their top five wage earners in each fiscal year, though some voluntarily disclose the pay for more than just the top five. ExecuComp records the individuals listed in the compensation table in the firm's proxy statement. The major challenge with using this data source to identify turnover is the possibility that an executive may disappear from the compensation table not because he leaves the firm, but because he is no longer one of the top five highest paid executives. For approximately one-third of the executives who disappear from a given firm in ExecuComp, there is no ambiguity because we can use the leftco variable (defined as the date the manager left the company) in the database to determine precisely which year the executive leaves.

To measure turnover of a top manager who has a missing leftco value, we develop a definition based on other data from the ExecuComp database. We first estimate the proportion of executives in ExecuComp who disappear from the top five in a given year and then reappear in a subsequent year, which would indicate that not appearing in the compensation table is not equivalent to exit. Of the non-CEO executives who disappear from the compensation table and are later observed at that same firm, 36.4% are gone for only one year, 84.7% are gone for no more than three years, and 93.8% are gone for no more than five years.

Given the proportions of disappearance and reappearance, we then establish a rule to define turnovers based on the number of consecutive years that an executive is absent from the top five data after dropping out of it. There is a clear trade-off in choosing the number of consecutive years that we require. Because our measure of turnover is forward-looking for every year, we lose our ability to use the future year in our subsequent analysis. Our data on top-five wage earners extends through 2012, so requiring three years of forward-looking data means that we must end our main analysis sample at 2009. If we were to require five years of forward-looking data to confirm turnover, we would have to end our main analysis sample at 2007, which would lead to the loss of two additional sample years. Thus, if we require more years of subsequent absence from the top-five before classifying an executive as turned over, we minimize instances where we classify a disappearance from the top-five compensation as a turnover, when it is really just a rotation in the most highly compensated individuals. But, the cost is that we must shorten our sample period to allow observation of the longer period of subsequent years. If we require fewer years of subsequent absence from the top five, we lose accuracy but we can use a longer sample period. Our subjective weighing of the costs and benefits leads us to classify a turnover event as all instances in which an executive leaves the firm's reported top earners list and does not reappear in the list by the third year.

To assess the validity of our turnover definition, we randomly select a sample of 50 turnover events fitting the definition, and hand-verify whether or not the executive had actually left the company using the more detailed executive lists that firms provide in their 10-Ks, supplemented with Internet searches of news items, LinkedIn, and similar sources. Our ExecuComp-based approach correctly identifies turnover in over 90% of the cases. Given potential remaining concerns about the algorithm, we also verify that our tabulated results are robust to different variations of our turnover classification by re-running our analyses (not shown) using definitions of turnover that range from disappearing from the top executive list from one to four years.

To further assess robustness, we run additional analyses, not shown, that use a much more stringent definition of turnover. Specifically, we collect all of the executive titles in ExecuComp (over 30,000 variations) and re-code those titles into one of 14 categories following work by Graffin et al. (2008). Based on these new titles, we then classify an event as an executive turnover only if a particular executive with a given title does not appear in the subsequent year, but someone else holding that same title code does appear. This approach only works in situations in which the executive's title code is unique (e.g. "Chief Financial Officer"). Events with non-unique title codes are right-censored in this analysis unless turnover is verified with the leftco variable provided by ExecuComp. If an executive with a unique title code leaves the firm's data in the same year that a new executive with the same title code enters, we define the former as having left the firm. Because this test requires unique title codes, it reduces our number of observed turnovers by roughly 900. When we count as turnover only these events, however, our results are unchanged from those tabulated.

For testing the hypotheses that link CEO narcissism and mass layoffs of lower-level employees, we need to identify firm layoffs. We obtain the firm identities and dates of mass layoffs over the 1997-2012 period from the Ravenpack database. Sources include Dow Jones Newswires, the Wall Street Journal, and Barron's. Using article similarity tags provided by RavenPack, we are able to identify not only the event date of a layoff announcement, but also whether the company has had any mention of a layoff over the last three months, across all news services tracked by the company. Only the first mention of each layoff is included in our analysis.

4.3.3 Relative Pay of the non-CEO Executive

Our hypothesis 1b is that a given executive's pay relative to the CEO will amplify the relation between narcissism and TMT turnover. We measure executive pay prominence using the difference between the CEO's salary and bonus pay and the TMT member's salary and bonus pay. We do not rely on stock grants or option grants (or associated option exercises) in compensation because such components are much lumpier through time than are salary and bonus. To normalize magnitudes across firms, we divide this difference by the TMT member's pay. We then use the log transform on the pay gap ratio due to skewness in the data that is generated by a handful of outlier CEOs (e.g. Steve Jobs); our results are similar when we winsorize the data instead of using a log transform.

The size of the gap in pay between the CEO and other TMT members could plausibly be partially driven by industry factors. For example, in more complex industries, where the CEO is expected to have a broader skill set, one would expect a general tendency for a higher pay gap. Thus, we adjust the log-transformed pay gap measure by the average industry log pay gap for executives with the same title using four-digit NAICS codes to classify industries. Finally, we take the negative of this number to ease in interpretation: higher numbers indicate pay closer to the CEO's pay level (thus greater prominence). For example, for firm i in industry j, the CFO's pay prominence is given as:

$$\text{Pay prominence for CFO}_i$$
$$= -\left\{ log\left(\frac{\text{CEO pay}_i - \text{CFO pay}_i}{\text{CFO pay}_i}\right) - \frac{1}{|J|}\sum_{j\in J} log\left(\frac{\text{CEO pay}_j - \text{CFO pay}_j}{\text{CFO pay}_j}\right)\right\}$$

4.3.4 Announcement Returns to TMT Departures and Mass Layoffs

For some of our tests, we examine how narcissism is related to announcement returns to two types of events: a non-CEO TMT member's departure and mass layoffs. We obtain announcement dates for TMT departures by searching firms' 8Ks for the TMT member's name during the year of the departure. We then read the 8K to verify that it announces the TMT member's departure. Upon such verification, we search Factiva for all stories in the press mentioning the executive's name in the month prior to the 8K filing to ensure there are no prior reports in the news of the departure. To ensure that the announcement returns are not driven by other events, we only include observations for which the 8K does not announce any additional information and for which there is no other news for the firm. For the purposes of searching for an announcement date for specific TMT departures, we randomly draw (without replacement) TMT departures identified through ExecuComp (using the procedure described above) until we obtain a sample of 100 uncontaminated announcement dates.[12] Randomly drawing and manually checking 462 8Ks netted 150 8Ks that unambiguously announced the departure of a TMT member. Of those 150 8Ks, 50 contained contaminating information not clearly related to the departure announcement, leaving a sample of 100 8Ks that announced the departure of a TMT member with no potentially confounding information. For the mass layoff announcement dates, we use the

[12]We lose one observation due to the unavailability of some control variables.

aforementioned dates from Ravenpack.

For each event, we compute the cumulative abnormal return over the three-day window centered on the event day based on the market model, with the parameters estimated over the prior 60 months and the CRSP value-weighted index as our proxy for the market portfolio.

4.3.5 Value of Human Capital and Unit Labor Costs

All else equal, the longer that a non-CEO executive has remained with the company, the more likely the executive is to have amassed a large amount of potentially valuable firm-specific human capital. We therefore use tenure with their current firm, as reported in ExecuComp, as a proxy for non-CEO TMT members' levels of human capital.

Data on the tenure of non-executive employees is not publicly available, so we cannot use tenure in our analysis of the mass layoffs. Instead, we appeal to microeconomic theory to argue that the wage an employee makes should, at a first order approximation, be equal to the marginal product of his labor. Hence, we can use wages as a proxy for the value of human capital. Average wages are not publicly available at the firm level, so we instead use the average real wage for the four-digit NAICS industry in the current year, as reported by the Bureau of Labor Statistics.

Since we relate average real wages to market reactions to layoffs, it is important that the firm's total labor costs do not confound our inferences. That is, a firm might initiate layoffs because its labor costs are high, and labor costs might be high because industry wages are high. To isolate the human capital value of labor from total labor costs, we also obtain from the Bureau of Labor Statistics unit labor costs for each industry, which is the total wage bill per dollar of output in the industry. We use this unit labor cost measure as a control variable in all regressions that use the average industry real wage as a test variable.

The distributions of both the average industry real wage and total industry real unit labor costs are highly skewed. We thus transform both these variables by adding a positive constant and taking the natural logarithm. In each case, we choose the constant term such that the resulting transformed variable has skewness as close to zero as possible.

4.3.6 Control Variables

Executive Controls. The past performance of a firm over an executive's tenure may influence that executive's turnover at the firm. If the value of a firm's stock declines over the executive's tenure, all else equal, one might expect an increased likelihood of management turnover. We therefore control for the firm's total buy-and-hold stock return over the executive's tenure at the firm. Executive turnover at narcissist-run firms could either be due to CEOs who terminate executives or to executives who choose to leave working under a narcissistic CEO. Hence, we control for executive age and sex (obtained from ExecuComp) and number of outside directorships (num. directorships, taken from BoardEx) to account for any differences in labor market opportunities that might contribute to an executive's decision to exit the company voluntarily. Since the departure of the most senior members of the TMT may have different implications for the firm, we include indicator variables for whether ExecuComp designates the executive as chief financial officer (CFO) or chief operating officer (COO). We control for total cash compensation as reported in ExecuComp because more highly compensated TMT members likely have a stronger incentive to stay with the firm. We control for whether the executive is also on the board (board member), since executives with board seats are plausibly more difficult to remove and/or more reluctant to leave. Finally, we control for the value of the executive's cumulative unvested compensation (pay on table) as reported in ExecuComp because having high unvested compensation would plausibly make an executive less likely to voluntarily leave.

Firm-level Controls. We control for several firm characteristics that are likely to be correlated with turnover and layoff decisions: firm size (log of total assets); return on assets or

ROA (the ratio current period net income to beginning-of-period total assets); the market-to-book ratio (the market value of equity plus total liabilities divided by total assets as of the end of the prior year); and contemporaneous sales growth. In addition, since changes in firm strategy may drive either TMT turnover or layoffs, we control for the absolute year-over-year change in advertising intensity and change in research and development intensity (where intensity is the ratio of the relevant expense to total sales), as well as the absolute year-over-year change in financial leverage (total debt to total assets) and number of acquisitions made by the firm in the prior year, as reported by Thomson One. Layoffs can be triggered not only by poor firm-specific performance, but also by industry-wide declines in profitability. Therefore, we control for industry ROA, defined as the median ROA for all firms in the same 4-digit NAICs code. In some alternative specifications, we control for firm and industry performance with the number of consecutive quarters in which either industry ROA or firm ROA has been negative up to and including the current quarter. All of the above variables, unless otherwise indicated, are constructed from Compustat data. Since the CEO's Optimism (alternatively called "overconfidence" in other studies) can plausibly impact layoffs, we control for it using the text-based measure of Aktas et al. (2016), which we construct using the same transcripts we use to construct our measure of CEO narcissism. Finally, since narcissists crave external accolades, we control for CEO Awards, as constructed by Boivie, Graffin, and Gentry (2016). It is defined as the number of times a CEO has been given special recognition by Forbes, BusinessWeek, Chief Executive Management, Industry Week, Institutional Investor, or Worth magazine over the past five years.

4.3.7 Descriptive Statistics

Descriptive statistics for our panel of non-CEO executive-firm-year observations, our sample of TMT turnover events, our firm-year sample, and our sample of mass layoffs are, respectively, in tables 4.1-4.6. All continuous variables are winsorized at the 1st and 99th percentiles to attempt to ensure outliers do not drive our results. In all panels, because it has been standardized, our narcissism measure has a mean close to zero and a standard deviation close to one. In tables 4.1 and 4.2, we see that the average non-CEO executive has been at the firm for around five years and has cash compensation around $500,000 per year. The large majority of TMT members have no outside directorships. Only 8% are women. The firm-level variables are largely in line with what has been reported in other studies of larger firms similar to those in our sample. In tables 4.3 and 4.4 we see that the average announcement return to a TMT departure is negative and economically large, 0.36%, but it is not statistically significant at conventional levels. There is also considerable cross-sectional variation in the CAR, as indicated by a standard deviation of 4.37 percentage points. In table 4.5 we see that layoffs are relatively rare, occurring in only 3.6% of firm-year observations in the panel. Finally, in table 4.6, we see that the mean stock price reaction to layoff news is 0.79%, and it is statistically different from zero (p=0.006). This result suggests that on average markets infer from layoff announcements bad news about future firm performance, as it seems implausible that layoffs themselves destroy value on average. Mean firm-level ROA is larger than the mean of our industry ROA measure because our sample consists of large firms, which tend to have higher-than-average profitability.

Table 4.1: Descriptive Statistics for TMT Turnover (Executive Level)

This table reports descriptive statistics for the executive level observations used in predicting top management team (TMT) turnover. All variables are described in the text. CEO Narcissism is a standardized and de-skewed text-based measure of narcissism that is derived from the ratio of first-person singular pronouns to first-person pronouns used in CEO speech; we standardize it in all regressions. Pay prominence is the closeness of TMT member pay to CEO pay, relative to industry norms based on the TMT member's role. Pay on table is the ratio of unvested compensation to total compensation. P25, P50, and P75 are, respectively, the 25^{th}, 50^{th}, and 75^{th} percentiles.

	Mean	Std. Dev	P25	P50	P75
CEO narcissism	0.0043	1.0035	-0.6189	0.0127	0.6469
TMT Tenure	5.1499	3.3487	3.0000	4.0000	7.0000
Pay prominence	-0.0111	0.8307	-0.2149	0.0000	0.1905
Pay on table	0.9859	1.9870	0.0000	0.2805	1.0637
Num directorships	0.0910	0.2504	0.0000	0.0000	0.0000
CFO	0.1070	0.3091	0.0000	0.0000	0.0000
COO	0.0130	0.1132	0.0000	0.0000	0.0000
Age	50.8585	6.9410	46.0000	51.0000	56.0000
Sex	0.0819	0.2742	0.0000	0.0000	0.0000
Cash compensation	493.4485	393.8862	292.9730	398.4750	556.6250
Board member	0.1130	0.3166	0.0000	0.0000	0.0000
Return over tenure	-0.0325	0.9446	-0.3059	-0.2047	-0.0098
CEO optimism	0.4716	0.3444	0.0000	0.5500	0.7500
CEO awards	0.1182	0.4920	0.0000	0.0000	0.0000
Observations	32347				

Table 4.2: Descriptive Statistics for TMT Turnover (Firm Level)
This table reports descriptive statistics for the firm level observations used in predicting top management team (TMT) turnover. All variables are described in the text. P25, P50, and P75 are, respectively, the 25^{th}, 50^{th}, and 75^{th} percentiles.

	Mean	Std. Dev	P25	P50	P75
ROA	0.0358	0.1092	0.0128	0.0454	0.0832
Ind. ROA	0.0004	0.0822	0.0009	0.0187	0.0371
Market to book	1.8023	1.0964	1.1097	1.4587	2.0752
log(total assets)	7.6342	1.6472	6.4624	7.5422	8.6769
Change in leverage	0.0454	0.0644	0.0055	0.0236	0.0581
Sales growth	0.0955	0.2419	-0.0132	0.0717	0.1672
Advertising intensity	0.0040	0.0087	0.0005	0.0014	0.0037
R&D intensity	0.0245	0.1143	0.0001	0.0023	0.0118
Num acquisitions	0.8323	1.6089	0.0000	0.0000	1.0000
Observations	32347				

Table 4.3: Descriptive Statistics for TMT CAR Regressions (Executive Level)

This panel reports descriptive statistics for the executive level observations used in predicting cumulative abnormal returns surrounding TMT departures; the observations constitute a random subsample from panel A. All variables are described in the text. CEO Narcissism is a standardized and de-skewed text-based measure of narcissism that is derived from the ratio of first-person singular pronouns to first-person pronouns used in CEO speech; we standardize it in all regressions. Pay prominence is the closeness of TMT member pay to CEO pay, relative to industry norms based on the TMT member's role. Pay on table is the ratio of unvested compensation to total compensation. P25, P50, and P75 are, respectively, the 25^{th}, 50^{th}, and 75^{th} percentiles.

	Mean	Std. Dev.	P25	P50	P75
CAR	-0.3605	4.3722	-2.3589	-0.1599	2.0438
CEO narcissism	0.0000	1.0000	-0.7134	-0.1116	0.5992
Tenure	3.0707	1.6241	2.0000	3.0000	4.0000
Pay prominence	1.9042	0.1074	1.8377	1.8859	1.9736
Pay on table	0.4110	1.5692	0.0000	0.0000	0.0097
Num directorships	0.0642	0.2233	0.0000	0.0000	0.0000
CFO	0.3434	0.4773	0.0000	0.0000	1.0000
COO	0.1212	0.3280	0.0000	0.0000	0.0000
Age	50.2727	7.9678	44.0000	50.0000	56.0000
Sex	0.0505	0.2201	0.0000	0.0000	0.0000
Cash compensation	262.1513	185.8729	142.6580	247.6360	321.0940
Return over tenure	0.3453	1.3583	-0.3577	0.0427	0.6032
CEO optimism	0.5181	0.3136	0.2759	0.6154	0.7692
CEO awards	0.0909	0.4064	0.0000	0.0000	0.0000
Observations	99				

Table 4.4: Descriptive Statistics for TMT CAR Regressions (Firm Level)

This panel reports descriptive statistics for the firm level observations used in predicting cumulative abnormal returns surrounding TMT departures; the observations constitute a random subsample from panel A. All variables are described in the text. CEO Narcissism is a standardized and de-skewed text-based measure of narcissism that is derived from the ratio of first-person singular pronouns to first-person pronouns used in CEO speech; we standardize it in all regressions. Pay prominence is the closeness of TMT member pay to CEO pay, relative to industry norms based on the TMT member's role. Pay on table is the ratio of unvested compensation to total compensation. P25, P50, and P75 are, respectively, the 25^{th}, 50^{th}, and 75^{th} percentiles.

	Mean	Std. Dev.	P25	P50	P75
ROA	0.0404	0.0736	0.0111	0.0376	0.0743
Ind. ROA	0.0237	0.0332	0.0090	0.0212	0.0379
Market to book	1.7082	1.0950	1.0706	1.2831	1.9547
log(total assets)	7.2677	1.7599	5.8301	7.0129	8.4433
Change in leverage	0.0392	0.0563	0.0027	0.0214	0.0538
Sales growth	0.0469	0.2856	-0.0918	0.0277	0.1739
Num acquisitions	0.6465	1.5539	0.0000	0.0000	0.0000
Advertising intensity	0.0022	0.0062	0.0000	0.0000	0.0009
R&D intensity	0.0053	0.0120	0.0000	0.0000	0.0023
Observations	99				

Table 4.5: Descriptive Statistics for Mass Layoff Regressions

This panel reports descriptive statistics for the observations used in predicting the announcement of a mass layoff at a firm. All variables are described in the text. CEO Narcissism is a standardized and de-skewed text-based measure of narcissism that is derived from the ratio of first-person singular pronouns to first-person pronouns used in CEO speech; we standardize it in all regressions. Real hourly wage is the average hourly wage for the firm's industry. The variable $\#(\text{ROA}< 0)$ is the consecutive number of quarters of negative return on assets sustained by the firm; $\#(\text{Ind. ROA} < 0)$ is the industry analog. Standard errors cluster on industry. P25, P50, and P75 are, respectively, the 25^{th}, 50^{th}, and 75^{th} percentiles.

	Mean	Std. Dev.	P25	P50	P75
Layoff annoucement	0.0358	0.1857	0.0000	0.0000	0.0000
CEO narcissism	-0.0000	1.0000	-0.5703	0.0076	0.6153
ROA	0.0111	0.0254	0.0034	0.0128	0.0235
Ind. ROA	0.0015	0.0207	0.0000	0.0058	0.0111
#(ROA < 0)	0.5499	1.6348	0.0000	0.0000	0.0000
#(Ind. ROA < 0)	4.1504	3.8426	0.0000	3.0000	9.0000
Real hourly wage	2.7126	0.4788	2.4167	2.7092	3.0094
CEO optimism	0.3738	0.3676	0.0000	0.3837	0.7027
CEO awards	0.1343	0.5743	0.0000	0.0000	0.0000
Market to book	2.9337	2.6551	1.4660	2.1890	3.4044
log(total assets)	7.5285	1.6030	6.3625	7.4429	8.5694
Change in leverage	0.0177	0.0298	0.0013	0.0075	0.0205
Sales growth	0.0333	0.1864	-0.0470	0.0231	0.0971
Num acquisitions	0.1374	0.4056	0.0000	0.0000	0.0000
Unit labor cost	-0.5932	0.1563	-0.7147	-0.5881	-0.4715
Observations	41565				

Table 4.6: Descriptive Statistics for Mass Layoff CAR
This panel reports descriptive statistics for the observations used in predicting the cumulative abnormal return surrounding a mass layoff announcement. All variables are described in the text. CEO Narcissism is a standardized and de-skewed text-based measure of narcissism that is derived from the ratio of first-person singular pronouns to first-person pronouns used in CEO speech; we standardize it in all regressions. Real hourly wage is the average hourly wage for the firm's industry. P25, P50, and P75 are, respectively, the 25^{th}, 50^{th}, and 75^{th} percentiles.

	Mean	Std. Dev.	P25	P50	P75
CAR	-0.7881	10.4562	-4.5888	-0.4194	3.6289
CEO narcissism	-0.0000	1.0000	-0.5346	0.0096	0.6540
ROA	-0.0021	0.0627	-0.0022	0.0079	0.0186
Ind. ROA	-0.0035	0.0244	-0.0045	0.0031	0.0083
CEO optimism	0.3700	0.3728	0.0000	0.3810	0.6923
CEO awards	0.3722	1.0899	0.0000	0.0000	0.0000
log(total assets)	8.1131	1.7705	6.7098	7.9063	9.4532
Market to book	2.7831	3.2416	1.2359	1.8821	3.0402
Change in leverage	0.0200	0.0327	0.0022	0.0090	0.0243
Sales growth	-0.0065	0.1716	-0.0868	-0.0046	0.0643
Real hourly wage	2.7781	0.4014	2.5444	2.7822	2.9985
Num acquisitions	0.1729	0.4871	0.0000	0.0000	0.0000
Unit labor cost	0.2105	0.0821	0.1399	0.2182	0.2717
Observations	1064				

4.4 Tests and Results

Testing the hypotheses discussed in section 4.2 require examining the relation between narcissism and the following: the hazard rate of top management turnover, the stock price reaction to news of departures of non-CEO top managers, the odds of a mass layoffs, and the stock price reaction to news of mass layoffs. Below we discuss results for each set of tests.

4.4.1 Hazard Rate of TMT Turnover

To test our first hypothesis on the relation between CEO narcissism and executive turnover (Hypothesis 1a), we estimate a Cox proportional hazard model that predicts the executive's departure as a function of the CEO's narcissism and many control variables:

$$h(i,t) = h_0(t)\text{exp}(\beta_1 \text{CEO Narcissism}_{i,t} + \gamma' \text{Controls}_{i,t}) \tag{4.1}$$

where $h(i,t)$ is the departure hazard of executive i who has been part of the TMT for t years and $h_0(t)$ is the baseline hazard rate corresponding to t. Controls is the vector of all the firm, CEO and executive-level control variables discussed in the prior section. We estimate the above equation using partial likelihood with standard errors clustered by CEO. Hence, our results are robust to arbitrary error correlation across executive-year observations with the same CEO in charge. Instead of coefficients, we report the exponential of each parameter estimate, which can be interpreted as proportional hazard ratios. Thus, if a variable increases (decreases) the likelihood of a TMT member departure, the hazard ratio will be significantly greater (less) than 1.0. Recall that Hypothesis 1a predicts that CEO narcissism is associated with an increase in the hazard of TMT turnover, implying that the hazard ratio for CEO narcissism should be greater than 1.0.

Results of the hazard regressions are in table 4.7. Consistent with hypothesis 1a, the coefficient on CEO Narcissism is statistically significant across the specifications and its exponential is greater than 1.0. Using the most complete specification in column (4), the proportional hazard ratio is 1.066. Since narcissism has been standardized, the point estimate implies that increasing CEO narcissism from one standard deviation below the mean to one standard deviation above increases the departure hazard rate by a factor of 1.136 $[1.136 = 1.0662^2]$, or 13.6%.

We also predict that the effect of narcissism on TMT turnover will be larger when the value of a non-CEO executive is perceived and recognized more highly by the firm (hypothesis 1b). As discussed in section 4.3, we use the executive's pay prominence as our measure of such perceived value. We test hypothesis 1b by estimating another Cox hazard model identical to equation (4.1) above, except that we also include an interaction of CEO Narcissism with pay prominence.

The results of our interaction specification are in column (2) of table 4.7. Supporting hypothesis 1b, the results show that higher pay prominence amplifies the effect of CEO narcissism on the likelihood a TMT member's turnover, and the effect is statistically significant at the 1% level. In addition, the direct effect of narcissism remains significant in this specification, with a p-value of 0.009 and a hazard ratio point estimate of 1.055. To ease the economic interpretation of the interaction term coefficients, all interacted continuous variables have been standardized in all specifications. Thus the product of the exponentials of the coefficients on the direct effect and the interaction term imply that, when an executive has pay prominence one standard deviation above the mean, the marginal effect of a standard deviation increase in narcissism is to increase the hazard of the executive's turnover by a factor of 1.15 $[1.15 = 1.055 * 1.088]$. This effect is large in an absolute sense, as well as relative to 1.055, the effect of a one standard deviation increase in CEO narcissism on the turnover hazard when the executive's pay prominence is at the mean.

News reports and 8K filings that disclose the departure of TMT members do not generally provide sufficient information to tell whether a departure is voluntary or involuntary (i.e., a termination). Nevertheless, to try to get a sense of whether CEO narcissism-induced TMT turnover is voluntary for a significant part of the sample, we examine how the relation between narcissism and turnover is related to pay on table. We define this variable

Table 4.7: TMT Turnover Prediction

This table shows the estimation results to equation (4.1) in the paper. The likelihood of turnover for a non-CEO top management team (TMT) member in a given year is modeled via a Cox proportional hazards function, where the dependent variable takes value 1 if the TMT Member leaves the firm and 0 otherwise. The baseline hazard of turnover is a function of the tenure of the TMT member; the covariates shown below increase or decrease the turnover risk. All variables are described in the text. A hazard ratio (exponentiated beta coefficient) over 1 indicates increased turnover likelihood, a hazard ratio below 1 indicates decreased likelihood. CEO Narcissism is a standardized and de-skewed text-based measure of narcissism that is derived from the ratio of first-person singular pronouns to first-person pronouns used in CEO speech (conference calls and interviews). Pay prominence is the closeness of TMT member pay to CEO pay, relative to industry norms based on the TMT member's role. Pay on table is the ratio of unvested compensation to total compensation. Standard errors cluster on the CEO, p-values are reported below the hazard ratios. The labels *, **, and *** denote p-values less than .01, .05, and .01, respectively. Coefficients on control variables are not reported for brevity.

CEO narcissism	1.042**	1.051**	1.062***	1.038*	1.066***
	(0.043)	(0.014)	(0.007)	(0.081)	(0.005)
... x Pay prominence		1.074***			1.072***
		(0.008)			(0.009)
... x Pay on table			0.981**		0.982**
			(0.040)		(0.041)
... x ROA				1.101	1.122
				(0.580)	(0.502)
... x Ind. ROA				0.971	0.955
				(0.896)	(0.842)
Pay prominence	0.922***	0.905***	0.922***	0.922***	0.905***
	(0.001)	(0.000)	(0.001)	(0.001)	(0.000)
Pay on table	1.027	1.027	1.028	1.027	1.027
	(0.255)	(0.251)	(0.234)	(0.258)	(0.242)
ROA	1.023	1.020	1.023	1.022	1.020
	(0.299)	(0.348)	(0.290)	(0.311)	(0.358)
Ind. ROA	1.007	1.007	1.006	1.007	1.006
	(0.781)	(0.787)	(0.802)	(0.774)	(0.800)
Controls	Yes	Yes	Yes	Yes	Yes
Pseudo R-squared	0.038	0.038	0.038	0.038	0.038
N	32347	32347	32347	32347	32347
Num executives	11480	11480	11480	11480	11480
Num turnovers	3563	3563	3563	3563	3563

as the dollar value of unvested compensation that a TMT member would forgo should he leave.[13] If a significant number of TMT turnovers induced by CEO narcissism are voluntary, we would expect the response of TMT turnover to CEO narcissism to be weaker when pay on table is larger. We thus run another specification, identical to equation 4.1, except we include the standardized pay on table and its interaction with CEO Narcissism as independent variables. Our point estimate of coefficient on the interaction implies a significant negative relation (p-value = 0.065). The exponential of this coefficient is 0.981, whereas the exponential of the direct effect of CEO narcissism is 1.064. These results imply that when pay left on the table is at its mean, a standard deviation increase in CEO narcissism increases the departure hazard rate by a factor of 1.064. However, when pay on table increases by a standard deviation, the effect of narcissism on the departure hazard declines to 1.04 $[1.04 = 1.064 * 0.981]$. Since pay left on the table significantly alters the effect of CEO Narcissism on turnover hazard, we conclude that a non-trivial number of CEO narcissism-induced TMT departures are voluntary. We caution the reader, however, that we cannot infer from this result whether the majority of narcissism-induced TMT departures are voluntary.

We also examine whether the effect of narcissism on the turnover hazard changes with firm and industry performance. We thus run a specification identical to equation (4.1), except that we include interactions of narcissism with both industry and firm ROA. The interaction terms are statistically indistinguishable from zero. In a specification that includes all of the interaction terms used in prior specifications (including with pay prominence and pay on table), all of our inferences remain unchanged, and the stand-alone effect of narcissism is largest in magnitude.

[13] We say could forgo because we cannot observe severance agreements for TMT members.

4.4.2 Stock Price Reaction to TMT Departures

We next turn our attention to stock market reactions to news of TMT turnover events. Recall that we predict that the effect of narcissism on the market's reaction to the news of TMT turnover depends on the value of the executive's firm-specific human capital, for which we use the executive's tenure as a proxy. We thus estimate the following equation using weighted least squares for the sample of turnover events:

$$\text{CAR}_i = \alpha + \beta_1 \text{CEO Narcissism}_i + \beta_2 \text{CEO Narcissism}_i \times \text{Tenure}_i + \gamma' \text{Controls}_i + \epsilon_i \quad (4.2)$$

Where CAR_i the cumulative abnormal return over the three-day window around event i and controls include the same variables as in the hazard regressions. The weight for each observation is the inverse of the standard forecast error associated with each cumulative abnormal return from a market model. The weighting scheme assigns more weight to observations where the CAR is estimated more precisely, and should adjust the regression residuals toward being homoscedastic. Nevertheless, we compute White robust standard errors.

Our estimates of parameters and p-values from equation (4.2) are in table 3. We include specifications with one interaction term at a time (in columns 2-5), as well as a specification that includes all interactions at the same time (column 6). In the most complete specification (column 6), the coefficient on narcissism is significantly positive (p-value = 0.015) and economically large at 1.984. This result implies that news of a non-CEO executive departure produces a CAR of 1.984 percentage points larger when CEO narcissism increases one standard deviation above the mean. In a Cox hazard model, the time since the executive began their position is absorbed in the baseline hazard; there is no equivalent feature in the OLS regressions that we use for the cumulative abnormal returns. We thus interact CEO narcissism with the tenure of the TMT member to assess whether the effects differ as a function of how long the TMT member has been in his position. The coefficient on the interaction term is significantly negative (p-value = 0.009). The point estimate of 0.797 implies that, for every additional year the departing TMT member has been with the firm, the effect of CEO narcissism on the CAR is less positive by 0.797 percentage points. The two coefficient estimates taken together imply that the effect of a one standard deviation increase in CEO narcissism on the announcement return becomes negative once a TMT member has been with the firm approximately 2.5 years or more. These results support Hypothesis 2: the effect of narcissism-induced TMT turnover on firm value is negative when human capital levels are high. To ensure that normal retirements are not driving our results, we run a similar specification wherein we drop all observations where TMT members are older than 62. The results (untabulated) are qualitatively similar to those reported.

Although the positive coefficient on the direct effect of narcissism on TMT departure announcement returns is consistent with CEO narcissism having a positive effect on firm value when it induces the turnover of TMT members with short tenure, it is also consistent with an alternative hypothesis. It is well-established that narcissists do not work well with others, so under a narcissistic CEO, TMT departures are potentially more likely to be driven by interpersonal conflicts. If investors are aware of such conflicts and if the conflicts adversely impact firm performance, the market could greet as good news the departure of a TMT member involved in such a conflict with the CEO, even if the conflict were totally the fault of the narcissistic CEO. If the CEO's narcissism caused the conflict in the first place, we should not conclude that the resulting positive association of CEO narcissism with turnover announcement returns indicates that CEO narcissism creates value by inducing such turnover. Hence, we cannot infer from the positive direct effect that CEO narcissism increases firm value by making the departure of short-tenure TMT members more likely.

Table 4.8: TMT Announcement Returns

This table shows the estimation results to equation (4.2) in the paper using weighted least squares. The dependent variable is the cumulative abnormal return for the firm over a three day window. CEO Narcissism is a standardized and de-skewed text-based measure of narcissism that is derived from the ratio of first-person singular pronouns to first-person pronouns used in CEO speech (conference calls and interviews). Pay prominence is the closeness of TMT member pay to CEO pay, relative to industry norms based on the TMT member's role. Pay on table is the ratio of unvested compensation to total compensation. Standard errors are robust to heteroskedasticity, p-values are reported below the coefficients. The labels *, **, and *** denote p-values less than .01, .05, and .01, respectively. Coefficients on control variables are not reported for brevity.

CEO narcissism	-0.351	1.717**	-0.228	-0.350	-0.327	1.984**
	(0.414)	(0.040)	(0.500)	(0.434)	(0.451)	(0.015)
... x Tenure		-0.722**				-0.797***
		(0.018)				(0.009)
... x Pay prominence			-0.439			-0.398
			(0.250)			(0.337)
... x Pay on table				-0.013		0.245
				(0.973)		(0.623)
... x ROA					-0.523	-0.685
					(0.246)	(0.129)
... x Ind. ROA					-0.081	0.164
					(0.833)	(0.663)
Tenure	-0.090	0.006	-0.066	-0.090	-0.085	0.027
	(0.650)	(0.974)	(0.727)	(0.654)	(0.667)	(0.887)
Pay prominence	-0.204	-0.363	-0.167	-0.203	-0.251	-0.401
	(0.642)	(0.404)	(0.708)	(0.651)	(0.570)	(0.389)
Pay on table	0.444	0.674	0.561	0.451	0.486	0.751
	(0.382)	(0.206)	(0.303)	(0.453)	(0.341)	(0.285)
ROA	-0.171	-0.572	-0.142	-0.167	0.131	-0.203
	(0.772)	(0.348)	(0.808)	(0.786)	(0.842)	(0.778)
Ind. ROA	-0.539	-0.363	-0.489	-0.539	-0.570	-0.380
	(0.154)	(0.329)	(0.195)	(0.158)	(0.169)	(0.347)
Controls	Yes	Yes	Yes	Yes	Yes	Yes

In contrast, the interpretation of the negative coefficient on the interaction term is clearer. Even if CEO narcissism causes TMT departures primarily by inducing interpersonal conflicts of which the market is aware prior to the departure, the fact that the market still reacts negatively to such departures when executives have longer tenure indeed indicates that the loss of the executive is a net negative for the firm. Therefore, although we cannot infer whether narcissism-induced departures of short-tenure TMT members are a net positive or negative for a firm, we can infer that narcissism-induced departures of long-tenure executives are unambiguously a net negative.

4.4.3 Odds of Mass Layoffs

Recall that our hypothesis 3 states that because narcissistic CEOs are less empathetic than other CEOs, they are more likely to initiate layoffs, and those layoffs are less dependent on poor performance. We thus model the probability of a layoff in any firm-quarter observation as a function of CEO narcissism and firm performance:

$$\begin{aligned}\Pr(\text{Layoff}_{i,t}) = F(&\alpha + \beta_1 \text{CEO Narcissism}_{i,t} + \beta_2 \text{ROA}_{i,t} + \beta_3 \text{Industry ROA}_{i,t} \\ &+ \gamma' \text{Controls}_{i,t}) + \epsilon_{i,t} \end{aligned} \tag{4.3}$$

$$\begin{aligned}\Pr(\text{Layoff}_{i,t}) = F(&\alpha + \beta_1 \text{CEO Narcissism}_{i,t} + \beta_2 \text{ROA}_{i,t} + \beta_3 \text{Industry ROA}_{i,t} \\ &+ \beta_4 \text{CEO Narcissism}_{i,t} \times \text{ROA}_{i,t} \\ &+ \beta_5 \text{CEO Narcissism}_{i,t} \times \text{Industry ROA}_{i,t} \\ &+ \gamma' \text{Controls}_{i,t}) + \epsilon_{i,t} \end{aligned} \tag{4.4}$$

Where $\text{ROA}_{i,t}$ is the standardized return on assets for firm i in quarter t, as defined above, and Industry $\text{ROA}_{i,t}$ is the median of the same variable across all firms within the same 4-digit NAICS code in the quarter. We include both firm and industry ROA because poor performance at either the firm or industry level could plausibly trigger layoffs. The vector Controls now includes all of the firm-level controls used in the TMT regressions, as well as the average hourly real wage and unit labor costs for the industry, which, respectively, serve as proxies for the value of human capital and total labor costs in the firm. We estimate the above equations with a logistic regression. Note from table 4.5 that Layoff takes a value of zero in more than 97% of the observations. Thus, we also estimate a complementary log-log regression to ensure that our results are not driven by the finite sample bias of logit models that might arise due to the sparseness of our dichotomous dependent variable. We estimate robust standard errors clustered by industry, so our results are robust to arbitrary heteroscedasticity, serial correlation, and any correlation between error terms of firms within the same industry across all quarters.

Our parameter estimates and p-values from equations (4.3) and (4.4) are in the first four columns of table 4.9. The coefficient on CEO Narcissism is statistically insignificant in the specifications without interactions, but it is positive and significant at the 5% level in the specifications with the interaction terms. We note that the coefficients and standard errors are qualitatively very close in the logit and complementary log-log specifications, so we conclude that the finite sample bias of logit regressions is not a material concern. Henceforth, we focus on the logit parameter estimates because they are easier to interpret. We also include specifications where narcissism is interacted with the average industry real wage (our proxy for human capital value), as well as specifications where all interactions are included simultaneously.

Table 4.9: Layoff Prediction

This table shows the estimation results to equation (4.4) in the paper. The dependent variable is an indicator that takes value one when the firm announces a layoff during the fiscal quarter and zero otherwise. The dichotomous outcome is estimated via maximum likelihood with a complementary log-log link function in the odd columns and logit link function in the even columns. CEO Narcissism is a standardized and de-skewed text-based measure of narcissism that is derived from the ratio of first-person singular pronouns to first-person pronouns used in CEO speech (conference calls and interviews). Real hourly wage is the average hourly wage for the firm's industry. Exponentiated beta coefficients are reported (in the case of logit, these are odds ratios) with p-values in parentheses. Standard errors cluster on industry. Regressions also control for year and fiscal quarter. The labels *, **, and *** denote p-values less than .01, .05, and .01, respectively. Coefficients on control variables are not reported for brevity.

	Complementary Log-Log	Logit	Complementary Log-Log	Logit	Complementary Log-Log	Logit	Complementary Log-Log	Logit
CEO narcissism	1.034	1.037	1.068**	1.070**	1.036	1.039	1.072**	1.074**
	(0.293)	(0.274)	(0.041)	(0.042)	(0.273)	(0.255)	(0.032)	(0.032)
... x ROA			1.049*	1.050*			1.050*	1.051*
			(0.064)	(0.080)			(0.060)	(0.077)
... x Ind. ROA			1.051**	1.054**			1.056**	1.059**
			(0.031)	(0.029)			(0.027)	(0.025)
... x Real hourly wage					1.029	1.028	1.044*	1.043*
					(0.207)	(0.241)	(0.052)	(0.067)
ROA	0.759***	0.746***	0.755***	0.743***	0.758***	0.745***	0.754***	0.742***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Ind. ROA	0.863***	0.857***	0.857***	0.851***	0.863***	0.857***	0.856***	0.851***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Real hourly wage	0.953	0.947	0.949	0.944	0.952	0.947	0.947	0.943
	(0.529)	(0.497)	(0.495)	(0.468)	(0.524)	(0.494)	(0.481)	(0.457)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	41565	41565	41565	41565	41565	41565	41565	41565

Focusing on the parameter estimates and standard errors of the logit specifications where narcissism is interacted with firm and industry ROA, the direct effect of narcissism is significant in the expected direction in all of them. Hence, narcissism does indeed increase the odds of a layoff. In addition, the exponentials of the interaction term coefficients are both greater than one, and using a Wald test we find they are jointly significantly different from 1.0 at the 5% level or better. The latter finding implies that narcissists are willing to initiate layoffs at higher levels of both industry and firm profitability, or stated differently, that narcissists initiate layoffs following smaller declines in profitability than do non-narcissists.

To get a sense of the economic magnitude of the effect of narcissism on layoff odds, note that in the most complete logit specification (in the last column in table 4.9), the exponential of the coefficient on narcissism is 1.074. This result implies that, for a firm with average industry and firm-specific ROA, a change in CEO Narcissism from one standard deviation below the mean to one standard deviation above increases the odds of a layoff by a factor of 1.153 $[1.153 = 1.0742^2]$, or 15.3%. Now consider the economic significance of the interaction effect. The exponential of the direct effect of firm ROA and industry ROA are 0.743 and 0.851, respectively, whereas the coefficient on the interaction terms are 1.050 for firm ROA and 1.054 for industry ROA. Taken together, these point estimates imply that a one-standard-deviation increase in firm ROA and industry ROA at a firm whose CEO has average narcissism reduces the odds of layoffs by a factor of 0.631 $[0.742 \times 0.851 = 0.631]$. However, the same increase in both firm and industry profitability decreases the odds of layoffs for a narcissistic CEO by a factor of only 0.702.

To determine whether narcissistic CEOs are quicker to initiate layoffs in response to poor performance, we run the following specification using both logit and complementary log-log regressions:

$$
\begin{aligned}
\Pr(\text{Layoff}_{i,t}) = F(&\alpha + \beta_1 \text{CEO Narcissism}_{i,t} \\
&+ \beta_2 N(\text{ROA}_{i,t} < 0) + \beta_3 N(\text{Industry ROA}_{i,t} < 0) \\
&+ \beta_4 \text{CEO Narcissism}_{i,t} \times N(\text{ROA}_{i,t} < 0) \\
&+ \beta_5 \text{CEO Narcissism}_{i,t} \times N(\text{Industry ROA}_{i,t} < 0) \\
&+ \gamma' \text{Controls}_{i,t}) + \epsilon_{i,t}
\end{aligned}
\tag{4.5}
$$

Where $N(\text{ROA} < 0)_{i,t}$ and $N(\text{Industry ROA} < 0)_{i,t}$ are the numbers of consecutive quarters with negative ROA that the firm or industry, respectively, have experienced up to and including quarter t. If narcissistic CEOs are faster to initiate layoffs in response to losses, we expect the exponential of the coefficient on the interaction term to be less than one. As before, we cluster standard errors by industry. The results are in table 4.10, and they show that indeed, the exponentials of the interaction term coefficient point estimates are both less than one. Using a Wald test, we also find they are jointly significantly different from 1.0 at the 5% level or better. We conclude that narcissistic CEOs do not require as many quarters of poor performance before deciding to initiate layoffs.

Table 4.10: Layoff Prediction (Alternative Performance Measure)

This table shows the estimation results to equation (4.5) in the paper. The dependent variable is an indicator that takes value one when the firm announces a layoff during the fiscal quarter and zero otherwise. The dichotomous outcome is estimated via maximum likelihood with a complementary log-log link function in the odd columns and logit link function in the even columns. CEO Narcissism is a standardized and de-skewed text-based measure of narcissism that is derived from the ratio of first-person singular pronouns to first-person pronouns used in CEO speech (conference calls and interviews). Real hourly wage is the average hourly wage for the firm's industry. Exponentiated beta coefficients are reported (in the case of logit, these are odds ratios) with p-values in parentheses. The variable $\#$(ROA<0) is the consecutive number of quarters of negative return on assets sustained by the firm; $\#$(Ind. ROA <0) is the industry analog. Standard errors cluster on industry. Regressions also control for year and fiscal quarter. The labels *, **, and *** denote p-values less than $.01$, $.05$, and $.01$, respectively. Coefficients on control variables are not reported for brevity.

	Complementary Log-Log	Logit	Complementary Log-Log	Logit	Complementary Log-Log	Logit	Complementary Log-Log	Logit
CEO narcissism	1.035	1.037	1.115**	1.120**	1.036	1.039	1.128**	1.134**
	(0.288)	(0.280)	(0.044)	(0.040)	(0.264)	(0.256)	(0.028)	(0.026)
... x #(ROA <0)			0.983	0.982			0.983	0.982
			(0.177)	(0.183)			(0.189)	(0.195)
... x #(Ind. ROA <0)			0.988	0.988			0.986*	0.986*
			(0.110)	(0.104)			(0.076)	(0.073)
... x Real hourly wage					1.029	1.029	1.042*	1.042*
					(0.235)	(0.251)	(0.081)	(0.089)
#(ROA <0)	1.148***	1.156***	1.154***	1.161***	1.149***	1.156***	1.154***	1.161***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
#(Ind. ROA <0)	1.038**	1.039**	1.040**	1.041**	1.038**	1.040**	1.040**	1.042***
	(0.015)	(0.014)	(0.012)	(0.011)	(0.015)	(0.013)	(0.010)	(0.009)
Real hourly wage	0.948	0.945	0.945	0.941	0.948	0.945	0.943	0.940
	(0.491)	(0.478)	(0.462)	(0.448)	(0.484)	(0.473)	(0.444)	(0.434)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	41565	41565	41565	41565	41565	41565	41565	41565

4.4.4 Stock Price Reaction to News of Mass Layoffs

Finally, we turn attention to our analysis of cumulative abnormal returns (CARs) around layoff announcement dates. We use OLS to estimate two regressions, one with and one without interaction terms:

$$\text{CAR}_{i,t} = \alpha + \beta_1 \text{CEO Narcissism}_{i,t} + \gamma' \text{Controls}_{i,t} + \epsilon_{i,t} \tag{4.6}$$

$$\begin{aligned}\text{CAR}_{i,t} = \alpha &+ \beta_1 \text{CEO Narcissism}_{i,t} + \beta_2 \text{CEO Narcissism}_{i,t} \times \text{Real Hourly Wage}_{i,t} \\ &+ \gamma' \text{Controls}_{i,t} + \epsilon_{i,t} \end{aligned} \tag{4.7}$$

where Controls includes the same variables as in equations (4.3) and (4.4), including the average real hourly wage for the industry. These regressions are presented in table 4.11. In the specification without the interaction term, we see that layoffs initiated by narcissists have announcement CARs statistically indistinguishable from that of layoffs announced by a non-narcissist. In the interaction specification, however, we observe that the coefficient on the interaction term is negative and significant at the 5% level. This result implies that how the market receives news of layoffs initiated by narcissistic vs. non-narcissistic CEOs depends on the level of the real wage in the industry, our proxy for the average value of the human capital in the firm. We illustrate the economic significance of the effect with figure 4.1. Note that when the real hourly wage is two standard deviations above the mean, a standard deviation increase in the CEO's narcissism will tend to decrease the CAR by around 0.75 percentage points. We interpret our evidence as supporting hypothesis 4 – the higher propensity of narcissists to initiate layoffs has a negative effect on firm value when the value of human capital is high, consistent with the view that narcissistic CEOs can destroy firm value through layoffs by underestimating the value of the resulting loss of human capital.

Table 4.11: Layoff Announcement Returns

This table shows the estimation results to equation (4.7) in the paper. The dependent variable is the cumulative abnormal return for the firm over a three day window. CEO Narcissism is a standardized and de-skewed text-based measure of narcissism that is derived from the ratio of first-person singular pronouns to first-person pronouns used in CEO speech (conference calls and interviews). Real hourly wage is the average hourly wage for the firm's industry. Standard errors are clustered on the industry, p-values are reported below the coefficients. The labels *, **, and *** denote p-values less than .01, .05, and .01, respectively. Coefficients on control variables are not reported for brevity.

CEO narcissism	0.113	0.010	0.194	0.102
	(0.581)	(0.962)	(0.343)	(0.628)
... x ROA		0.539		0.454
		(0.179)		(0.247)
... x Ind. ROA		0.026		-0.021
		(0.803)		(0.821)
... x Real hourly wage			-0.486**	-0.469**
			(0.024)	(0.021)
ROA	0.149	0.118	0.164	0.135
	(0.766)	(0.791)	(0.742)	(0.765)
Ind. ROA	0.414***	0.416***	0.400***	0.408***
	(0.000)	(0.000)	(0.000)	(0.000)
Real hourly wage	0.201	0.199	0.254*	0.250*
	(0.127)	(0.136)	(0.080)	(0.086)
Controls	Yes	Yes	Yes	Yes
N	1064	1064	1064	1064

4.5 Conclusion

A large literature in finance demonstrates that psychological characteristics of CEOs have important effects on firm value. There has been little large sample work on the psy-

Figure 4.1: Marginal Effect of CEO Narcissism on Layoff Announcement Return

This figure plots the predicted effect (using the results from table 6) of CEO Narcissism on the cumulative abnormal return for a layoff announcement over the range of (standardized) real hourly wage. 90% confidence intervals are included around the predicted effect.

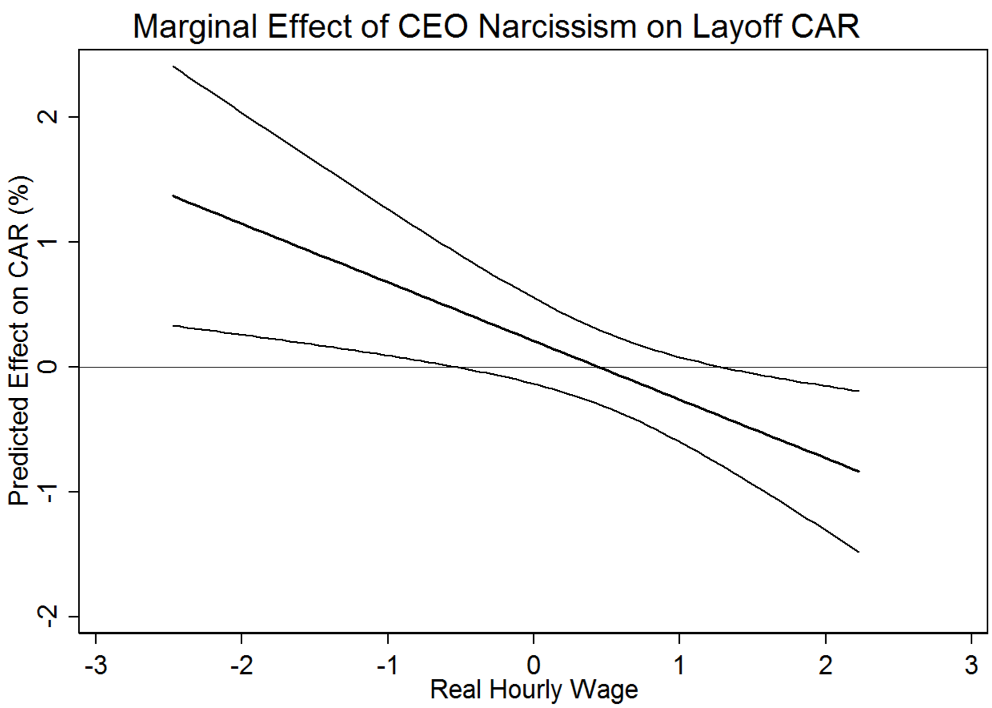

chological trait of narcissism, however, which prior research has demonstrated to be particularly common among individuals in high power positions, such as CEOs. In addition, there has been little work in finance examining how a CEO's management of human capital can impact firm value. Since narcissism tends to strongly impact how individuals interact with and view others, the degree to which CEO narcissism affects firm value through the human capital channel is a particularly salient research question to ask. Our results suggest that CEO narcissism matters. Narcissistic CEOs can potentially benefit shareholders in that their lack of empathy makes them more likely to terminate executives and workers when appropriate. Their tendency to underestimate the worth of others, however, leads them to underestimate the human capital cost of terminations and voluntary departures of others. Our results imply that CEO narcissism tends to destroy firm value through the turnover of other executives and through mass layoffs when the value of human capital is high.

5. SUMMARY

Economic theory suggests that these inter-firm links lead to behavior different from simplified settings in which firms operate in isolation without, for example, product market rivals and shared human capital. I empirically show that, in a multipartite economy, linkages between and within firms produce substantial cross-sectional variation in firm outcomes. In section 1, firm-specific breaches of loan covenants were shown to have implications for optimal investment policy of firms with similar default risk. In particular, lender-induced conservatism following covenant violation created an opportunity for predation by rival firms. In section 2, unanticipated data breaches led to shocks in a director's understanding of cybersecurity risk, and this resulted in changes in cybersecurity monitoring at the director's other firms. In section 3, a CEO's narcissistic personality affected his valuation of human capital at the firm. Depending on the circumstance, this created suboptimal levels of top management team turnover and employee layoffs.

REFERENCES

Aguilar, L. (2014). *Boards of Directors, Corporate Governance and Cyber-Risks: Sharpening the Focus*. URL: `https://www.sec.gov/news/speech/2014-spch061014laa` (visited on 05/12/2017).

Aktas, N., E. DeBodt, H. Bollaert, and R. Roll (2016). "CEO Narcissism and the Takeover Process: From Private Initiation to Deal Completion". *Journal of Financial and Quantitative Analysis* 51, pp. 113–137.

Atanasov, V. and B. S. Black (2016). "Shock-Based Causal Inference in Corporate Finance Research". *Critical Finance Review* 5, pp. 207–304.

Banerjee, S., M. Humphery-Jenner, and V. Nanda (2014). "Restraining Overconfident CEOs Through Improved Governance: Evidence from the Sarbanes-Oxley Act". *Review of Financial Studies* 28.10, pp. 2812–2858.

Banerjee, S., L. Dai, M. Humphery-Jenner, and V. Nanda (2015). "Top Dogs: Overconfident Executives and New CEO Selection". *SSRN Working Paper* 2371435.

Benmelech, E. and C. Frydman (2015). "Military CEOs". *Journal of Financial Economics* 117, pp. 43–59.

Billett, M. and Y. Qian (2008). "Are Overconfident CEOs Born or Made? Evidence of Self-Attribution Bias from Frequent Acquirers". *Management Science* 54.6, pp. 1037–1051.

Billett, M. T., B. Esmer, and M. Yu (2016). "Creditor Control and Product Market Competition". *SSRN Working Paper* 2307031.

Bjorhus, J. (2014). *Clean Reviews Preceded Target's Data Breach, and Others*. URL: `http://www.govtech.com/security/Clean-Reviews-Preceded-Targets-Data-Breach-and-Others.html` (visited on 05/12/2017).

Boehmer, E., C. M. Jones, and X. Zhang (2015). “Potential Pilot Problems: Treatment Spillovers in Financial Regulatory Experiments”. *SSRN Working Paper* 2621598.

Boivie, S., S. D. Graffin, and R. Gentry (2016). “Understanding the Direction, Magnitude, and Joint Effects of Reputation When Multiple Actors’ Reputations Collide”. *Academy of Management Journal* 59.1, pp. 188–206.

Bolton, P. and D. Scharfstein (1990). “A Theory of Predation Based on Agency Problems in Financial Contracting”. *American Economic Review* 80.1, pp. 93–106.

Bouwman, C. (2014). “Managerial Optimism and Earnings Smoothing”. *Journal of Banking and Finance* 41.April, pp. 283–303.

Bushman, B. and S. Thomaes (2011). “When the Narcissistic Ego Deflates, Narcissistic Aggression Inflates”. *The Handbook of Narcissisn and Narcissistic Personality Disorder*. Ed. by K. W. Campbell and J. D. Miller. Hoboken, New Jersey: John Wiley & Sons, pp. 319–329.

Campbell, W. K. (1999). “Narcissism and Romantic Attraction”. *Journal of Personality and Social Psychology* 77.6, pp. 1254–1270.

Carlson, E. N., L. P. Naumann, and S. Vazire (2011). “Getting to Know a Narcissist Inside and Out”. *The Handbook of Narcissism and Narcissistic Personality Disorder*. Ed. by K. W. Campbell and J. D. Miller. Hoboken, New Jersey: John Wiley & Sons, pp. 285–299.

Carlson, E. N., S. Vazire, and O. Tf (2011). “You Probably Think this Paper’s About You: Narcissists’ Perceptions of Their Personality and Reputation”. *Journal of Personality and Social Psychology* 101.1, pp. 185–201.

Cerqueti, R., F. Fiordelisi, and P. R. Rau (2015). “Corporate Culture and Enforcement Actions in Banking”. *AEA Conference Paper* 2016.

Cerulli, G. (2015). “Identification and Estimation of Treatment Effects in the Presence of Neighbourhood Interactions”. *Working paper*.

Chatterjee, A. and D. Hambrick (2007). "It's All About Me: Narcissistic Chief Executive Officers and Their Effects on Company Strategy and Performance". *Administrative Science Quarterly* 52.3, pp. 351–386.

Chava, S. and M. R. Roberts (2008). "How Does Financing Impact Investment? The Role of Debt Covenants". *Journal of Finance* 63.5, pp. 2085–2121.

Chi, J. D. and X. Su (2016). "Product Market Threats and the Value of Corporate Cash Holdings". *Financial Management* 45.3, pp. 705–735.

DeAngelo, H. and L. DeAngelo (2007). "Capital Structure, Payout Policy, and Financial Flexibility". *SSRN Working Paper* 916093.

Demerjian, P. R. and E. L. Owens (2014). "Measuring the Probability of Financial Covenant Violation in Private Debt Contracts". *Journal of Accounting and Economics* 61.2-3, pp. 433–447.

Denis, D. J. (2011). "Financial Flexibility and Corporate Liquidity". *Journal of Corporate Finance* 17.3, pp. 667–674.

Denis, D. J. and S. B. McKeon (2012). "Debt Financing and Financial Flexibility Evidence from Proactive Leverage Increases". *Review of Financial Studies* 25.6, pp. 1897–1929.

Deshmukh, S., A. M. Goel, and K. M. Howe (2013). "CEO Overconfidence and Dividend Policy". *Journal of Financial Intermediation* 22.3, pp. 440–463.

Dessaint, O. and A. Matray (2015). "Do Managers Overreact to Salient Risks? Evidence from Hurricane Strikes". *SSRN Working Paper* 2358186.

Emmons, R. A. (1987). "Narcissism: Theory and Measurement". *Journal of Personality and Social Psychology* 52.1, pp. 11–17.

Engelberg, J., P. Gao, and C. Parsons (2012). "The Price of a CEO's Rolodex". *Review of Financial Studies* 26.1, pp. 79–114.

Ersahin, N., R. Irani, and H. Le (2015). "Creditor Control Rights and Resource Allocation within Firms". *Working Paper*.

Ertimur, Y., F. Ferri, and D. Maber (2012). “Reputation Penalties for Poor Monitoring of Executive Pay: Evidence from Option Backdating”. *Journal of Financial Economics* 104.1, pp. 118–144.

Farwell, L. and R. Wohlwend-Lloyd (1998). “Narcissistic Processes: Optimistic Expectations, Favorable Self-Evaluations, and Self-Enhancing Attributions”. *Journal of Personality* 66.1, pp. 65–83.

Ferracci, M., G. Jolivet, and G. van den Berg (2014). “Evidence of Treatment Spillovers Within Markets”. *Review of Economics and Statistics* 96.5, pp. 812–823.

Ferreira, D., M. A. Ferreira, and B. Mariano (2015). “Creditor Control Rights and Board Independence”. *SSRN Working Paper* 2021522.

Fich, E. M. and A. Shivdasani (2007). “Financial Fraud, Director Reputation, and Shareholder Wealth”. *Journal of Financial Economics* 86.2, pp. 306–336.

Fisher, R. (1925). *Statistical Methods for Research Workers*. Edinburgh: Oliver and Boyd.

Flores, C., A. Flores-Lagunes, A. Gonzalez, and T. C. Neumann (2012). “Estimating the Effects of Length of Exposure to Instruction in a Training Program: The Case of Job Corps”. *Review of Economics and Statistics* 94.1, pp. 153–171.

Foster, J. D. and W. K. Campbell (2007). “Are There Such Things as "Narcissists" in Social Psychology? A Taxometric Analysis of the Narcissistic Personality Inventory”. *Personality and Individual Differences* 43.6, pp. 1321–1332.

FTI (2014). *Managing Cyber Risk: Job #1 for Directors and General Counsel*. URL: `http://www.fticonsulting.com/insights/fti-journal/managing-cyber-risk-job-1-for-directors-and-general-counsel` (visited on 05/12/2017).

Gabriel, M. T., J. W. Critelli, and J. S. Ee (1994). “Narcissistic Illusions in Self Evaluations of Intelligence and Attractiveness”. *Journal of Personality* 62.1, pp. 143–155.

Galasso, A. and T. S. Simcoe (2011). “CEO Overconfidence and Innovation”. *Management Science* 57.8, pp. 1469–1484.

Gamba, A. and A. Triantis (2012). “The Value of Financial Flexibility”. *Journal of Finance* 63.5, pp. 2263–2296.

Gerstner, W., A. Konig, A. Enders, and D. C. Hambrick (2013). “CEO Narcissism, Audience, Engagement, and Organizational Adoption of Technological Discontinuities”. *Administrative Science Quarterly* 58.2, pp. 257–291.

Gitelman, A. I. (2005). “Estimating Causal Effects From Multilevel Group-Allocation Data”. *Journal of Educational and Behavioral Statistics* 30.4, pp. 397–412.

Goel, A. M. and A. Thakor (2008). “Overconfidence, CEO Selection, and Corporate Governance”. *Journal of Finance* 63, pp. 2737–2784.

Goyal, V. and C. Park (2002). “Board Leadership and CEO Turnover”. *Journal of Corporate Finance* 8.1, pp. 49–66.

Graffin, S. D., J. B. Wade, J. F. Porac, and R. C. McNamee (2008). “The Impact of CEO Status Diffusion on the Economic Outcomes of Other Senior Managers”. *Organization Science* 19.3, pp. 457–474.

Graham, J. R. and C. R. Harvey (2001). “The Theory and Practice of Corporate Finance: Evidence from the Field”. *Journal of Financial Economics* 60.2-3, pp. 187–243.

Graham, J. R., C. R. Harvey, and M. Puri (2013). “Managerial Attitudes and Corporate Actions”. *Journal of Financial Economics* 109.1, pp. 103–121.

Graham, J. R. and K. Narasimhan (2004). “Corporate Survival and Managerial Experiences During the Great Depression”. *SSRN Working Paper* 489694.

Ham, C., N. Seybert, and S. Wang (2017). “Narcissism is a Bad Sign: CEO Signature Size, Investment, and Performance”. *SSRN Working Paper* 2144419.

Hansen, L. P. (1982). “Large Sample Properties of Generalised Method of Moments Estimators”. *Econometrica* 50.4, pp. 1029–1054.

Hege, U. and C. Hennessy (2010). “Acquisition Values and Optimal Financial (In)flexibility”. *Review of Financial Studies* 23.7, pp. 2865–2899.

Hennessy, C. and I. Strebulaev (2015). “Beyond Random Assignment: Credible Inference of Causal Effects in Dynamic Economies”. *NBER Working Paper* 20978.

Hirano, K. and G. Imbens (2004). *The Propensity Score with Continuous Treatments*. Ed. by A Gelman and X.-L. Meng. Vol. 0226164. Chichester, UK: John Wiley & Sons, Ltd, pp. 1–13.

Hirshleifer, D., A. Low, and S. H. Teoh (2012). “Are Overconfident CEOs Better Innovators?” *Journal of Finance* 67.4, pp. 1457–1498.

Hoberg, G. and G. Phillips (2010). “Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis”. *Review of Financial Studies* 23.10, pp. 3773–3811.

Hoberg, G. and G. Phillips (2015). “Text-Based Network Industries and Endogenous Product Differentiation”. *Journal of Political Economy* 124.5, pp. 1423–1465.

Hogan, R., G. J. Curphy, and J. Hogan (1994). “What We Know About Leadership: Effectiveness and Personality”. *American Psychologist* 49.6, pp. 493–504.

Hong, G. and S. W. Raudenbush (2006). “Evaluating Kindergarten Retention Policy”. *Journal of the American Statistical Association* 101.475, pp. 901–910.

Horton, R. S. and C. Sedikides (2009). “Narcissistic Responding to Ego Threat: When the Status of the Evaluator Matters”. *Journal Of Personality* 77.5, pp. 1493–1525.

Hudgens, M. G. and M. E. Halloran (2008). “Toward Causal Inference With Interference”. *American Statistician* 103.482, pp. 832–842.

Huson, M, R Parrino, and L Starks (2001). “Internal Monitoring Mechanisms and CEO Turnover: a Long-Term Perspective”. *Jounral of Finance* 56.6, pp. 2265–2297.

Imai, K. and M. Ratkovic (2014). “Covariate Balancing Propensity Score”. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 76.1, pp. 243–263.

John, O. P. and R. W. Robins (1994). “Accuracy and Bias in Self-Perception: Individual Differences in Self-Enhancement and the Role of Narcissism”. *Journal of Personality and Social Psychology* 66.1, pp. 206–219.

Judge, T. A., J. A. LePine, and B. L. Rich (2006). “Loving Yourself Abundantly: Relationship of the Narcissistic Personality to Self- and other Perceptions of Workplace Deviance, Leadership, and Task and Contextual Performance”. *Journal of Applied Psychology* 91.4, pp. 762–776.

Kaplan, S., M. Klebanov, and M. Sorensen (2012). “Which CEO Characteristics and Abilities Matter?” *Journal of Finance* 67.3, pp. 973–1007.

Knutson, R. (2017). *Why Verizon Decided to Stick With Yahoo Deal After Big Data Breaches*. URL: `https://www.wsj.com/articles/why-verizon-decided-to-still-buy-yahoo-after-big-data-breaches-1487679768` (visited on 05/12/2017).

Koh, P. S. and D. M. Reeb (2015). “Missing R&D”. *Journal of Accounting and Economics* 60.1, pp. 73–94.

Kolasinski, A. C. and X. Li. (2013). “Can Strong Boards and Trading Their Own Firm’s Stock Help CEOs Make Better Decisions? Evidence from Corporate Acquisitions by Overconfident CEOs”. *Journal of Financial and Quantitative Analysis* 48.4, pp. 1173–1206.

Koziol, J. A. and M. D. Perlman (1978). “Combinintg Independent Chi-Squared Tests”. *Journal of the American Statistical Association* 73.364, pp. 753–763.

Langevin, J. and J. Himes (2015). *Re: Disclosure Effectiveness Review*. URL: `https://langevin.house.gov/sites/langevin.house.gov/files/documents/06-17-15_Langevin_Himes_Letter_to_SEC.pdf` (visited on 05/12/2017).

Leary, M. T. and M. R. Roberts (2014). "Do Peer Firms Affect Corporate Financial Policy?" *Journal of Finance* 69.1, pp. 139–178.

Maccoby, M (2003). *The Productive Narcissist: The Promise and Peril of Visionary Leadership*. New York: Broadway Books.

— (2004). "Narcissistic Leaders: the Incredible Pros, the Inevitable Cons". *Harvard Business Review* 82.1, pp. 92–101.

Malmendier, U and G Tate (2008). "Who Makes Acquisitions? CEO Overconfidence and the Market's Reaction". *Journal of Financial Economics* 89, pp. 20–43.

— (2015). "Behavioral CEOs: the Role of Managerial Overconfidence". *Journal of Economic Perspectives* 29, pp. 37–60.

Malmendier, U. and S. Nagel (2011). "Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?" *Quarterly Journal of Economics* 126.1, pp. 373–416.

Malmendier, U. and G. Tate (2005). "Does CEO Overconfidence Affect Corporate Investment? CEO Overconfidence Measures Revisited". *European Financial Management* 11, pp. 649–659.

Malmendier, U., G. Tate, and J. Yan (2011). "Overconfidence and Early-life Experiences: The Effect of Managerial Traits on Corporate Financial Policies". *Journal of Finance* 66.5, pp. 1687–1733.

Manski, C. F. (2013). "Identification of Treatment Response with Social Interactions". *The Econometrics Journal* 16.1, S1–S23.

Morf, C. C. and F. Rhodewalt (2001). "Unraveling the Paradoxes of Narcissism: a Dynamic Self-Regulatory Processing Model". *Psychological Inquiry* 12.177-196.

Murphy, K. and J. L. Zummerman (1999). "Financial Performance Surrounding CEO Turnover". *Journal of Accounting and Economics* 16.1-2, pp. 377–419.

NACD (2017). *Cyber-Risk Oversight*. URL: https://www.nacdonline.org/files/FileDownloads/NACDCyber-RiskOversightHandbook2017.pdf (visited on 05/12/2017).

Nevicka, B. et al. (2011). “Reality at Odds with Perceptions: Narcissistic Leaders and Group Performance”. *Psychological Science* 22, pp. 1259–1264.

Nini, G., D. C. Smith, and A. Sufi (2009). “Creditor Control Rights and Firm Investment Policy”. *Journal of Financial Economics* 92, pp. 400–420.

— (2012). “Creditor Control Rights, Corporate Governance, and Firm Value”. *Review of Financial Studies* 25.6, pp. 1713–1761.

Otto, C. A. (2014). “CEO Optimism and Incentive Compensation”. *Journal of Financial Economics* 114.2, pp. 366–404.

Parino, R. (1997). “CEO Turnover and Outside Succession: a Cross-Sectional Analysis”. *Journal of Financial Economics* 46.2, pp. 165–197.

Patel, P. C. and D. Cooper (2014). “The Harder They Fall, the Faster They Rise: Approach and Avoidance Focus in Narcissistic CEOs”. *Strategic Management Journal* 35.10, pp. 1528–1540.

Pennebaker, J. W. and L. A. King (1999). “Linguistic Styles: Language Use as an Individual Difference”. *Journal of Personality and Social Psychology* 77.6, pp. 1296–1312.

Poitevin, M. (1989). “Financial Signalling and the "Deep-Pocket" Argument”. *The RAND Journal of Economics* 20.1, pp. 26–40.

Rapp, M. S., T. Schmid, and D. Urban (2014). “The Value of Financial Flexibility and Corporate Financial Policy”. *Journal of Corporate Finance* 29, pp. 288–302.

Raskin, R. and R. Shaw (1988). “Narcissism and the Use of Personal Pronouns”. *Journal of Personality* 56.2, pp. 393–404.

Raskin, R. and H. Terry (1988). "A Principal-Components Analysis of the Narcissistic Personality Inventory and Further Evidence of its Construct Validity". *Journal of Personality and Social Psychology* 54.5, pp. 890–902.

Reed, J., S. Collins, and M. Warner (2017). *Cybersecurity Disclosure Act of 2017*. URL: https://www.congress.gov/bill/115th-congress/senate-bill/536 (visited on 05/12/2017).

Ren, S. et al. (2010). "Nonparametric Bootstrapping for Hierarchical Data". *Journal of Applied Statistics* 37.9, pp. 1487–1498.

Roberts, M. R. and A. Sufi (2009). "Control Rights and Capital Structure: An Empirical Investigation". *Journal of Finance* 64.4, pp. 1657–1695.

Roll, R. (1986). "The Hubris Hypothesis of Corporate Takeovers". *Journal of Business* 59.2, pp. 197–216.

Rosenthal, S. A. and T. L. Pittinsky (2006). "Narcissistic Leadership". *The Leadership Quarterly* 17.6, pp. 617–633.

Rubin, D. B. (1974). "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology* 66.5, pp. 688–701.

Scally, D. (2014). *What Directors Think: 2014 Survey*. URL: www.nyse.com/publicdocs/nyse/listing/What_Directors_Think_2014.pdf (visited on 05/12/2017).

SEC. *Concept Release on Business and Financial Disclosure Required by Regulation S-K*. URL: https://www.sec.gov/comments/s7-06-16/s70616-25.pdf (visited on 05/12/2017).

Sobel, M. E. (2006). "What Do Randomized Studies of Housing Mobility Demonstrate?" *Journal of the American Statistical Association* 101.476, pp. 1398–1407.

Srinivasan, S. (2005). "Consequences of Financial Reporting Failure for Outside Directors: Evidence from Accounting Restatements and Audit Committee Members". *Journal of Accounting Research* 43.May, pp. 291–334.

Stathis, S. (2015). *Ocean Tomo Releases 2015 Annual Study of Intangible Asset Market Value*. URL: http://www.oceantomo.com/blog/2015/03-05-ocean-tomo-2015-intangible-asset-market-value/ (visited on 05/12/2017).

Wagner, W. G., J. Pfeffer, and I. O'Reilly C. A. (1984). "Organizational Demography and Turnover in Top-Management Groups". *Administrative Science Quarterly* 29, pp. 74–92.

Wallace, H. M. and R. F. Baumeister (2002). "The Performance of Narcissists Rises and Falls with Perceived Opportunity for Glory". *Journal of Personality and Social Psychology* 82.5, pp. 819–834.

Walsh, J. P. (1988). "Top Management Turnover Following Mergers and Acquisitions". *Strategic Management Journal* 9.2, pp. 173–183.

Walsh, J. P. and J. W. Ellwood (1991). "Mergers, Acquisitions, and the Pruning of Managerial Deadwood". *Strategic Management Journal* 12.3, pp. 201–217.

Wiersema, M. F. and K. A. Bantel (1993). "Top Management Team Turnover as an Adaptation Mechanism: the Role of the Environment". *Strategic Management Journal* 14, pp. 485–504.

Xuan, Y. (2009). "Empire-Building or Bridge-Building? Evidence from New CEO's Internal Capital Allocation Decisions". *Review of Financial Studies* 22.12, pp. 4919–4948.

Yim, S. (2013). "The Acquisitiveness of Youth: CEO Age and Acquisition Behavior". *Journal of Financial Economics* 108.1, pp. 250–273.

Zeigler-Hill, V., E. M. Myers, and C. B. Clark (2010). "Narcissism and Self-Esteem Reactivity: the Role of Negative Achievement Events". *Journal of Research in Personality* 44.2, pp. 285–292.

Zhang, Z. (2016). "Bank Interventions and Trade Credit : Evidence from Debt Covenant Violations". *SSRN Working Paper* 2482249.

Zhu, D. H. and G. Chen (2015a). “CEO Narcissism and the Impact of Prior Board Experience on Corporate Strategy”. *Administrative Science Quarterly* 60.1, pp. 31–65.

— (2015b). “Narcissism, Director Selection, and Risk-Taking Spending”. *Strategic Management Journal* 36.13, pp. 2075–2098.

APPENDIX A

VARIABLE DEFINITIONS FOR SECTION 1

This section defines the Compustat variables used in Section 1.

Outcomes:

- Investment[1]: `capxy/ppentq`
- Debt issuance[2]: `[(dlttq+dlcq)-(L.dlttq+L.dlcq)]/atq`
- Sales growth: `(saleq-L.saleq)/saleq`

Controls

- Leverage ratio: `(dlttq+dlcq)/atq`
- Net worth to assets: `seqq/atq`
- Current ratio: `actq/lctq`
- Interest expense scaled by avg. assets: `xintq/[(atq+L.atq)/2]`
- Operating income scaled by avg. assets: `oibdpq/[(atq+L.atq)/2]`
- Market to book: `[(prccq*cshoq)+(atq-ltq+txditcq)+atq]/atq`
- Return on assets: `niq/atq`
- Size: `log(atq)`
- PP&E to assets: `ppentq/atq`
- Herfindahl index: $\sum_{i \in \text{ industry } j} \texttt{saleq}_i^2$

[1] I take first differences of `capxy` in fiscal quarters $2-4$.
[2] I use `L.x` to denote the first lag of x.

APPENDIX B

FIRST STAGE COVARIATE BALANCE FOR SECTION 1

This section shows how to calculate a single test statistic for covariate balance following the propensity score calculation used in the group-by-group matching (section 2.2.3.1). Maximizing the log-likelihood function of a parametrically specified propensity score – $Pr(\tau_i = 1|x_i) = \pi_\beta(x_i)$ – gives the first-order condition

$$\frac{1}{N}\sum_{i=1}^{N} \frac{\tau_i \pi'_\beta(x_i)}{\pi_\beta(x_i)} - \frac{(1-\tau_i)\pi'_\beta(x_i)}{1-\pi_\beta(x_i)} \tag{B.1}$$

for x_i a $K \times 1$ vector. When the dimension of β is also $K \times 1$, as is the case when using a logistic model, (B.1) shows the K moment conditions one could use to calculate the propensity score via GMM. Following Imai and Ratkovic (2014), covariate balance implies the moment conditions

$$\mathrm{E}\left[\frac{\tau_i f(x_i)}{\pi_\beta(x_i)} - \frac{(1-\tau_i) f(x_i)}{1-\pi_\beta(x_i)}\right] = \mathbf{0} \tag{B.2}$$

for *any* vector-valued measurable function of x_i.

The insight of Imai and Ratkovic (2014) is to take $f(x_i) = x_i$ and apply Hansen's test for overidentifying restrictions (Hansen 1982) to the K moment conditions specified in (B.2) with the just-identified model given by the moments in (B.1). Thus for a group of firms c, matching between treated and untreated firms yields the hypothesis:

$$H_0^c : \mathrm{E}\left[\frac{\tau_i x_i}{\pi_\beta(x_i)} - \frac{(1-\tau_i) x_i}{1-\pi_\beta(x_i)}\middle| i \in c\right] = \mathbf{0}_{K\times 1}. \tag{B.3}$$

Because matching between treated and untreated firms occurs at a group level, there are many groups $\{c_1, ..., c_m\}$ that each must satisfy covariate balance. If not, one or more of the estimated groups used in the second step may have predicted outcomes that are driven by observable differences between treated and untreated firms. To test whether this is the case, I use Fisher's method and aggregate the p-values from B.3 into a single hypothesis (Fisher 1925):

$$H_0 : \cap_{j=1}^{m} \left(H_0^{c_j} : \mathbf{E} \left[\frac{\tau_i x_i}{\pi_\beta(x_i)} - \frac{(1-\tau_i)x_i}{1-\pi_\beta(x_i)} \middle| i \in c_j \right] = \mathbf{0}_{K\times 1} \right) \tag{B.4}$$

against the alternative that there is at least one group, c_j, for which Hansen's test for overidentification rejects the null. The joint hypothesis H_0 is rejected whenever $-2\Sigma_{j=1}^{m}\log(p_j) \geq \chi^2_{2m,\alpha}$ for α the level of significance and p_j the p-value from the $H_0^{c_j}$.

Because the test statistic J used for (B.3) is itself distributed χ^2 under the null, an alternative approach is to aggregate these individual statistics and reject the joint null whenever

$$\Sigma_{j=1}^{m} J_j > \chi^2_{\Sigma k_j,\alpha} \tag{B.5}$$

for k_j the degrees of freedom of the j^{th} test (Koziol and Perlman 1978). I report both the overidentification test (B.4) and the sum test (B.5) when testing for covariate balance.

APPENDIX C

WORD LISTS FOR SECTION 2

Cybersecurity words. I search 10-K item 1A filings for the following words related to cybersecurity risk: cyber, data?breach*, cyber?security, cyber?attack*, computer?hack*, hack*, information?security, unauthorized?access, security?breach*. I use ? as a wildcard to mean either "" or " " since some firms use the spelling "cybersecurity" while others use "cyber security" or "cyber-security" (I convert hyphens to spaces in pre-processing). The wildcard * denotes any number of trailing alphabetic characters so that plurals (e.g. "data breaches") are picked up in the analysis.

Boardex committee names. The following is a list of all committee names in the Boardex database that I use to classify a director as having a technology role: cyber security, technology, technology environmental social responsibility, technology & development, technology & products, technology advisory, technology risk, technology strategy and innovation, technology strategy and investment, technology and acquisition, technology and competition, technology and corporate responsibility, technology and environment, technology and quality, technology and reserves, technology and safety, technology and science, technology and strategy, technology and transactions, e-commerce and technology, health it, health safety environment and technology, it, it oversight, it steering, information systems, information systems steering committee, information technology, information technology and security, risk and information security, safety environment and technology, science and technology, science and technology advisory, scientific and technology, technical, technical health safety and environmental, technical advisory, technical services, technical and commercial oversight, technical and operations, technical and projects, technical and reserves, technical and resources, technical and safety.

The following is a list of all committee names in the Boardex database that I use to classify a director as having a risk role: corporate risk, enterprise risk management, governance & risk, governance nominating & risk oversight, governance and risk, risk, risk compliance and planning, risk & credit, risk assessment, risk capital and subsidiaries, risk management, risk management & compliance, risk management and finance, risk oversight, risk oversight and management, risk policy, risk policy and capital, risk review, risk review investment and loan, risk and capital, risk and credit policy, risk and public policy, risk and regulatory, risk and return, risk and safety health and environment.