Three essays in finance

Phua, Kenny Jing Wen

2018

Phua, K. J. W. (2018). Three essays in finance. Doctoral thesis, Nanyang Technological University, Singapore.

<http://hdl.handle.net/10356/73254>

<https://doi.org/10.32657/10356/73254>

**NANYANG
TECHNOLOGICAL
UNIVERSITY**

SINGAPORE

THREE ESSAYS IN FINANCE

**PHUA JING WEN KENNY
NANYANG BUSINESS SCHOOL**

2018

THREE ESSAYS IN FINANCE

PHUA JING WEN KENNY

PHUA JING WEN KENNY

Nanyang Business School

A thesis submitted to the Nanyang Technological University
in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

ACKNOWLEDGEMENTS

I am most grateful to my adviser Wei Chishen for his astute guidance in my doctoral program. My thesis advisory committee members—Hwang Chuan Yang and Angie Low—provided outstanding advice and help in my academic development. The fruition of this thesis is also made possible by the invaluable feedback provided by numerous participants in seminars and conferences. The formidable ladies of the Ph.D. office—Bee Hua, Karen, Tsai Ting, and Nisha—provided amazing support in administrative matters.

Contents

Chapter One

Are overconfident CEOs better leaders? Evidence from stakeholder commitments pg. 1 - 70

Chapter Two

Do security analysts learn from their colleagues? pg. 71 - 153

Chapter Three

Commonality in analyst coverage and information diffusion in networks pg. 154 - 206

Are overconfident CEOs better leaders? Evidence from stakeholder commitments

Kenny Phua
Nanyang Technological University

T. Mandy Tham
No Affiliation

Chishen Wei*
Singapore Management University

Abstract

We find evidence that the leadership of overconfident chief executive officers (CEOs) induces stakeholders to take actions that contribute to the leader's vision. By being intentionally overexposed to the idiosyncratic risk of their firms, overconfident CEOs exhibit a strong belief in their firms' prospects. This belief attracts suppliers beyond the firm's observable expansionary corporate activities. Overconfident CEOs induce more supplier commitments including greater relationship-specific investment and longer relationship duration. Overconfident CEOs also induce stronger labor commitments as employees exhibit lower turnover rates and greater ownership of company stock in benefit plans.

JEL Classification Code: G32, J53, J54, L14, L22

Keywords: CEO overconfidence, Leadership, Customer-supplier, Employee ownership

“I told people you weren’t betting on a device. You were betting on Steve Jobs.”

~ Randall Stephenson (AT&T CEO)

1. Introduction

Managerial overconfidence can significantly affect corporate activities and outcomes. Studies show that overconfident CEOs affect investment decisions, merger and acquisition choices, and accounting practices.¹ Yet, some of the most successful leaders, such as Jack Welch of General Electric and Steve Jobs of Apple Inc., displayed managerial overconfidence during their tenure as CEO. Recent studies have uncovered important benefits to employing overconfident CEOs such as higher research & development (R&D) productivity and innovation output (e.g., Galasso and Simcoe, 2011; Hirshleifer, Low, and Teoh, 2012). We add to this growing literature by asking: Are overconfident CEOs better leaders?

Our definition of leadership follows Hermalin (1998), where a leader’s actions motivate key stakeholders, such as employees and suppliers, to exert greater effort. Leadership is distinct from formal authority because stakeholders’ actions are voluntary. In this context, suppliers choose to invest in relationship-specific assets and to sell their products to customers. Employees select their employment and level of effort on the job. To motivate stakeholder actions, a leader must have strong self-belief and belief in the firm’s prospects under her leadership. Recent psychology studies show that overconfident people are more respected, more influential, and viewed as more competent by their peers (e.g., Anderson, Brion, Moore, and Kennedy, 2012; Kennedy, Anderson, and Moore, 2013). These psychological underpinnings motivate our hypothesis that the leadership of overconfident CEOs may attract greater stakeholder commitments.

Short of conducting interviews or running experiments, leadership and influence are typically unobservable and difficult to quantify and measure. Instead, we test our *leadership hypothesis* by focusing on the observable actions of two key stakeholders of the

¹ For example, see Malmendier and Tate (2005, 2008), Billet and Qian (2008), Kolasinski and Li (2013), Schrand and Zechman (2012), Ahmed and Duellman (2013), Banerjee, Humphery-Jenner, Nanda, and Tham (2017).

firm, suppliers and employees. This is an ideal test setting because voluntary participation represents a significant leadership outcome, particularly when stakeholder effort and commitment are critical to a firm’s success. From the stakeholder’s view, commitment and effort require costly investments that may be firm-specific and have low redeployability value outside of the relationship.² Therefore, stakeholders will only make such investments if they strongly believe in the leadership of the firm’s CEO.

The success of the original iPhone is an example of the importance of stakeholder commitment. AT&T (then Cingular) helped Apple Inc. to secretly develop the iPhone and made heavy concessions to become the exclusive iPhone carrier in the U.S. market, effectively tying their fate to the iPhone’s future prospects.³ This example shows the close interdependency between a firm’s success and the commitments made by a firm’s stakeholders towards product design and quality. More notably, the decision of Randall Stephenson, the CEO of AT&T, to commit relationship-specific investment to Apple Inc. was not motivated by the prospect of the iPhone per se, but by his belief in the leadership of Steve Jobs, the CEO of Apple Inc. at the time.

Following extant literature, we measure CEO overconfidence using vested-in-the-money stock options of CEOs in the ExecuComp database.⁴ While holding vested-in-the-money stock options may underdiversify the CEO’s wealth (Hall and Murphy, 2002), it may also display leadership for at least three reasons. First, having “skin in the game” conveys the CEO’s strong belief in the firm’s prospects because the CEO’s human capital is already tied to the firm. Second, it demonstrates a willingness to lead by example. Hermalin (1998) argues that leading by example may be a credible form of leadership and cites historical examples including Dr. Martin Luther King, Jr. marching at the head of civil rights marches. Third, holding vested-in-the-money options may reflect a commitment by the manager to exert greater effort (Gervais, Heaton, and Odean, 2011). The manager believes the firm has valuable growth opportunities and therefore works

² A classic example is the steel dies used to form the body of a car model which have low outside practical value (Klein, Crawford, and Alchian, 1978).

³ “Life after the iPhone: How AT&T’s bet on Apple mobilized the company” *Forbes*, January 21, 2013.

⁴ For example, see: Campbell et al. (2011), Malmendier, Tate, and Yan (2011), Hirshleifer, Low, and Teoh (2012), Banerjee, Humphery-Jenner, and Nanda (2015).

harder. Stakeholders understand that this type of leadership and dedication to hard work translates into higher firm value, which makes the bilateral relationship more valuable.

Our first finding indicates that overconfident CEOs are more likely to initiate and to expand dependent supplier networks.⁵ Controlling for firm characteristics, CEO overconfidence increases the probability of adding a dependent supplier by +0.9%. This effect is economically large given that on average, only 7.2% of firm-year observations in our sample experience an increase in dependent suppliers. These findings suggest that CEO overconfidence facilitates the development of important bilateral relationships with suppliers.

A concern with this interpretation is the possibility of an omitted variable that drives the CEO to hold vested-in-the-money options and induces growth in the firm’s supplier network. We partially address this issue by estimating conditional logit regressions that include firm strata, which capture unobserved firm heterogeneity, and industry-year strata, which capture industry growth cycles. We also employ an overconfidence measure by Kolasinski and Li (2013) that requires active managerial choice rather than inaction or inertia. While these robustness tests do not entirely eliminate endogeneity concerns, they indicate that our results are unlikely driven by industry growth cycles, unobserved firm heterogeneity, and insider information.

Our second set of tests focuses on stakeholder commitments. Employees are arguably the most important stakeholders and their commitments can impact firm performance. First, we measure employee commitment using turnover because committed employees are likely to stay on with their employer. Using options cancellation data as a proxy of employee turnover (e.g., Carter and Lynch, 2004; Babenko and Sen, 2014), we find that firms led by overconfident CEOs are associated with a 2.4% decrease in employee turnover. Second, we measure employee commitment based on the amount of employer stock held in employee retirement benefit plans. We expect that committed employees are more likely to own more company stock. Using data from filings of Internal Revenue

⁵ We follow Statement of Financial Accounting Standards (SFAS) 14 in classifying dependent suppliers as firms that generate at least 10% of revenues from a customer firm. Sample firms that are not reported as customers in this data set are assumed to have no dependent suppliers.

Service (IRS) Form 5500, we find that the allocation of benefit plan assets to company stock is 2.7% higher in firms led by overconfident CEOs. Overall, our evidence indicates that overconfident CEOs are able to induce greater employee commitment.

We also expect that the leadership of an overconfident CEO is important when supplier commitments are particularly valuable. For example, specialized inputs, such as the components of an iPhone, require customized supplier investment. These relationship-specific investments create value by improving productivity (Dyer and Nobeoka, 2000), enhancing core competency (Parmigiani, 2007), and stimulating inter-project spillovers (Kang et al., 2009). But unless suppliers have strong conviction in the leadership of their customer firm’s CEO, they may be reluctant to develop customized inputs for customer firms due to costly initial investment and low ex post redeployability of the inputs if the relationship terminates prematurely (Titman, 1984).

We test this prediction by examining firms in durable manufacturing industries because these industries produce unique products that require greater relationship-specific inputs from suppliers (Titman and Wessels, 1988; Banerjee, Dasgupta, and Kim, 2008). Consistent with this hypothesis, our results are stronger in the durable manufacturing industries. The effects are also more pronounced for firms in high ‘contract intensity’ industries where contract intensity represents the depth of relationship-specific investment between supplier and customer industries (Nunn, 2007).

We further find direct evidence of relationship-specific investment in two additional tests. First, CEO overconfidence is associated with greater suppliers’ R&D intensity, which is a proxy for relationship-specific investments (e.g., Kale and Shahrur, 2007; Raman and Shahrur, 2008). This effect is more pronounced for smaller suppliers who typically have less bargaining power. The leadership of overconfident CEOs potentially mitigates suppliers’ concerns regarding hold-up problems that arise from weaker bargaining power. Second, we measure the depth of supply chain commitment using the duration of a customer-supplier relationship (Fee, Hadlock, and Thomas, 2006). Estimates from the hazard model suggest that the customer-supplier relationship is less likely to

terminate when the customer CEO is overconfident. Together, this set of findings suggests that overconfident CEOs induce greater supplier commitments.

We also design an additional test based on extant evidence that overconfident CEOs have higher investment-cash flow sensitivities (Malmendier and Tate, 2005). This implies that when cash flow is low, overconfident CEOs tend to refrain from investment expenditures. If suppliers are solely focused on signals related to expansionary activities, they are less likely to supply to firms led by overconfident CEOs with low cash flow. However, after separating firms based on cash flow, we find that overconfident CEOs of low cash-flow firms also attract suppliers. This implies that the anticipation of future expansionary corporate activities is unlikely to explain our findings.

While our set of findings provides support for the leadership hypothesis, our evidence is potentially consistent with a ‘dark-side’ view of overconfidence under which CEOs overpay to acquire these commitments. In this scenario, the leadership outcomes that we document may lower profitability and hurt the company’s bottom line. However, additional tests uncover no such evidence. Instead, firms led by overconfident CEOs have higher future gross profitability, enjoy lower input costs, and generate higher risk-adjusted returns relative to their competitors. This lends credence to the leadership hypothesis and is consistent with a bright-side view that overconfident CEOs enhance firm value.

An alternative imperfectly rational interpretation of our findings is that stakeholder commitments are not due to the leadership channels of CEO overconfidence discussed earlier, but rather some combination of CEO and stakeholder irrationality. Perhaps stakeholders naively accept the CEO’s overconfidence, which lacks actual leadership actions that improve the value of the bilateral relationship. However, it seems unlikely that suppliers would make costly relationship-specific investments unless they are convinced that the CEO will exert efforts to ensure a long-term sustainable relationship. Similarly, employees are unlikely to make costly commitments in their firm without a strong belief that their CEO will work hard to deliver value.

Lastly, we address the reporting-bias issue of the customer-supplier data set that firms may selectively disclose customer identities (Ellis, Fee, and Thomas, 2012).⁶ Suppose suppliers that commit greater relationship-specific investments are more likely to report the identities of their customers and also take stronger cues from the overconfidence of their customer firms’ CEOs. This example would work in favor of us finding empirical support for our leadership hypothesis. However, the authors find the opposite inclination that suppliers with greater proprietary costs are less inclined to reveal customers’ identities. This evidence works against us finding our results.

A key contribution of our study is to provide empirical evidence supporting theories on leadership by examining stakeholder actions. Hermalin (1998) and Komai, Stegeman, and Hermalin (2007) show that by setting an example, managers may signal private information, and motivate subordinates to work harder. Almazan, Chen, and Titman (2017) show that “top-down” capital allocation may optimally create higher levels of investment expenditure to motivate effort from employees. We find that stakeholders provide greater commitment to the firm when the CEO is overconfident. In particular, the finding on supplier commitment emphasizes that CEO leadership reaches beyond the boundaries of the firm to include external stakeholders.

A large literature shows that overconfident CEOs have a significant impact on firm outcomes. An open question is why boards appoint overconfident CEOs (e.g., Goel and Thakor, 2008; Gervais, Heaton, and Odean, 2011) when evidence suggests that CEO overconfidence causes investment distortions, costly mergers and acquisitions (M&A), and loose accounting practices.⁷ Recent studies find a ‘bright side’ of CEO overconfidence. Overconfident CEOs produce higher R&D productivity, generate better innovative output, and convert growth opportunities into firm value (e.g., Galasso and Simcoe, 2011; Hirshleifer, Low, and Teoh, 2012). Our evidence is also consistent with a ‘bright side’ of

⁶ SFAS 131 which was issued in June 1997 requires firms to disclose sales to each material customer, but not the identity of the customer.

⁷ See, Malmendier and Tate (2005, 2008), Billet and Qian (2008), Kolasinski and Li (2013), Schrand and Zechman (2012), Ahmed and Duellman (2013), Banerjee, Humphery-Jenner, Nanda, and Tham (2017).

CEO overconfidence. We find evidence that overconfident CEOs deliver important leadership outcomes in the form of valuable commitments from key stakeholders.

2. Sample selection and data

We start with firms in the ExecuComp database with available CEO stock option data. Following the standard literature, we remove utilities (Standard Industry Classification (SIC): 4000-4999) and financial firms (SIC: 6000-6999). Next, we identify customer and supplier pairs from the business segment files of Compustat. In accordance with SFAS 14, public firms are required to disclose sales to their principal customers, defined as customers that contribute to at least 10% of the total revenue of the firm or if sales to a customer are material to the business of the firm. Principal customer names are manually matched to Compustat GVKEYs following the approach in Fee, Hadlock, and Thomas (2006).⁸

We identify ExecuComp firms that are reported as customers by firms in the customer-supplier data set. *Suppliers increase (decrease)* is an indicator equal to one if the year-on-year change in the number of dependent suppliers is positive (negative), and zero otherwise. *Start of supplier network* equals one if a firm has at least one dependent supplier in year t , and none in year $t-1$, and equals zero otherwise.

Financial variables and stock return data are obtained from Compustat and Center for Research in Security Prices (CRSP). We collect insider trades from Thomson Insider and acquisition data from Securities Data Company (SDC) Platinum. Our sample period starts from 1993 and ends in 2011, which is the last year that we have information on customer-supplier pairs. Using these databases, we construct the following controls: firm size, leverage, return on assets (ROA), ROA volatility, sales growth, market-to-book ratio, investment, R&D, cash holdings, past stock return, CEO ownership, CEO tenure,

⁸ For customer names that are abbreviated, we hand-match and use industry affiliations to determine whether the customer is in Compustat. For the remaining unmatched customers, we check their corporate websites in the Directory of Corporate Affiliation (DCA) database to determine if the customer is a subsidiary of a listed firm. If so, we assign the customer to its parent's GVKEY. To ensure accuracy, we discard any customer name that cannot be unambiguously matched to a GVKEY.

acquisitions, and financing activities.⁹ We have 1,921 unique firms from 1993 to 2011, for a panel totaling 14,745 firm-year observations.

2.1. Measures of CEO overconfidence

Our primary measure of CEO overconfidence is a stock option-based measure, motivated by Malmendier and Tate (2005). Using CEO stock options data from ExecuComp, we compute the average moneyness of the CEO’s option holdings annually following the approach in Campbell et al. (2011), Malmendier, Tate, and Yan (2011), Hirshleifer, Low, and Teoh (2012), and Banerjee, Humphery-Jenner, and Nanda (2015). Specifically, we obtain the number and value of the CEO’s vested stock options to construct *CEO overconfidence* as the ratio of average value per option to average strike price, where the average value per option is the total value of the CEO’s option holdings (ExecuComp: opt_unex_exer_val) scaled by the number of such options (ExecuComp: opt_unex_exer_num). The average strike price is the firm’s stock price at the end of the fiscal year (Compustat: prcc_f) less the average value per option. We define *Confident CEO (options)* as an indicator equal to one if the *CEO overconfidence* measure is at least 67% in-the-money on at least two occasions in the past five years (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013).¹⁰

We also create two indicator variables to capture when overconfidence is revealed so we can precisely connect to the timing of stakeholder actions. *CEO overconfidence up* is an indicator equal to one when *CEO overconfidence* moves into the top quartile from year $t-1$ to year t , and zero otherwise. *CEO overconfidence down* is an indicator equal to one when *CEO overconfidence* falls out of the top quartile in year t from year $t-1$.

Since unexercised options represent a non-action, it may reflect inattention or inertia rather than overconfidence. We estimate a measure of overconfidence that tracks the CEO’s unprofitable insider purchases of firm stock and thus captures trading actions rather than inaction or inattention (Kolasinski and Li, 2013). *Overconfidence trade* is an

⁹ The Appendix provides full variable construction details.

¹⁰ Since variation in moneyness is a direct function of stock prices, we control for past stock returns in our tests following Malmendier, Tate, and Yan (2011), Hirshleifer, Low, and Teoh (2012), and Malmendier and Tate (2015).

indicator equal to one if the CEO purchases shares over the next two years that experience negative buy-and-hold returns benchmarked against the Fama-French size-decile portfolio. To capture the timing of overconfidence, we construct an indicator variable, *Overconfidence trade up*, equal to one if *Overconfidence trade* is unity in year t but not in year $t-1$, and zero otherwise.

We also use a media-based measure of overconfidence following Banerjee, Humphery-Jenner, and Nanda (2015). The measure is based on keyword search for references to confidence and non-confidence in media releases.¹¹ *CEO media positivity* is equal to one in the year if the number of ‘confident’ articles exceeds the number of ‘non-confident’ articles, and zero otherwise. Missing *CEO media positivity* values are set to zero.

2.2. Measures of supplier commitment

We measure supplier commitment using relationship-specific investment (RSI) and relationship duration. Following Kale and Shahrur (2007) and Raman and Shahrur (2008), we proxy for RSI using suppliers’ R&D expenditure. Specifically, we define *Supplier R&D intensity* as the product of the supplier’s R&D expenditure and the fraction of sales to the customer, divided by total assets of the supplier. The normalization allows for comparability across suppliers of different sizes. We measure relationship duration using the duration of customer-supplier relationships (Fee, Hadlock, and Thomas, 2006). In Section 3.3.3, we conduct survival analysis using the Andersen-Gill extension of the Cox model on the termination of the relationship at a given point in time (Andersen and Gill, 1982).

2.3. Measures of employee commitment

We measure employee commitment using: 1) *Employee turnover* and 2) *Employee ownership* of the company stock in their benefits plans. We proxy for *Employee turnover* using the rate of stock option cancellations, forfeitures, and expirations (e.g., Carter and Lynch, 2004; Babenko and Sen, 2014) because firm-level employee turnover is not publicly

¹¹ References to confidence are *overconfident*, *overconfidence*, *optimistic*, and *optimism*. References to non-confidence are *reliable*, *cautious*, *conservative*, *practical*, *frugal*, and *steady*. The search is done through Factiva database for articles referring to the CEOs in *The New York Times*, *Business Week*, *Financial Times*, *The Economist*, *Forbes Magazine*, *Fortune Magazine*, and *The Wall Street Journal*.

available.¹² Cancellations, forfeitures, and expiration of stock options may be reasonable proxies for labor turnover because these typically occur upon employee separation from the firm. Carter and Lynch (2004) find that the option-based measure of employee turnover is positively correlated ($\rho=0.66$, $p\text{-value}<0.01$) with actual industry-level employee turnover data from the Saratoga Institute. To address concerns that employee retirement induces noise in this measure, our tests include firm age to proxy for workforce age because Ouimet and Zarutskie (2014) find that young firms disproportionately employ younger workers in firm-level U.S. Census data.

We gather stock option and option cancellation data from Compustat, available from the year 2004 onwards, and construct *Employee turnover* as the ratio of stock option cancellations (Compustat: optca)¹³ to the number of non-executive employee stock options outstanding at the beginning of the year defined as the total number of stock options outstanding (Compustat: optosey) minus the sum of the number of exercisable and non-exercisable executive stock options outstanding (ExecuComp: opt_unex_exer_num, opt_unex_unexer_num). To capture the turnover of rank-and-file employees, we further classify employee stock option plans as broad-based if the volume of option grants to non-executive employees exceeds 0.5% of the number of shares outstanding (e.g., Oyer and Schaefer, 2005; Hochberg and Lindsey, 2010; Chang et al., 2015).

Our second measure captures the view that committed employees are more likely to overweight company stock in their retirement benefit accounts when they believe in the company’s future prospects under the CEO’s leadership. From the 2004–2011 Form 5500 Schedule H filings of employee retirement plans (i.e., defined benefit and contribution plans), we collect the dollar value of company stock held in employee retirement plans and create two measures—*Employee stock holdings (%)* and *Per-employee stock holdings (\$)*. *Employee stock holdings (%)* is the aggregate dollar-value of holdings of employer stock divided by the total benefit plan assets (item 1f of Form 5500 Schedule H). *Per-*

¹² *Fortune Magazine* publishes the employee turnover rates of companies in the Fortune 100 list, but this limited sample includes many private companies and is therefore not in our sample.

¹³ Compustat item *OPTCA* is an aggregation of cancelled, forfeited, expired, terminated, and lapsed stock options. This aggregation is consistent with the finding of Carter and Lynch (2004) that many firms do not provide a breakdown of cancelled, forfeited, and expired options.

employee stock holdings (\$) is the average employee dollar value of employer stock holdings calculated as the dollar value of employer stock held divided by the number of active plan participants (item 6a2 of Form 5500 Annual Return/Report). We match approximately 82% of our sample to the Form 5500 filings. Our tests also include controls for employee and plan-specific characteristics (i.e., cash/stock match, employee wealth) that may affect the employee’s choice of stock ownership in benefit plans.

2.4. *Summary statistics*

Panel A of Table 1 shows that 53.3% of our customer-years are led by overconfident CEOs. Changes in CEO overconfidence (*CEO overconfidence up*) are only found in 6.8% of sample observations. 57% of customer-supplier relationships have overconfident customer CEOs. Rows 6 and 7 show that the arrival of dependent suppliers (mean=7.2%) is about as frequent as the departure of dependent suppliers (mean=7.5%). The start of supplier networks (Row 8) are rarer, occurring in 3.3% of the sample observations.

[Insert Table 1]

Panel B of Table 1 partitions the sample firms into firms with dependent suppliers (i.e., customer firms) and firms without dependent suppliers (i.e., standalone firms). There is no difference in likelihood of an overconfident CEO among these two groups, however CEOs of customer firms have a higher confidence level (i.e., greater vested-in-the-money option holdings). Relative to the average standalone firm, the average customer firm is about six times larger based on total assets, has higher valuation based on the market-to-book ratio (3.82 vs 3.05), has higher ROA (6.6% vs 5.0%), has lower ROA volatility (4.6% vs 5.5%), and has higher sales growth (13.2% vs 12.1%). The average customer firm also makes more investment (7.1% vs 6.5%) and has higher R&D intensity (4.8% vs 4.2%), but has higher leverage (24.6% vs 21.9%) and lower cash levels (9.8% vs 12.7%). Stock performance between the two types of firms is not statistically different.

3. **Overconfident CEOs and leadership outcomes**

This section tests the hypothesis that overconfident CEOs induce leadership outcomes. In our setting, leadership outcomes are voluntary stakeholder actions rather than mandates from formal authority. In Section 3.1, we test whether overconfident CEOs are

able to attract key suppliers. Sections 3.2 and 3.3 analyze employee and supplier commitments.

3.1. Do overconfident CEOs attract suppliers?

Suppliers provide necessary inputs and are vital to firm success. They also represent a key leadership outcome because they are independent entities that cooperate voluntarily. We examine whether CEO overconfidence attracts suppliers by estimating a logistic regression model using Eq. (1):

$$\text{Suppliers Increase}_{i,t} = a + \beta_1 \text{CEO overconfidence Measure}_{i,t-1} + \beta_2 \varphi_{i,t-1} + \varepsilon_{i,t} \quad (1)$$

Suppliers increase is an indicator equal to one if the firm experiences an increase in the number of suppliers from year $t-1$ to t , and zero otherwise. The results are also similar using an ordinary least squares (OLS) model with the number of suppliers as the dependent variable. We estimate logit and conditional logit models using various *CEO overconfidence* measures described in Section 2 and the Appendix. φ represents a vector of control variables that are commonly used in the customer-supplier literature (see, e.g., Banerjee, Dasgupta, and Kim, 2008; Raman and Shahrur, 2008). In all tests, we control for past stock returns because it is related to the moneyiness of the CEO's vested stock options. The baseline regression specification includes year indicators to control for macroeconomic trends and industry indicators to capture differences across industries, following the approach in Malmendier, Tate, and Yan (2011). We cluster standard errors at the firm level.

[Insert Table 2]

The evidence in Table 2 suggests that suppliers respond strongly to the presence of an overconfident CEO. The univariate logit results in Column 1 of Panel A show a positive and significant effect of the *Confident CEO (options)* measure on *Suppliers increase*. Column 2 shows that the finding remains robust with the inclusion of firm characteristics as controls. We are careful to include past stock returns, investment (capital expenditure), an acquisition indicator, and SEO proceeds as controls because *Confident CEO (options)* may be related to recent stock performance and suppliers may be attracted to

expansionary corporate activities. Evaluated at the means of the other independent variables, CEO overconfidence has a marginal effect of +0.9% on the likelihood of the firm experiencing an expansion of its supplier network. This represents a 12.5% increase over the average frequency of *Suppliers increase* (7.2%) in our sample.

We conduct a sharper test of the leadership hypothesis by examining the timing of CEO overconfidence. This test is important for two reasons. First, a corollary of the leadership hypothesis is that supplier actions should occur during periods when the CEO provides leadership. To measure the timing of leadership, we use the *CEO overconfidence up* indicator variable. Second, this test may help address plausible alternative explanations such as underlying industry shocks that may simultaneously increase demand and motivate overconfident CEOs to provide leadership. To capture time-varying industry-wide changes, we adopt a conditional logit model that stratifies observations along the industry-year dimension based on two-digit SIC. While it cannot completely rule out endogeneity concerns, the conditional logit model sidesteps the incidental parameters problem. However, it inevitably causes fluctuations in sample size across specifications.¹⁴ Standard errors are clustered at the industry level.

The conditional logit estimates with industry-year strata suggest that suppliers react during times of leadership action. In Column 3, the positive and significant coefficient estimate on *CEO overconfidence up* suggests that the timing of leadership is important as suppliers are more likely to form relationships during times when the customer CEO provides leadership. The inclusion of industry-year strata mitigates the possibility that time-varying industry shocks drive our findings. Another concern is unobserved underlying firm heterogeneity. To capture unobserved firm heterogeneity, we exploit the timing of the *CEO overconfidence up* measure by estimating a conditional logit model with firm strata. Column 4 shows that the coefficient estimate on *CEO overconfidence up* remains positive and significant with the inclusion of firm strata.

¹⁴ For example, when the stratification is industry fixed effects, the conditional logit model constrains the coefficient estimates to be equivalent across regressions in all industries. If a certain stratum experiences no variation in the dependent variable, all observations in that stratum are eliminated. This condition may cause variation in sample sizes across different specifications.

Overall, our evidence is consistent with the view that overconfidence attracts suppliers. While we cannot rule out all possible endogeneity explanations, the conditional logit estimates suggest that unobserved firm heterogeneity is not behind our findings.

3.1.1. *Alternative measures of CEO overconfidence*

It is possible that our findings are sensitive to the construction of CEO overconfidence measures. Also, prior studies find that the degree of CEO overconfidence may be important (see, e.g., Campbell et al., 2011; Hirshleifer, Low, and Teoh, 2012). Therefore, we re-estimate the conditional logit model using the continuous measure of *CEO overconfidence* (Banerjee, Humphery-Jenner, and Nanda, 2015).

Panel B shows that our findings are not sensitive to the measurement of CEO overconfidence. Column 1 shows that the coefficient estimate on the continuous measure *CEO overconfidence* is positive and significant with the inclusion of firm strata. This result indicates that the degree of CEO option holdings affects the probability of gaining a dependent supplier. Next, we use a measure of overconfidence based on CEO insider transactions (*Overconfidence trade up*) following the approach in Kolasinski and Li (2013).¹⁵ Column 2 reports a positive and significant coefficient estimate on *Overconfidence trade up*, suggesting that the results are similar using the insider transaction measure. Together, these results suggest our findings are not sensitive to the measurement of CEO overconfidence.

3.1.2. *Do suppliers respond to the withdrawal of leadership?*

Our tests until this point focus on positive leadership outcomes as measured by an expansion in the dependent supplier network. A direct corollary of the leadership hypothesis is that suppliers should also react when the CEO withdraws leadership. Examining the withdrawal of leadership may provide a more powerful test if there are lingering concerns regarding construction of the overconfidence measures. We measure withdrawal of leadership using the *CEO overconfidence down* measure, which is the analog of the *CEO overconfidence up* measure, and replace the dependent variable with *Supplier*

¹⁵ See Section 2 and the Appendix for full details of the construction of this measure.

decrease, which is an indicator equal to one if the firm experiences a decrease in the number of suppliers from year $t-1$ to t , and zero otherwise.

The evidence in Columns 3 and 4 of Panel B suggests that suppliers respond strongly to the withdrawal of leadership by customer CEOs. Controlling for industry-year fixed effects in Column 3, the coefficient estimate on *CEO overconfidence down* is positive and statistically significant. This implies that a firm is more likely to lose dependent suppliers when an overconfident CEO withdraws leadership. The results are similar in Column 4 with the inclusion of firm strata.

3.1.3. *Initiations of supplier networks*

To help address a potential econometric issue that supplier relationships tend to be persistent, we examine the boundary scenario where firms attract their first dependent supplier. For example, Fee, Hadlock, and Thomas (2006) estimate a 76% probability that a customer-supplier relationship continues in the subsequent year.¹⁶ We find a similar pattern in our sample in which 75% of customer-supplier relationships continue into the next year.

To examine the initiations of supplier networks, we estimate our earlier logit regression specifications using *Start of supplier network* as the dependent variable. *Start of supplier network* is an indicator equal to one if a firm adds at least one dependent supplier in year t , but had no dependent supplier in year $t-1$.

[Insert Table 3]

The evidence in Table 3 shows a positive association between CEO overconfidence and the start of a supplier network. The coefficient estimate on *Confident CEO (options)* in Column 1 is positive and statistically significant, implying that firms led by overconfident CEOs are more likely to experience supplier network initiations. Column 2 reports similar results with the inclusion of firm characteristics as controls. Evaluated at the means of the other independent variables, *Confident CEO (options)* has a marginal

¹⁶ Holding control variables at their sample means in a logistic regression, Fee, Hadlock, and Thomas (2006) estimate that the probability of a relationship termination is about 24%.

effect of +0.6% on the likelihood of supplier network initiations. This represents an 18% increase over the average sample frequency of supplier initiation (3.3%).

Similar to Table 2, large, high sales growth, high R&D expenditure firms are more likely to initiate supplier networks. The estimated coefficients have weaker statistical significance potentially because *Start of supplier network* represents a boundary scenario and occurs less frequently than *Suppliers increase* in the sample. Nevertheless, the economic effect of *Confident CEO (options)* on supplier network initiation remains economically large.

The results are also similar using conditional logit models. Using the *CEO overconfidence up* measure, Column 3 shows that customers attract their first set of suppliers when the CEO provides strong leadership. Column 4 shows that the coefficient estimate on *CEO overconfidence up* remains significantly positive with the inclusion of firm strata. This suggests that our findings are not driven by the persistence of customer-supplier relationships.

The overall findings in Section 3.1 support the leadership hypothesis. The results indicate that CEO overconfidence has strong effects on the expansions and initiations of firms' supplier networks. Specifically, we find that the timing of leadership is important as supplier actions occur during times when the overconfident CEO provides leadership. Our results are also robust to various measures of CEO overconfidence and the use of industry-year and firm stratification in conditional logit models.

3.2. Do overconfident CEOs influence employee commitments?

Employees are arguably the most important stakeholders and their commitments can impact firm performance. We focus on employee commitment using 1) *employee turnover* and 2) *employee ownership* in company stock within benefit plans.

Employee turnover represents a leadership outcome because it reflects the ability of a company to retain productive labor and to induce employees' commitment to the vision of the CEO. The employee's decision to invest retirement funds in the company stock reflects a costly monetary commitment because it underdiversifies the employees'

retirement wealth¹⁷ (e.g., Meulbroek, 2005; Benartzi et al., 2007). We test the relation between CEO overconfidence and employee commitment by estimating Eq. (2) in fixed-effect panel Tobit models.

$$Employee\ Commitment_{i,t} = \alpha + \beta_1 Confident\ CEO\ (options)_{i,t-1} + \beta_2 \varphi_{i,t-1} + \varepsilon_{i,t} \quad (2)$$

Employee commitment is measured using *Employee turnover* or *Employee stock holdings* and φ represents a vector of control variables. We include industry-year fixed effects to control for unobserved heterogeneity across industries within the year. Since *Employee turnover* is bounded by zero, we estimate one-sided panel Tobit regressions and switch to two-sided panel Tobit regressions for *Employee stock holdings* (%).

[Insert Table 4]

The results in Table 4 show that employees are less likely to turn over when the CEO is overconfident. The significantly negative association between *Confident CEO (options)* and *Employee turnover* in Column 1 suggests that *Employee turnover* is lower under the leadership of an overconfident CEO. Since some firms only issue options to executives, we also perform the analysis using only firms with broad-based plans that include rank-and-file employees (e.g., Oyer and Schaefer, 2005; Hochberg and Lindsey, 2010; Chang et al., 2015). Column 2 shows that the results are similar using the sample of firms with broad-based plans. We also find that *Employee turnover* is lower at larger, more productive (higher ROA), growth firms (higher market-to-book) with longer-serving CEOs. This is consistent with the view that employees are likely to stay longer with firms that have stable management and good growth potential. As in Aldatmaz, Ouimet, and Van Wesep (2018), past stock performance is associated with higher *Employee turnover*. A potential explanation is that better stock performance partially reflects industry growth and better outside labor opportunities, inducing employees to switch employers (e.g., Oyer, 2004).

Retaining talent is particularly important in industries where human capital is valuable. Therefore, we examine whether the effects of CEO overconfidence are stronger

¹⁷ Meulbroek (2005) estimates that an employee who has 25% of her wealth allocated to company stock over a decade sacrifices 42% of the stock's market value on average.

in such industries. Following Coff (2002), we measure the human-capital intensity of an industry as its proportion of knowledge workers each year based on the distribution of occupation-types from the annual American Community Survey (ACS) census microdata (Ruggles et al., 2015).¹⁸ In Columns 3 and 4, we split the sample based on the sample median of the human-capital intensity measure. The effect of CEO overconfidence on employee turnover is stronger among firms in industries that are more reliant on human capital. The Welch-Satterthwaite *t*-test indicates a significantly larger coefficient estimate at the 10% level on *Confident CEO (options)* for high human-capital intensity firms than for low human-capital intensity firms in Column 4.

Next, we examine the employees' decision to hold employer stock in their benefit plan.¹⁹ We estimate Eq. (2) using *Employer stock holdings* and include controls for financial variables following Rauh (2006) and all control variables used in Table 4. Additionally, we add employee and plan-specific characteristics (i.e., cash/stock match, employee wealth) that may affect the employee's choice of company stock in retirement benefit plans. We also include firm age to proxy for workforce age because Ouimet and Zarutskie (2014) find that young firms disproportionately employ younger workers in firm-level U.S. Census data.

[Insert Table 5]

The results in Table 5 show that overconfident CEOs attract higher *Employee stock holdings*. Results from a two-sided panel Tobit regression in Column 1 show a positive and significant association between *Confident CEO (options)* and *Employee stock holdings (%)*. The coefficient estimate suggests that an overconfident CEO increases *Employee stock holdings (%)* by 2.7%. For reference, the employee benefit plans in our sample have

¹⁸ We obtain the 2004–2012 American Community Survey (ACS) census microdata from IPUMS-USA. The person-weighted database classifies occupations into seven broad categories, 1) management, professional, and related occupations, 2) service occupations, 3) sales and office occupations, 4) farming, fishing, and forestry occupations, 5) construction, extraction, and maintenance occupations, 6) production, transportation, and material moving occupations, and 7) military-specific occupations. Following the definition of professional workers in Coff (2002), we define knowledge occupations as the ones in Category (1) of the ACS database.

¹⁹ We collect data on benefit plans from filings of Form 5500 between year 2001 and year 2011 (see Section 2.3 for full details).

a mean (median) size of about \$419 million (\$102 million). We repeat this analysis on a subsample of benefit plans that have positive holdings of company stock in the preceding year to ensure that the results are not driven by plan-years with zero holdings of company stock. Column 2 shows that our results are robust in this subsample analysis. Columns 3 and 4 show that the results are similar using a per-employee dollar measure of stock ownership, *Employee stock holdings (\$)*. The coefficient estimates from Column 4 imply that CEO overconfidence increases the per-employee value of stock ownership in benefit plans by about \$1,570 annually. Economically, this represents a 21.3% increase relative to the sample mean of \$7,355.

Overall, the findings in this section are consistent with the view that overconfident CEOs attract greater employee commitment. They are better able to retain employees, particularly when talent is valuable. Their employees are also more likely to ‘buy-in’ by holding more company stock in their benefit plans.

3.3. Do overconfident CEOs attract greater supplier commitments?

In this section, we focus on supplier commitment along three dimensions. First, we examine the effect of CEO overconfidence in settings where relationship-specific investment (RSI) is valuable. Second, the supplier’s commitment to produce relationship-specific inputs should be revealed in higher supplier R&D investment. Therefore, we test whether the overconfidence of customer firm CEOs is associated with higher supplier R&D intensities. Third, terminations of customer-supplier relationships often stem from hold-up problems or contractual frictions (Fee, Hadlock, and Thomas, 2006). We test the leadership hypothesis by examining whether customer CEO overconfidence predicts lower risks of relationship terminations.

3.3.1. Importance of overconfident CEOs in industries requiring intensive relationship-specific investment

Businesses in certain industries are particularly reliant upon suppliers to create customized inputs for their final products. Supplier relationships in these industries are on average more valuable because these inputs are costlier to develop and more difficult to redeploy if the relationship terminates. Moreover, cultivating supplier relationships in

industries requiring intensive RSI may reflect a leadership outcome due to potential hold-up problems (e.g., Alchian, 1984; Tirole, 1988). For example, firms in durable goods manufacturing industries (henceforth, durable industries) produce more unique products that can only be sold to few customers (e.g., Titman, 1984; Titman and Wessels, 1988). To the extent that firms in durable industries require greater customized inputs from suppliers, CEO overconfidence may be particularly important in order to attract valuable suppliers onboard the firm and to induce greater supplier RSI.

We test this hypothesis by narrowing our focus to only firms in manufacturing industries and separately analyze durable and non-durable industries. We employ the same empirical tests used in Tables 2 and 3 for both expansion and initiation of a supplier network.

[Insert Table 6]

Panel A of Table 6 reports the results from logit regressions using the *Confident CEO (options)* measure of CEO overconfidence. As before, we include industry dummies and year dummies following Malmendier, Tate, and Yan (2011) and cluster standard errors by firm. The regressions include the full set of control variables as in Tables 2 and 3 but are suppressed to conserve space. Column 1 presents results from the durable industries, while Column 2 presents results from the non-durable industries.

The evidence suggests that the leadership of overconfident CEOs is more valuable for the expansion of a supplier network in durable manufacturing relative to non-durable manufacturing industries. The coefficient estimate on *Confident CEO (options)* is positive and statistically significant at the 10% level in durable industries, while it is negative and insignificant among the non-durable industries as shown in Column 2. We also find similar patterns for the initiation of a supplier network. Column 3 shows that, among the durable industries, the coefficient estimate on *Confident CEO (options)* is positive and significant at the 10% level indicating that overconfident CEOs are more likely to initiate a supplier network. Column 4 shows that there is no such relation in the non-durable industries. This evidence is consistent with the view that leadership by overconfident CEOs is particularly important when suppliers are particularly valuable.

The above analysis has weaker statistical power due to smaller subsamples and the broad classification of durable/non-durable manufacturing industries. To estimate a sharper test, we examine the timing of supplier actions using the *CEO overconfidence up* measure. In Panel B, we use the conditional logit model with firm stratification and find stronger links between CEO overconfidence and supplier network outcomes in the durable manufacturing industry. Column 1 shows that the loading on *CEO overconfidence up* is positive and significant among the durable industries, but is insignificant among the non-durable industries as shown in Column 2. Also, Column 3 shows a strong association between CEO overconfidence and the initiation of a supplier network among the durable industries, but not in the non-durable subsample in Column 4. We present firm strata specifications to capture unobserved firm heterogeneity although the patterns are similar when we employ industry-year strata to control for industry shocks.

The evidence is consistent with the view that the leadership of overconfident CEOs is more important when relationship-specific investment is particularly valuable to firm success. To more precisely evaluate the importance of RSI, we classify firms based on a measure of “contract intensity” within each industry. Contract intensity refers to the proportion of inputs in an industry that are not traded on an exchange nor reference-priced (Nunn, 2007). Therefore, industries with higher contract intensity have more customized goods. To keep the analysis comparable to the durable/non-durable analysis above, we split manufacturing firms into two groups based on the median industry-level contract intensity. We present results using both the *Confident CEO (options)* and *CEO overconfidence up* measures of CEO overconfidence.

The evidence suggests that the leadership of overconfident CEOs is particularly valuable in industries with high contract intensities. Panel C presents the results using *Confident CEO (options)*. Column 1 indicates that in industries with high contract intensities, the estimated coefficient on *Confident CEO (options)* is positive and statistically significant at the 10% level. This suggests that CEO overconfidence is associated with supplier network expansions in high contract-intensity industries. This association is not found among firms in low contract-intensity industries. The results are statistically stronger when examining supplier network initiation. Column 3 shows that in

high contract-intensity industries, overconfident CEOs are more likely to initiate supplier networks, while no such relation exists in low contract-intensity industries (Column 4).

The results are also similar when we use the *CEO overconfidence up* measure of CEO overconfidence. Panel D shows the coefficient estimates are positive and significant on *CEO overconfidence up* in high contract-intensity industries (Columns 1 and 3), but not in low contract-intensity industries (Columns 2 and 4).

Collectively, the evidence indicates that in manufacturing industries which require relationship-specific inputs, CEO overconfidence is particularly important in initiating and growing supplier networks. In these industries, the cost of supplier commitment tends to be higher and contracting imperfections are likely to be more salient, making leadership particularly important. Our evidence in this section supports this view.

3.3.2. *Supplier commitments: R&D intensity*

The previous section examines supplier commitments by exploiting variation in the degree of relationship-specific investment (RSI) across industries. In this section, we search for more direct evidence of RSI at the individual customer-supplier level.

We use two proxies for relationship-specific investment motivated from the existing literature: 1) *Supplier R&D intensity*, which captures the amount of research and development at the supplier firm (Kale and Shahrur, 2007; Raman and Shahrur, 2008), and 2) *Relationship duration*, which captures the likelihood of relationship termination at a given point in time (Fee, Hadlock, and Thomas, 2006).

Our first test uses *Supplier R&D intensity*, which we construct as follows. First, we scale the supplier's R&D activity by its relationship sales with the customer.²⁰ To allow for comparability across suppliers of different size, we normalize by total assets and apply the natural logarithm transformation to minimize the influence of outliers.²¹ Using this measure, we estimate OLS regressions following Eq. (3).

²⁰ For example, if a firm has \$100 in R&D activity and 40% of its sales are to a single customer, then we attribute \$100*0.4=\$40 to this particular relationship.

²¹ The results are both qualitatively and quantitatively similar using the raw values.

$$\text{Supplier R\&D Intensity}_{i,t} = a + \beta_1 \text{Confident CEO (options)}_{i,t-1} + \beta_2 \varphi_{i,t-1} + \beta_3 \gamma_{i,t-1} + \varepsilon_{i,t} \quad (3)$$

φ represents a vector of supplier control variables while γ represents a vector of customer control variables. All regression specifications include industry fixed effects to capture differences in R&D activity across industries and time fixed effects to capture macroeconomic trends. We cluster standard errors by customer-year because each customer may have multiple suppliers in a given year. The regressions include customer characteristics and supplier characteristics that are potentially associated with R&D activity.

[Insert Table 7]

The results in Table 7 suggest that suppliers of customer firms with overconfident CEOs produce greater *Supplier R&D intensity*. Column 1 shows a positive and significant coefficient estimate on *Confident CEO (options)* with the inclusion of customer firm characteristics. The results are similar after controlling for supplier firm characteristics in Column 2, alleviating concerns that supplier heterogeneity is driving our findings. The estimates from Column 2 suggest that *Confident CEO (options)* is associated with a 7.2% increase in supplier R&D investment. As a comparison benchmark, we log-transform the estimated coefficient on customer firm size (0.054) and find that a 100% increment in customer firm size increases supplier R&D investment by 5.4%. This suggests that the effect of *Confident CEO (options)* on *Supplier R&D intensity* is on the same order of magnitude as the effect of doubling firm size.

Consistent with Raman and Shahrur (2008) and Kale, Kedia, and Williams (2015), we find that suppliers tend to commit more R&D investments if the customer firm has more cash holdings and lower leverage. The negative association with investments of customer firms suggests a substitution between in-house production of inputs (i.e., vertical integration) and procurement from external suppliers (Banerjee, Dasgupta, and Kim, 2008).

Suppliers tend to have less bargaining power when they are smaller than their major customers. While locking in a larger customer may secure revenues, smaller suppliers run

the risk of ex post hold-up problems due to their relatively weaker bargaining positions. Therefore, smaller suppliers may require additional assurances, in particular when relationship-specific investments are required. CEO overconfidence can be a solution, albeit imperfect, to the hold-up problem. By holding in-the-money options, customer CEOs may be less likely to hold up the smaller suppliers because delays along the supply-chain may affect the CEO’s personal wealth substantially.

To test this hypothesis, we split our sample by the total assets of suppliers relative to that of their customers. A supplier is classified as small (big) if the *Supplier-customer size ratio* is lower (higher) than the SIC two-digit industry median value. The evidence presented in Column 3 suggests that overconfident CEOs strongly affect supplier R&D intensity among smaller suppliers, but have little impact among larger suppliers as shown in Column 4. Our evidence supports the view that the leadership of overconfident CEOs convinces smaller suppliers to make relationship-specific investments when hold-up problems are more likely to exist.

3.3.3. *Supplier commitments: relationship duration*

Another dimension of supplier commitment is the duration of the customer-supplier relationship which relates to contractual frictions that may arise from RSI-related issues or hold-up problems (Fee, Hadlock, and Thomas, 2006). If the leadership of overconfident CEOs helps to alleviate such frictions, we expect that these relationships are less likely to terminate, *ceteris paribus*.

To model the duration of customer-supplier relationships, we use survival analysis to estimate a hazard function. Specifically, we estimate the probability that a customer-supplier relationship terminates within a time interval, conditional on the survival of the relationship up till the beginning of that interval. Since a customer-supplier relationship may terminate and then restart in the future, we employ the Andersen-Gill extension of the standard Cox model.

We define *Overconfident CEO relationship* equal to one if a *Confident CEO (options)* customer CEO is present between the start and termination of the relationship, and zero

otherwise. *Overconfident CEO relationship* is re-computed for relationships with subsequent restarts and terminations. All other independent variables are based on relationship and firm characteristics in the first year of the customer-supplier relationship. Our definition of customer-supplier relationship duration follows Fee, Hadlock, and Thomas (2006) closely. The start of a relationship is defined to be the first year in which both customer and supplier are linked in the Compustat business segment file. We follow the relationship till the year in which the link is broken (termination year). If both firms are present in Compustat in the termination year, we determine that the customer-supplier relationship is terminated. If at least one of the firms disappears from Compustat in the termination year, the relationship is right-censored because we cannot determine if a relationship is terminated mutually. Since our sample period ends in 2011, we also classify all surviving relationships in 2011 as being right-censored. As it may take at least two years for a customer CEO to be classified as a *Confident CEO (options)* CEO, we only include relationships that begin on or after 1995. We report hazard ratios, which represents the ratio of hazard rates corresponding to two levels of the variable.²² For example, a hazard ratio above one (below one) implies that the variable increases (decreases) the chance of relationship termination at that point in time.

[Insert Table 8]

The results in Table 8 suggests that the supplier relationships of overconfident CEOs have statistically lower termination risk on average. Column 1 shows that the hazard of a relationship termination is about 10.8% lower in an *Overconfident CEO relationship* relative to a *Non-overconfident CEO relationship*. Column 2 shows that the effects are stronger with the inclusion of supplier and customer control variables which are unreported to conserve space. The hazard of a relationship termination is about 17.5% lower in an *Overconfident CEO Relationship* relative to a *Non-Overconfident CEO Relationship*. The results are similar after including *Supplier-customer size ratio*, *Supplier R&D intensity*, and *Sales dependency* as controls in Column 3. Consistent with Fee,

²² Likewise, the hazard ratio of a continuous variable represents the ratio of hazard rates corresponding to a unit change in the said variable.

Hadlock, and Thomas (2006), we find that customer-supplier relationships face lower termination risks when *Sales dependency* is higher.

Next, we verify that hazard proportionality holds in the model. Under hazard proportionality, the magnitude of treatment effects does not vary across time so interaction terms should not predict the termination risks of customer-supplier relationships.²³ Column 4 confirms that hazard proportionality is not violated in our model after adding interaction terms of all independent variables with time. *Overconfident CEO relationship* continues to predict a lower hazard of relationship termination (-29.2%), but the effect of its interaction with time is not statistically distinguishable from zero. This result corroborates our earlier findings and also implies that the effect of customer CEO leadership is stable throughout the relationship duration.

Finally, we restrict our sample to customer-supplier relationships without any restarts (i.e., only one termination). Column 5 shows that our conclusions remain unchanged. In sum, our findings suggest that leadership of overconfident CEOs helps to ease contractual frictions between customers and suppliers, resulting in more durable relationships.

4. Additional test and discussions

In this section, we provide additional evidence in support of the leadership hypothesis and discuss alternative explanations for our findings.

4.1. Do supplier commitments represent a dark side of CEO overconfidence?

The results in previous sections support a bright-side view that overconfident CEOs create valuable stakeholder relationships along various dimensions. However, if overconfident CEOs are over-optimistic and overpay for supplier commitments, these ‘achievements’ may ultimately hurt the company’s bottom line and sacrifice firm value. This would represent a dark-side view of CEO overconfidence.

²³ The standard Cox framework assumption of hazard proportionality implies that hazard ratios are constant through time. The Andersen-Gill extension relaxes this assumption, making it possible to include interaction terms between independent variables and time.

We create two tests to examine the possibility that overconfident CEOs overpay for their supplier commitments. First, we examine *Input costs* and *Markup percentage* which capture gross profitability and buying power (Fee and Thomas, 2004). Overpaying suppliers would result in higher input costs and lower markup percentages. Second, we examine stock returns to directly measure valuation effects. The dark-side view would imply that these supplier commitments come at a significant cost, destroying firm value and producing poor stock performance.

4.1.1. *Input costs and gross profitability*

We test whether overconfident CEOs overpay for their supplier commitments by estimating OLS regressions following Eq. (4).

$$Margin_{i,t} = a + \beta_1 Confident\ CEO\ (options)_{i,t-1} + \beta_2 \varphi_{i,t-1} + \varepsilon_{i,t}. \quad (4)$$

Margin represents either *Input costs* or *Markup percentage*. Following Fee and Thomas (2004), we measure buying power using *Input Costs* defined as the cost of goods sold (COGS) scaled by total sales. *Markup percentage* is defined as total sales less cost of goods sold (COGS), normalized by COGS. φ represents a vector of control variables. In all specifications, we include industry-year fixed effects to capture unobserved heterogeneity in *Input costs* or *Markup percentage* across industries within each year. Since we are interested in examining the possibility of overpaying suppliers, we only consider observations with at least one dependent supplier in each of both year t and year $t-1$. After applying this constraint, our subsample comprises 422 unique firms (2,244 firm-years) across 29 industries (two-digit SIC) from the years 1994 to 2011.

[Insert Table 9]

The results in Table 9 are inconsistent with the dark-side view. The estimates from Column 1 show that a firm led by an overconfident CEO has 2.3% lower *Input costs* relative to the sample mean ($-0.014/0.606 = -2.3\%$). While it may be difficult to benchmark overpayment of asset-specific inputs, overconfident CEO do not appear to pay more for input costs. Column 2 shows that controlling for input costs, firms led by overconfident CEOs command higher prices for their products. Economically, *Confident*

CEO (options) is associated with 15.0% higher *Markup percentage* relative to the sample mean ($0.164/1.093=15\%$).²⁴

Overall, the evidence is inconsistent with the view that supplier commitments acquired by overconfident CEOs destroy firm value. While we cannot definitively rule out overpayment for supplier commitments, the evidence generally points in the opposite direction. Firms led by overconfident CEOs tend to have greater buying power and gross profitability, as measured by *Input costs* and *Markup percentage*. These findings support the bright-side view that CEO overconfidence generates valuable leadership outcomes.

4.1.2. Stock performance of firms with dependent suppliers

In our second test, we use stock returns to examine the valuation effects of CEO overconfidence in firms with dependent suppliers. If overconfident CEOs secured supplier commitments at excessively high costs, this will likely generate poor future stock performance. We test the overpayment hypothesis by forming stock portfolios and tracking their performance using the following procedure. First, we restrict our sample to observations with at least one disclosed supplier at the fiscal year end. We allocate firms into either the *Confident* (*Non-confident*) portfolio if the CEO is (is not) overconfident at the fiscal year end. The firm stays in the portfolio for the next 12 months. We calculate the average returns of the stocks in the two portfolios each month and average alphas using the Fama-French five factors and the momentum factor.

[Insert Figure 1]

Fig. 1 presents the results. The evidence strongly rejects the overpayment hypothesis. Panel A presents the average monthly alphas of the *Confident* portfolio, *Non-Confident* portfolio, and the difference in returns between the two portfolios. The *Confident* portfolio yields a statistically significant and positive monthly alpha of 0.38%, while the alpha of the *Non-confident* portfolio is not statistically distinguishable from zero. The long-short portfolio yields a statistically significant and positive monthly alpha of 0.27%. Our results

²⁴ The sample means referred to in Columns 1 and 2 are computed only among firms with at least one supplier in the current year and at least one supplier in the previous year. Due to this constraint, the sample size in Table 9 is smaller than that in Table 2.

are also similar using characteristics-adjusted returns (Daniel, Grinblatt, Titman, and Wermers, 1997), and are available upon request. The evidence suggests that overconfident CEOs are unlikely to overpay suppliers.

We also present the factor loadings of the *Confident*, *Non-confident*, and the long-short portfolio in Panel B. Notably, both overconfident CEO and non-overconfident CEO firms have similar loadings on small-minus-big (SMB), suggesting that they have similar factor exposures to the size premium. However, the factor loadings on high-minus-low (HML) are statistically insignificant for firms led by overconfident CEOs but positive for their counterparts. This suggests that stocks in the *Non-confident* portfolio tend to be value firms. Firms in the *Non-confident* portfolio also tend to be more conservative in their corporate investment activities as they have a significantly positive loading on conservative-minus-aggressive (CMA). Firms in the *Confident* portfolio are not necessarily aggressive in their investment activities as the factor loading on CMA is close to zero (0.029).

Panels C and D examine manufacturing firms in the durable/non-durable industries and high-RSI/low-RSI industries, respectively. Firms in the *Confident* portfolios do not underperform in these subsamples. In the durable industries (Panel C) and high-RSI industries (Panel D), the *Confident* portfolio has a statistically significant and positive alpha. Overall, these patterns are inconsistent with the view that overconfident CEOs overpay for supplier inputs.

4.2. *Do suppliers anticipate expansionary corporate activities?*

Our evidence so far suggests that observable expansionary corporate activities do not explain our findings as we have included controls for investments, M&A activities, equity issuance, and R&D expenditures. However, we recognize that the inclusion of these controls does not rule out the possibility that suppliers anticipate future expansionary activities.

To address this, we employ a test that distinguishes the effect of overconfidence from corporate activity.²⁵ Malmendier and Tate (2005) find that overconfident CEOs have higher investment-cash flow sensitivity which implies that overconfident CEOs tend to restrain investment when cash flows are low. Therefore, if suppliers ignore overconfidence and focus solely on expansionary investment signals, the relation between overconfidence and supplier commitments is likely to be weaker in the low cash-flow subsample. We implement this test by re-estimating our main supplier test based on Eq. (1), splitting the sample based on high and low cash flow in year $t-1$.²⁶ We also include the Kaplan-Zingales index as a control for financial constraints (Kaplan and Zingales, 1997).

[Insert Table 10]

The evidence in Table 10 supports our leadership hypothesis as CEO overconfidence continues to attract supplier relationships in both low and high cash-flow firms. Columns 1 and 2 show a positive and significant relation between *Confident CEO (options)* and *Suppliers increase* in both the low and high cash-flow subsamples, respectively. The Welch-Satterthwaite t -test reveals that the coefficients on *Confident CEO (options)* in Columns 1 and 2 are not statistically different (p -value=0.663). These results suggest that supplier actions are not solely motivated by anticipated expansionary corporate activities. As Malmendier and Tate (2015) find a weaker relation between CEO overconfidence and investment-cash flow sensitivity during the Great Recession years (2008–2009), we show that our results are robust to the exclusion of those years in Columns 3 and 4.

Overall, the evidence suggests that the overconfidence of customer CEOs is important and directly influences suppliers beyond observable and anticipated expansionary corporate activities.

4.3. Does the manner of overconfidence matter?

The option holdings behavior of overconfident CEOs provides leadership through three potential channels: 1) communicating a strong belief in the firm’s prospects, 2)

²⁵ We thank the referee for suggesting this test.

²⁶ We measure cash flows following Chang, Dasgupta, Wong, and Yao (2014) and provide details in the Appendix.

leading by example, and 3) displaying commitment. Throughout the paper, we measure overconfidence using an options-based measure but overconfidence can also be demonstrated through positive public statements to the news media. Following Banerjee, Humphery-Jenner, and Nanda (2015), we use a news-based measure of overconfidence, *CEO media positivity*, to capture verbal communication of overconfidence. When the interests of both parties are perfectly aligned, such verbal communication may be sufficient to induce commitments from stakeholders since stakeholders deduce that the CEO has no incentives to lie (e.g., Farrell and Rabin, 1996).

To test whether the manner of leadership is important, we include *CEO media positivity* as a key independent variable and re-estimate our previous tests.

[Insert Table 11]

The results in Table 11 suggest that stakeholders do not respond to the verbal leadership measure of CEO overconfidence. Column 1 shows a positive but statistically insignificant association between *CEO media positivity* and *Suppliers increase*. The results are similar after including *Confident CEO (options)* in Column 2. Our tests on employee commitments paint a similar picture. Columns 3 and 4 show that the association between verbal leadership and employee turnover is negative but statistically insignificant. Using employee holdings of company stock, we find a positive effect of verbal leadership in Column 5 but this effect is statistically indistinguishable from zero. With the inclusion of the *Confident CEO (options)* measure in Column 6, the coefficient estimate on *CEO media positivity* remains statistically insignificant.

Our findings suggest that stakeholders appear to be influenced only by CEO leadership displayed through options-based overconfidence. A possible explanation is that stakeholders rely predominantly on verifiable behavior because the interests of the parties are not perfectly aligned.

4.4. *Do supplier actions drive CEO overconfidence?*

Under the leadership hypothesis, suppliers are more inclined to make costly commitments after they observe managerial overconfidence in customer firms. Motivated

by concerns of reverse causality, one may propose an alternative mechanism where supplier actions drive the overconfidence of the customer firm CEO.

The alternative mechanism is improbable for three reasons. First, the nature of the SFAS No. 131 reporting requirement is such that customer firms are generally significantly larger than their identified suppliers (Fee and Thomas, 2004). This size difference makes it unlikely that actions initiated by suppliers evoke managerial overconfidence in customer firms. Second, it is more economically intuitive that the customer firm CEO originates a project vision and subsequently finds suppliers to support it. Third, the alternative mechanism is not consistent with findings in Chu, Tian, and Wang (2017). In their study of corporate innovation along the supply chain, they find causal evidence that feedback flows from customers to suppliers. While we cannot categorically rule out the effect of supplier actions on customer firm CEO overconfidence, these reasons suggest reverse causality is unlikely to be a severe concern.

5. Conclusion

CEOs are hired for their vision and leadership talent. Yet, leadership has many dimensions and definitions, which makes the systematic analysis of this ability difficult. We hypothesize that corporate boards hire overconfident CEOs because they are better leaders. Our study provides a tractable empirical setting where stakeholder commitments represent valuable leadership outcomes. Leadership is distinct from formal authority because stakeholders' actions are voluntary.

Our evidence suggests that overconfident CEOs generate tangible leadership outcomes in the form of stakeholder commitments. Overconfident CEOs attract supplier relationships, particularly in durable goods manufacturing industries and high relationship-specificity industries, where such relationships are valuable. Suppliers are also more likely to provide relationship-specific products to the customer firm when the customer CEO is overconfident. The supplier relationships of overconfident CEOs have lower termination risk.

The leadership effects of CEO overconfidence also extend to employees. Overconfident CEOs are better able to retain talent, particularly in industries where human capital is more valuable. Their employees are also more likely to ‘buy-in’ as they hold more company stock in their benefit plan accounts. Together, these findings are consistent with the view that CEO overconfidence leads to greater employee commitments.

Overall, our results reveal a bright side of CEO overconfidence that is consistent with recent findings by Galasso and Simcoe (2011) and Hirshleifer, Low, and Teoh (2012). Our findings may also help to explain boards’ appointments of overconfident CEOs. Gervais, Heaton, and Odean (2011) show that boards are more likely to appoint overconfident CEOs at riskier, growth firms creating potential benefits for both the firm and manager. Boards may be more likely to appoint overconfident CEOs at firms where developing good stakeholder relationships and attracting key commitments are especially important. Our evidence suggests that overconfident CEOs are better able to deliver these leadership outcomes. Future research may explore this and other additional positive dimensions of CEO overconfidence.

References

- Ahmed, A., Duellman, S., 2013. Managerial overconfidence and accounting conservatism. *Journal of Accounting Research* 51, 1–30.
- Alchian, A., 1984. Specificity, specialization, and coalitions. *Journal of Institutional and Theoretical Economics* 140, 34–49.
- Aldatmaz, S., Ouimet, P., Van Wesep, E. D., 2018. The option to quit: the effect of employee stock options on turnover. *Journal of Financial Economics* 127, 163–151.
- Almazan, A., Chen, Z., Titman, S., 2017. Firm investment and stakeholder choices: a top-down theory of capital budgeting. *Journal of Finance* 72, 2179–2228.
- Anderson, C., Brion, S., Moore, D., Kennedy, J., 2012. A status-enhancement account of overconfidence. *Journal of Personality and Social Psychology* 103, 718–735.
- Andersen, K., Gill, R., 1982. Cox's regression model for counting processes: a large sample study. *The Annals of Statistics* 10, 1100–1120.
- Babenko, I., Sen, R., 2014. Money left on the table: An analysis of participation in employee stock purchase plans. *Review of Financial Studies* 27, 3658–3698.
- Banerjee, S., Dasgupta, S., Kim, Y., 2008. Buyer-supplier relationships and the stakeholder theory of capital structure. *Journal of Finance* 63, 2507–2552.
- Banerjee, S., Humphery-Jenner, M., Nanda, V., 2015. Restraining overconfident CEOs though improved governance: evidence from the Sarbanes-Oxley Act. *Review of Financial Studies* 28, 2812–2858.
- Banerjee, S., Humphery-Jenner, M., Nanda, V., Tham, M., 2017. Executive overconfidence and securities class actions. *Journal of Financial and Quantitative Analysis* Forthcoming.
- Benartzi, S., Thaler, R., Utkus, S., Sunstein, C., 2007. The law and economics of company stock in 401 (k) plans. *Journal of Law and Economics* 50, 45–79.
- Billett, M., Qian, Y., 2008. Are overconfident CEOs born or made? Evidence of self-attribution bias from frequent acquirers. *Management Science* 6, 1037–1051.
- Campbell, T., Gallmeyer, M., Johnson, S., Rutherford, J., Stanley, B., 2011. CEO optimism and forced turnover. *Journal of Financial Economics* 101, 695–712.
- Carter, M., Lynch, L., 2004. The effect of stock option repricing on employee turnover. *Journal of Accounting and Economics* 37, 91–112.
- Chang, X., Dasgupta, S., Wong, G., Yao, J., 2014. Cash-flow sensitivities and the allocation of internal cash flow. *Review of Financial Studies* 27, 3628–3657.
- Chang, X., Fu, K., Low, A., Zhang, W., 2015. Non-executive employee stock options and corporate innovation. *Journal of Financial Economics* 115, 168–188.
- Chu, Y., Tian, X., Wang, W., 2017. Learning from customers: corporate innovation along the supply chain. *Management Science* Forthcoming.
- Coff, R., 2002. Human capital, shared expertise, and the likelihood of impasse in corporate acquisitions. *Journal of Management* 28, 107–128.
- Daniel, K., Grinblatt, M., Titman, S., Wermers, R., 1997. Measuring mutual fund performance with characteristic-based benchmarks. *Journal of Finance* 52, 1035–1058.
- Dyer, J., Nobeoka, K., 2000. Creating and managing a high performance knowledge-sharing network: the Toyota case. *Strategic Management Journal* 21, 345–367.

- Ellis, J., Fee, C., Thomas, S., 2012. Proprietary costs and the disclosure of information about customers. *Journal of Accounting Research* 50, 685–727.
- Farrell, F., Rabin, M., 1996. Cheap talk. *Journal of Economic Perspectives* 10, 103–118.
- Fee, C., Thomas, S., 2004. Sources of gains in horizontal mergers: evidence from customer, supplier, and rival firms. *Journal of Financial Economics* 74, 423–460.
- Fee, C., Hadlock, C., Thomas, S., 2006. Corporate equity ownership and the governance of product market relationships. *Journal of Finance* 61, 1217–1251.
- Galasso, A., Simcoe, T., 2011. CEO overconfidence and innovation. *Management Science* 57, 1469–1484.
- Gervais, S., Heaton, J., Odean, T., 2011. Overconfidence, compensation contracts, and capital budgeting. *The Journal of Finance* 66, 1735–1777.
- Goel, A., Thakor, A., 2008. Overconfidence, CEO selection, and corporate governance. *Journal of Finance* 63, 2737–2784.
- Hall, B., Murphy, K., 2002. Stock options for undiversified executives. *Journal of Accounting and Economics* 33, 3–42.
- Hermalin, B., 1998. Toward an economic theory of leadership: leading by example. *American Economic Review* 88, 1188–1206.
- Hirshleifer, D., Low, A., Teoh, S., 2012. Are overconfident CEOs better innovators? *Journal of Finance* 67, 1457–1498.
- Hochberg, Y., Lindsey, L., 2010. Incentives, targeting, and firm performance: an analysis of non-executive stock options. *Review of Financial Studies* 23, 4148–4186.
- Kale, J., Shahrur, H., 2007. Corporate capital structure and the characteristics of suppliers and customers. *Journal of Financial Economics* 83, 321–365.
- Kale, J., Kedia, S., Williams, R., 2015. The effect of CEO's risk-taking incentives on relationship-specific investments by customers and suppliers. Unpublished working paper. Northeastern University.
- Kang, M., Mahoney, J., Tan, D., 2009. Why firms make unilateral investments specific to other firms: the case of OEM suppliers. *Strategic Management Journal* 30, 117–135.
- Kaplan, S., Zingales, L., 1997. Do investment–cash flow sensitivities provide useful measures of financing constraints? *Quarterly Journal of Economics* 1, 169–215.
- Kennedy, J., Anderson, C., Moore, D., 2013. When overconfidence is revealed to others: testing the status-enhancement theory of overconfidence. *Organizational Behavior and Human Decision Processes* 122, 266–279.
- Klein, B., Crawford, R., Alchian, A., 1978. Vertical integration, appropriable rents, and the competitive contracting process. *Journal of Law and Economics* 21, 297–326.
- Kolasinski, A., Li, X., 2013. Can strong boards and trading their own firm's stock help CEOs make better decisions? Evidence from acquisitions by overconfident CEOs. *Journal of Financial and Quantitative Analysis* 48, 1173–1206.
- Komai, M., Stegeman, M., Hermalin, B., 2007. Leadership and information. *American Economic Review* 97, 944–947.
- Loughran, T., Ritter, J., 2004. Why has IPO underpricing changed over time? *Financial Management* 3, 5–37.
- Malmendier, U., Tate, G., 2005. CEO overconfidence and corporate investment. *Journal of Finance* 60, 2661–2700.

- Malmendier, U., Tate, G., 2008. Who makes acquisitions? CEO overconfidence and the market's reaction. *Journal of Financial Economics* 89, 20–43.
- Malmendier, U., Tate, G., Yan, J., 2011. Overconfidence and early-life experiences: the effect of managerial traits on corporate financial policies. *Journal of Finance* 66, 1687–1733.
- Malmendier, U., Tate, G., 2015. Behavioral CEOs: the role of managerial overconfidence. *Journal of Economic Perspectives* 29, 37–60.
- Meulbroeck, L., 2005. Company stock in pension plans: how costly is it? *Journal of Law and Economics* 48, 443–474.
- Nunn, N., 2007. Relationship-specificity, incomplete contracts, and the pattern of trade. *Quarterly Journal of Economics* 122, 569–600.
- Ouimet, P., Zarutskie, R., 2014. Who works for startups? The relation between firm age, employee age, and growth. *Journal of Financial Economics* 112, 386–407.
- Oyer, P., 2004. Why do firms use incentives that have no incentive effects? *Journal of Finance* 59, 1619–1650.
- Oyer, P., Schaefer, S., 2005. Why do some firms give stock options to all employees? An empirical examination of alternative theories. *Journal of Financial Economics* 76, 99–133.
- Parmigiani, A., 2007. Why do firms both make and buy? An investigation of concurrent sourcing. *Strategic Management Journal* 28, 285–311.
- Raman, K., Shahrur, H., 2008. Relationship-specific investments and earnings management: evidence on corporate suppliers and customers. *The Accounting Review* 83, 1041–1081.
- Rauh, J., 2006. Own company stock in defined contribution pension plans: a takeover defense? *Journal of Financial Economics* 81, 379–410.
- Ruggles, S., Genadek, K., Goeken, R., Grover, J., Sobek, M., 2015. Integrated public use microdata series: version 6.0 (machine-readable database). University of Minnesota.
- Schrand, C., Zechman, S., 2012. Executive overconfidence and the slippery slope to financial misreporting. *Journal of Accounting and Economics* 53, 311–329.
- Tirole, J., 1988. *The Theory of Industrial Organization*. The MIT Press, Cambridge, MA.
- Titman, S., 1984. The effect of capital structure on a firm's liquidation decision. *Journal of Financial Economics* 13, 137–151.
- Titman, S., Wessels, R., 1988. The determinants of capital structure choice. *Journal of Finance* 43, 1–19.

Appendix A. Variable definitions

2-Year stock returns	Buy-and-hold returns of firm stock from beginning of year $t-1$ to end of year t . Source: CRSP
Acquisitions	Indicator variable that equals one if the firm performs at least one acquisition within the $(t-1, t)$ window, and equals zero otherwise. Source: SDC Platinum
Annual employer cash match	Indicator variable that equals one if the firm offers matching cash contributions to its employees in the previous year, and equals zero otherwise. Source: Item (2a1a) from Schedule H of Form 5500
Annual employer stock match	Indicator variable that equals one if the firm offers matching non-cash contributions to its employees in the previous year, and equals zero otherwise. Source: Item (2a2) from Schedule H of Form 5500
Broad-based plan	An employee stock option plan is broad-based if the number of options granted to non-executive employees exceeds 0.5% of the number of shares outstanding. Source: Oyer and Schaefer (2005), Hochberg and Lindsey (2010), Chang et al. (2015), ExecuComp, Compustat
Cash	Cash holdings of firm, scaled by total assets. Source: Compustat
CEO media positivity	Indicator variable that equals one if the number of articles containing references to confidence is more than the number of articles containing references to non-confidence during the year, and equals zero otherwise (Banerjee, Humphery-Jenner, and Nanda, 2015). Source: Banerjee, Humphery-Jenner, and Nanda (2015)
CEO overconfidence	Average value of the CEO's options scaled by the average strike price. The numerator is the value of the CEO's vested and unexercised options (ExecuComp: opt_unex_exer_val) scaled by the number of such options (ExecuComp: opt_unex_exer_num). The denominator is the difference between the firm's stock price at the end of the fiscal year (Compustat: prcc_f) and the numerator. Source: Campbell et al. (2011), Malmendier, Tate, and Yan (2011), Hirshleifer et al. (2012), ExecuComp, CRSP
CEO overconfidence up	Indicator variable that equals one if CEO overconfidence in year t is in the top quartile of the sample and if CEO overconfidence in year $t-1$ is not in the top quartile of the sample, and equals zero otherwise. Source: ExecuComp, CRSP
CEO stock ownership	CEO's percentage share ownership. Source: ExecuComp
CEO tenure	If the date of appointment as CEO is available in ExecuComp, variable equates to the number of years elapsed since the appointment date. Otherwise, variable equates to the number of years elapsed since the earliest date where the CEO first appears in the database. Source: ExecuComp
Confident CEO (options)	Indicator variable that equals one if a CEO's vested option holdings are at least 67% in-the-money on at least two instances (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013), and equals zero otherwise. Variable switches from zero to one from the first such instance. Source: ExecuComp, CRSP
Dividends-to-price ratio	Ratio of common equity dividends (dvc) to market value of equity (csho*prcc_f). Compustat data items are contained in parentheses. Source: Compustat
Durable firm	Industries whose SIC codes are between 3400 and 3999. Source: Titman and Wessels (1988)
Employee benefit plan assets	Total assets of the employee benefit plan. Source: Item (1f) from Schedule H of IRS Form 5500.
Employee stock holdings	We extract the dollar value of employer securities held in employee benefit plans from item (1d1) of the firm's Form 5500 Schedule H. As a simplification, we assume that employer securities in employee benefit plans only comprise the common stock of the firm. Employee stock holdings is the dollar value of employer stock held in the benefit plan scaled by the number of active plan participants (item 6a2 of Form 5500 Annual Return/Report) in the year. We also construct a variant that is scaled by the total assets of the benefit plan (item 1f of Schedule H) in the year. Source: Form 5500 Annual Return/Report, Schedule H of Form 5500

Appendix A. (Continued)

Employee turnover	The ratio of the number of non-executive options that are cancelled, terminated, forfeited, expired, or lapsed in the year to the number of non-executive options outstanding at the beginning of the year. We construct the number of non-executive options outstanding by subtracting the number of exercisable and non-exercisable executive options outstanding (ExecuComp: opt_unex_exer_num, opt_unex_unexer_num) from the total number of options outstanding (Compustat: optosey). The construction of Employee turnover begins in 2004 when Compustat begins coverage of option data. Compustat does not provide a breakdown of the quantities of options that are cancelled, terminated, forfeited, expired, and lapsed. Instead, it aggregates these quantities into a single variable – OPTCA. We apply a logarithmic transformation to Employee turnover. Source: Carter and Lynch (2004), ExecuComp, Compustat
Financial constraints	Kaplan-Zingales index. Definitions are from Kaplan and Zingales (1997). $-1.001909 (CF/Capital) + 0.2826389 (Q) + 3.139193 (Leverage) - 39.3678 (Dividend/Capital) - 1.314759 (Cash / Capital)$. Source: Kaplan and Zingales (1997)
Firm age	The number of years elapsed since the founding year of the firm. Data on the founding years of firms are publicly available on Jay Ritter’s website and are used in Loughran and Ritter (2004). If the founding year of the firm is not available, Firm age is the number of years elapsed since the earliest year where the firm is covered in Compustat and has a non-missing stock price. Source: Loughran and Ritter (2004), Compustat
Industry relationship-specificity	Proportion of inputs used in industry that are neither traded on organized exchanges nor reference-priced. Source: Nunn (2007)
Input costs	Ratio of cost of goods sold to total sales. Source: Compustat
Investment	Capital expenditure of firm, scaled by total assets. Source: Compustat
Leverage	Sum of long-term debt and short-term debt, scaled by total assets. Source: Compustat
High cash flow	Cash flow is the sum of income before extraordinary items (ibc), extraordinary items and discontinued operations (xidoc), depreciation and amortization (dpc), deferred taxes (txdc), equity in net loss (esubc), gains in sale of property, plant & equipment (PP&E) and investment (sppiv), other funds from operation (fopo), the exchange rate effect (exre), and less changes in working capital (wcapc), scaled by total assets (at). Compustat data items are contained in parentheses. A firm is marked as High cash flow if its Cash flow is higher than the year median. Source: Chang et al. (2014), Compustat
High human capital intensity	The ratio of knowledge workers to the total number of workers in each industry-year. The distribution of worker-types (or occupations) is obtained from the annual American Community Survey (ACS) census microdata. Data in the ACS census microdata are person-weighted. Knowledge workers are individuals whose occupations belong to the Management, Professional, and Related Occupations class. An industry is marked as High human capital if its Human capital ratio is higher than the sample median. Source: Integrated Public Use Microdata Series: Version 6.0
High RSI	An industry is marked as High RSI if its industry relationship-specificity values are higher than the sample median. Source: Nunn (2007)
Market-to-book ratio	Ratio of market value of equity to book value of equity. Source: Compustat
Markup percentage	Ratio of the difference between total sales and cost of goods sold to cost of goods sold. Source: Compustat
Missing R&D	Indicator variable that equals one if R&D expenses are missing in the Compustat database, and equals zero otherwise. Source: Compustat
Non-durable firm	Industries whose SIC codes are between 2000 and 3399. Source: Titman and Wessels (1988)
Overconfidence (OC) trade	Indicator variable that equals one in year t if the CEO’s stock purchases over the next two years have negative 180-day buy-and-hold abnormal returns (BHARs) on average, and equals zero otherwise. Source: Kolasinski and Li (2013)

Appendix A. (Continued)

Overconfidence trade up	Indicator variable that equals one if OC trade (see above) is zero in year $t-1$ and is unity in year t , and equals zero otherwise. Source: Kolasinski and Li (2013), Thomson Reuters Insiders
Overconfident CEO relationship	For a given customer-supplier relationship, indicator variable that equals one if a Confident CEO (options) customer CEO is incumbent in any year between the start (inclusive) and end (exclusive) of the relationship, and equals zero otherwise. Source: ExecuComp, CRSP, Compustat
Per-employee non-stock wealth	Total assets of benefit plan less Employee stock holdings, scaled by the number of active plan participants in the year. Source: Form 5500 Annual Return/Report, Schedule H of Form 5500
Price-to-earnings ratio	Ratio of market value of equity (csho*prcc_f) to net income (ibc). Compustat data items are contained in parentheses. Source: Compustat
R&D intensity	R&D expenses of customer firm, scaled by total assets. Missing R&D expenses are set to zero. Source: Compustat
Return on assets	Ratio of net income to assets. Source: Compustat
ROA volatility	3-year standard deviation of ROA from year $t-2$ to year t . Source: Compustat
Sales growth	Difference between sales in year t and year $t-1$, scaled by sales in year $t-1$. Source: Compustat
SEO proceeds	Total dollar value of seasoned equity offering (SEO) proceeds in the year, scaled by lagged total assets of the firm. Source: SDC Platinum, Compustat.
Start of supplier network	Indicator variable that equals one if firm has at least one dependent supplier in year t and has no dependent supplier in year $t-1$, and equals zero otherwise. Source: Compustat
Supplier R&D intensity	For each supplier-customer pair in a year, we first compute the supplier's R&D activity that is attributable to the customer by multiplying supplier's R&D expenses by Supplier sales dependency (see above for definition). Thereafter, we normalize the resulting value by the total assets of the supplier. To facilitate presentation, values are inflated by a factor of 1,000. Source: Compustat
Supplier sales dependency	The ratio of supplier-customer sales to total supplier sales for each supplier-customer pair in a year. Source: Compustat
Supplier-customer size ratio	For each customer-supplier pair in a year, Supplier-customer size ratio is computed as the ratio of total assets of the supplier to that of the customer. Supplier-customer size ratio is defined to be small (big) if it is lower (higher) than the SIC two-digit industry median value. Source: Compustat
Suppliers decrease	Indicator variable that equals one if the number of dependent suppliers in year t is lower than the number of dependent suppliers in year $t-1$, and equals zero otherwise. Source: Compustat
Suppliers increase	Indicator variable that equals one if the number of dependent suppliers in year t is higher than the number of dependent suppliers in year $t-1$, and equals zero otherwise. Source: Compustat
Total assets	Total assets of firm. Source: Compustat
Working capital	Difference between total current assets (act) and total current liabilities (lct), scaled by total assets (at). Compustat data items are contained in parentheses. Source: Compustat

Fig. 1. Monthly portfolio alphas. The figure reports monthly alphas of portfolios of firms led by overconfident/ non-overconfident CEOs. Our sample includes firms that are covered in the ExecuComp database and business segment files of Compustat. We remove utilities (Standard Industry Classification (SIC): 4000 – 4999) and financial firms (SIC: 6000 – 6999). Our sample spans the years 1993–2011 and is restricted to observations with at least one supplier in each of both year t and year $t-1$. We allocate each firm to the *Confident* (*Non-confident*) portfolio if the CEO is (is not) identified as *Confident CEO (options)* at the end of its fiscal year. The firm remains in the portfolio over the next 12 months. We calculate average returns of the stocks in each portfolio and benchmark them against the Fama-French five factors and the momentum factor. *Confident CEO (options)* is a dummy variable which equals one if the CEO’s vested option holdings are at least 67% in-the-money on at least two instances (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013), and equals zero otherwise. *Confident CEO (options)* switches from zero to one from the first such instance. Detailed variable definitions are contained in Appendix A. Panel A presents the monthly alphas of the *Confident* portfolio, the *Non-confident* portfolio, and the long-short portfolio (*Diff*). Panel B presents the loadings and standard errors of each portfolio on the six risk factors. Panel C repeats the monthly alpha analysis separately for firms in the *durable/non-durable* industries. Following Titman and Wessels (1988), *durable* industries have SIC codes 3400 – 3999 and *non-durable* industries have SIC codes 2000 – 3399. Panel D separates firms in *high-RSI/low-RSI* industries. We determine industry-level relationship specificity according to the year 1997 contract intensity variable from Nunn (2007). An industry is classified as *high (low)-RSI* if its industry contract intensity is above (below) the sample median. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Portfolio monthly alphas of firms led by overconfident/non-overconfident CEOs and with at least one supplier

Panel B. Portfolio loadings on Fama-French five factors + momentum factor

	(1)	(2)	(3)
Portfolio	Confident	Non-confident	Difference
α	0.381*** (0.127)	0.116 (0.120)	0.265* (0.146)
MKT – RF	1.091*** (0.032)	1.076*** (0.031)	0.015 (0.037)
SMB	0.469*** (0.046)	0.468*** (0.043)	0.000 (0.053)
HML	-0.107 (0.065)	0.125** (0.062)	-0.232*** (0.075)
RMW	0.155** (0.063)	0.236*** (0.060)	-0.082 (0.073)
CMA	0.029 (0.076)	0.157** (0.072)	-0.128 (0.087)
MOM	-0.232*** (0.024)	-0.260*** (0.023)	0.028 (0.028)
Months	234	234	234
R -squared	0.909	0.906	0.294

Panel C. Portfolio monthly alphas of firms led by overconfident/non-overconfident CEOs in durable/non-durable industries and with at least one supplier

Panel D. Portfolio monthly alphas of firms led by overconfident/non-overconfident CEOs in high/low-RSI industries and with at least one supplier

Table 1

Descriptive statistics

Panel A presents means, standard deviations, percentiles, and the number of observations for variables used in our baseline regression in Column 2 of Table 2 Panel A. Panel B compares the variable means of firm-year observations of customer firms and non-customer firms. Our sample spans the years 1993–2011 and includes firms that are covered in the ExecuComp database and business segment files of Compustat. We remove utilities (SIC: 4000 – 4999) and financial firms (SIC: 6000 – 6999). Detailed variable definitions are contained in Appendix A. Financial variables are winsorized at the 1st and 99th percentiles to reduce the influence of outliers. ***, **, * represent statistical significance of differences in means at the 1%, 5%, and 10% levels, respectively.

Table 1 (continued)*Panel A: Summary statistics for sample firms*

	N	Mean	S.D.	P10	P25	P50	P75	P90
1. CEO overconfidence	14745	0.321	0.262	0.000	0.075	0.289	0.516	0.703
2. Confident CEO (options)	14745	0.533	0.499	0	0	1	1	1
3. Overconfident CEO relationship	4974	0.570	0.495	0	0	1	1	1
4. CEO overconfidence up	14179	0.068	0.253	0	0	0	0	0
5. CEO media positivity	7923	0.810	0.392	0	1	1	1	1
6. Suppliers increase	14745	0.072	0.258	0	0	0	0	0
7. Suppliers decrease	14745	0.075	0.264	0	0	0	0	0
8. Start of supplier network	14745	0.033	0.180	0	0	0	0	0
9. Total assets	14745	4088	8760	190	408	1066	3225	10219
10. Market-to-book ratio	14745	3.204	2.904	1.075	1.580	2.365	3.703	5.947
11. Return on assets	14745	0.053	0.108	-0.052	0.019	0.062	0.105	0.158
12. ROA volatility	14745	0.054	0.069	0.007	0.014	0.030	0.064	0.125
13. Sales growth	14745	0.124	0.269	-0.123	0.001	0.087	0.198	0.387
14. Investment	14745	0.066	0.068	0.013	0.025	0.045	0.081	0.138
15. R&D intensity ²⁷	14745	0.043	0.068	0	0	0.010	0.062	0.134
16. Leverage	14745	0.224	0.201	0	0.050	0.202	0.335	0.464
17. Cash	14745	0.121	0.144	0.009	0.024	0.070	0.164	0.298
18. 2-Year stock returns	14745	0.373	1.279	-0.474	-0.186	0.164	0.605	1.236
19. CEO tenure	14745	1.751	0.855	0.693	1.099	1.792	2.398	2.833
20. CEO stock ownership	14745	2.021	4.917	0	0	0.21	1.35	5.63
21. Acquisitions indicator	14745	0.355	0.478	0	0	0	1	1
22. SEO proceeds	14745	0.006	0.037	0	0	0	0	0
23. Employee turnover	7742	0.094	0.097	0.010	0.027	0.063	0.126	0.220
24. Employee benefit plan (\$M)	7887	419.390	1224.709	13.994	34.966	102.420	302.880	879.596
25. Total-employee stock held (%)	3577	0.139	0.127	0.011	0.043	0.108	0.196	0.306
26. Per-employee stock held (\$1000)	3577	7.355	7.704	0.271	1.318	4.481	10.966	19.909
27. Durable indicator	14745	0.336	0.472	0	0	0	1	1
28. Non-durable indicator	14745	0.270	0.444	0	0	0	1	1
29. High-RSI indicator	8360	0.531	0.499	0	0	1	1	1
30. Customer-supplier size ratio	12002	0.081	0.386	0.001	0.002	0.008	0.039	0.146
31. Supplier R&D intensity	12002	23.213	48.114	0.654	2.083	7.652	21.417	52.577
32. Input costs	2244	0.606	0.218	0.269	0.451	0.657	0.769	0.865
33. Markup percentage	2244	1.093	1.481	0.155	0.300	0.522	1.219	2.723

²⁷ Missing values of *R&D* are assigned to be zero. Reported statistics exclude missing *R&D* values.

Table 1 (continued)*Panel B: Comparison between customer firm and stand-alone firm observations*

	Customer firm		Standalone firm		Difference
	N	Mean	N	Mean	
Confident CEO (options)	2904	0.545	11841	0.531	0.014
CEO overconfidence	2904	0.343	11841	0.315	0.028***
CEO overconfidence up	2804	0.065	11375	0.069	-0.004
Total assets	2904	12394	11841	2051	10343***
Market-to-book ratio	2904	3.824	11841	3.052	0.771***
Return on assets	2904	0.066	11841	0.050	0.015***
ROA volatility	2904	0.046	11841	0.055	-0.009***
Sales growth	2904	0.132	11841	0.121	0.010**
Investment	2904	0.071	11841	0.065	0.006***
R&D intensity	2904	0.048	11841	0.042	0.006***
Leverage	2904	0.246	11841	0.219	0.026***
Cash	2904	0.098	11841	0.127	-0.029***
2-Year stock returns	2904	0.407	11841	0.365	0.042
CEO tenure	2904	0.930	11841	0.959	-0.028***
CEO stock ownership	2904	1.123	11841	2.241	-1.118***

Table 2

The effect of CEO overconfidence on the increase/decrease in suppliers

This table presents results from logit and conditional logit regressions. *Suppliers increase* is a dummy variable that equals to one if a firm has more dependent suppliers in year t than it has in $t-1$, and equals to zero otherwise. *Suppliers decrease* is defined symmetrically to *Suppliers increase*. Dependent suppliers are defined according to SFAS 14. Columns 1 and 2 present results from a logit model while Columns 3 and 4 present results from a conditional logit model. *Confident CEO (options)* is a dummy variable which equals one if the CEO's vested option holdings are at least 67% in-the-money on at least two instances (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013), and equals zero otherwise. *Confident CEO (options)* switches from zero to one from the first such instance. *CEO overconfidence up* is a dummy variable that equals one if the continuous *CEO overconfidence* measure in year t is in the top quartile of the sample but not in the top quartile of the sample in year $t-1$, and equals zero otherwise. *CEO overconfidence down* is defined symmetrically to *CEO overconfidence up*. CEO overconfidence is the per-option value of CEO's vested and unexercised options scaled by average strike price (Campbell et al., 2011; Malmendier, Tate, and Yan, 2011; Hirshleifer, Low, and Teoh, 2012). *Overconfidence trade up* is an indicator that switches on if *Overconfidence trade* is equal to zero in year $t-1$ and equal to one in year t . *Overconfidence trade* is a dummy variable that equals one in year t if the CEO's purchases over the next two years have negative 180-day BHARs on average, and equals zero otherwise (Kolasinski and Li, 2013). BHARs are benchmarked against returns of Fama-French size-decile portfolios. The sample in Column 2 of Panel B begins from year 1996 (inclusive) onwards because the *Overconfidence trade* measure starts in 1996. Detailed variable definitions are in Appendix A. Table 1 contains a description of the sample. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 2 (continued)*Panel A: Baseline specifications of supplier commitment*

	(1)	(2)	(3)	(4)
	LOGIT: Suppliers increase	LOGIT: Suppliers increase	C-LOGIT: Suppliers increase	C-LOGIT: Suppliers increase
Confident CEO (options) $t-1$	0.261*** (0.085)	0.165** (0.080)		
CEO overconfidence up $t-1$			0.516*** (0.114)	0.368*** (0.122)
CEO stock ownership $t-1$		-0.006 (0.010)	-0.008 (0.011)	0.004 (0.013)
Log (Total assets $t-1$)		0.798*** (0.034)	0.810*** (0.042)	-0.113* (0.067)
Leverage $t-1$		-0.391* (0.210)	-0.400* (0.230)	0.045 (0.305)
Market-to-book ratio $t-1$		0.042*** (0.014)	0.040** (0.016)	0.013 (0.009)
Return on assets $t-1$		0.234 (0.439)	0.415 (0.380)	0.648 (0.607)
Sales growth $t-1$		0.317** (0.148)	0.411*** (0.139)	0.391** (0.187)
R&D intensity $t-1$		5.049*** (0.748)	4.895*** (0.637)	4.035*** (0.817)
Missing R&D $t-1$		0.229 (0.152)	0.202 (0.144)	0.348 (0.611)
Log (1 + CEO tenure $t-1$)		-0.025 (0.109)	0.017 (0.106)	0.001 (0.107)
Cash $t-1$		0.282 (0.346)	0.311 (0.336)	-0.584 (0.455)
ROA volatility $t-1$		0.914 (0.676)	0.926* (0.500)	0.392 (0.620)
Log (2-Year stock returns $t-1$)		-0.022 (0.075)	-0.009 (0.103)	0.036 (0.100)
Investment $t-1$		1.192* (0.669)	1.048 (0.947)	1.922*** (0.739)
Acquisitions $t-1$		0.014 (0.077)	0.006 (0.083)	-0.057 (0.089)
SEO proceeds $t-1$		-0.449 (1.076)	-0.875 (0.765)	-1.715* (0.897)
Num. suppliers $t-1$		-0.015** (0.006)	-0.015** (0.007)	-0.126*** (0.031)
Industry dummies	Yes	Yes	No	No
Year dummies	Yes	Yes	No	No
Observations	14,745	14,745	11,040	5,285
Pseudo R -squared	0.071	0.206	0.188	0.034
Industry-year strata	No	No	Yes	No
Firm strata	No	No	No	Yes
Firm cluster	Yes	Yes	No	No
Industry cluster	No	No	Yes	Yes

Table 2 (continued)*Panel B: Alternative specifications of supplier commitment*

	(1)	(2)	(3)	(4)
	C-LOGIT: Suppliers increase	C-LOGIT: Suppliers increase	C-LOGIT: Suppliers decrease	C-LOGIT: Suppliers decrease
CEO overconfidence $t-1$	0.904*** (0.198)			
Overconfidence trade up $t-1$		0.199** (0.099)		
CEO overconfidence up $t-1$			0.153 (0.129)	0.092 (0.127)
CEO overconfidence down $t-1$			0.403** (0.168)	0.301** (0.144)
CEO stock ownership $t-1$	-0.002 (0.012)	0.007 (0.010)	-0.026*** (0.008)	-0.014 (0.018)
Log (Total assets $t-1$)	-0.114* (0.067)	-0.154** (0.068)	0.702*** (0.108)	0.269*** (0.091)
Leverage $t-1$	0.198 (0.251)	0.130 (0.289)	-0.228 (0.211)	0.418 (0.395)
Market-to-book ratio $t-1$	-0.007 (0.010)	0.017 (0.011)	0.025** (0.011)	-0.001 (0.021)
Return on assets $t-1$	0.472 (0.461)	0.716* (0.420)	-0.727* (0.390)	-1.153** (0.482)
Sales growth $t-1$	0.359** (0.155)	0.435** (0.180)	-0.672*** (0.133)	-0.409* (0.229)
R&D intensity $t-1$	3.267*** (0.753)	2.661*** (0.777)	5.067*** (0.651)	0.658 (1.622)
Missing R&D $t-1$	0.184 (0.596)	0.344 (0.625)	0.148 (0.175)	-0.058 (0.367)
Log (1 + CEO Tenure $t-1$)	-0.063 (0.087)	0.003 (0.086)	0.149 (0.125)	0.081 (0.158)
Cash $t-1$	-0.429 (0.402)	-0.563 (0.374)	-0.019 (0.364)	-0.102 (0.635)
ROA volatility $t-1$	0.391 (0.591)	0.011 (0.501)	-0.359 (0.779)	-1.952*** (0.451)
Log (2-Year stock returns $t-1$)	-0.128 (0.097)	-0.062 (0.083)	-0.095 (0.091)	0.098 (0.102)
Investment $t-1$	1.171 (0.736)	1.247** (0.635)	1.556* (0.855)	0.259 (1.041)
Acquisitions $t-1$	-0.091 (0.082)	-0.074 (0.088)	0.053 (0.077)	-0.095 (0.061)
SEO proceeds $t-1$	-1.088 (0.779)	-0.092 (0.882)	-1.689 (1.102)	-1.447 (1.453)
Num. suppliers $t-1$	-0.131*** (0.031)	-0.127*** (0.030)	0.055 (0.068)	0.547*** (0.173)
Observations	5,575	5,629	11,684	5,583
Pseudo R -squared	0.035	0.029	0.213	0.130
Firm strata	Yes	Yes	No	Yes
Industry-year strata	No	No	Yes	No
Industry cluster	Yes	Yes	Yes	Yes

Table 3.

The effect of CEO overconfidence on the start of supplier network

This table presents results from logit and conditional logit regressions. *Start of supplier network* is a dummy variable that equals one when the firm has at least one dependent supplier in year t and none in $t-1$, and equals zero otherwise. Dependent suppliers are defined according to SFAS 14. *Confident CEO (options)* is a dummy variable which equals one if the CEO's vested option holdings are at least 67% in-the-money on at least two instances (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013), and equals zero otherwise. *Confident CEO (options)* switches from zero to one from the first such instance. *CEO overconfidence up* is a dummy variable that equals one if the continuous *CEO overconfidence* in year t is in the top quartile of the sample but not in the top quartile of the sample in year $t-1$, and equals zero otherwise. Columns 1 and 2 present results from a logit model while Columns 3 and 4 present results from a conditional logit model. Detailed variable definitions are in Appendix A. Table 1 contains a description of the sample. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 3 (continued)

	(1) LOGIT: Start Supp. Network	(2) LOGIT: Start Supp. Network	(3) C-LOGIT: Start Supp. Network	(4) C-LOGIT: Start Supp. Network
Confident CEO (options) $t-1$	0.242** (0.097)	0.179* (0.106)		
CEO overconfidence up $t-1$			0.434*** (0.129)	0.307** (0.149)
CEO Stock Ownership $t-1$		-0.018 (0.013)	-0.017 (0.013)	0.005 (0.015)
Log (Total Assets $t-1$)		0.317*** (0.036)	0.327*** (0.046)	-0.243*** (0.085)
Leverage $t-1$		0.072 (0.255)	0.138 (0.317)	0.153 (0.334)
Market-to-Book Ratio $t-1$		0.024 (0.018)	0.015 (0.021)	0.027* (0.016)
Return on Assets $t-1$		0.072 (0.526)	0.500 (0.425)	0.333 (0.630)
Sales Growth $t-1$		0.435** (0.173)	0.502** (0.202)	0.552** (0.253)
R&D intensity $t-1$		3.221*** (0.911)	3.167*** (0.587)	2.443*** (0.868)
Missing R&D $t-1$		-0.035 (0.169)	-0.057 (0.176)	0.560 (0.518)
Log (1 + CEO Tenure $t-1$)		0.007 (0.136)	0.099 (0.173)	0.143 (0.160)
Cash $t-1$		-0.134 (0.406)	-0.206 (0.370)	-0.380 (0.684)
ROA Volatility $t-1$		0.376 (0.856)	0.562 (0.568)	-1.061 (0.861)
Log (2-Year Stock Returns $t-1$)		-0.010 (0.091)	-0.030 (0.113)	-0.017 (0.084)
Investment $t-1$		-0.302 (0.809)	-0.285 (1.659)	-0.676 (1.283)
Acquisitions $t-1$		-0.012 (0.100)	-0.005 (0.097)	-0.126 (0.106)
SEO Proceeds $t-1$		-1.145 (1.454)	-1.110 (1.272)	-1.936* (1.037)
Industry dummies	Yes	Yes	No	No
Year dummies	Yes	Yes	No	No
Observations	14,731	14,731	9,054	4,040
Pseudo R -squared	0.027	0.056	0.045	0.018
Industry-year strata	No	No	Yes	No
Firm strata	No	No	No	Yes
Firm cluster	Yes	Yes	No	No
Industry cluster	No	No	Yes	Yes

Table 4

The effect of CEO overconfidence on employee turnover

This table presents results from one-sided panel Tobit regressions. Following Carter and Lynch (2004), *Employee turnover* is the ratio of the number of options that are cancelled, terminated, forfeited, expired, or lapsed in the year to the number of options outstanding at the beginning of the year. The construction of *Employee turnover* begins in 2004 when Compustat begins coverage of option cancellation data (*optca*). *Confident CEO (options)* is a dummy variable that equals one if the CEO's vested option holdings are at least 67% in-the-money on at least two instances (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013), and equals zero otherwise. *Confident CEO (options)* switches from zero to one from the first such instance. The sample in Column 2 comprises only firm-years with broad-based plans. An employee stock option plan is classified as broad-based if the number of options granted to non-executive employee exceeds 0.5% of the number of shares outstanding (e.g., Oyer and Schaefer, 2005; Hochberg and Lindsey, 2010; Chang et al., 2015). We split the sample in Columns 3 and 4 by the sample median of human capital intensity from the American Community Survey (ACS) census microdata. Detailed variable definitions are in Appendix A. Table 1 contains a description of the sample. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 4 (continued)

	(1) P-TOBIT: Employee turnover	(2) P-TOBIT: Employee turnover	(3) P-TOBIT: Employee turnover	(4) P-TOBIT: Employee turnover
Sample	All firms	Broad-based plans	High human-capital intensity	Low human-capital intensity
Confident CEO (options) $t-1$	-0.024*** (0.003)	-0.015*** (0.003)	-0.027*** (0.004)	-0.019*** (0.005)
Log (Firm age $t-1$)	0.003 (0.002)	-0.003 (0.002)	0.006* (0.004)	0.003 (0.003)
CEO stock ownership $t-1$	0.001** (0.000)	0.001 (0.000)	0.001** (0.001)	0.001 (0.000)
Log (Total assets $t-1$)	-0.013*** (0.001)	-0.010*** (0.001)	-0.009*** (0.002)	-0.017*** (0.002)
Leverage $t-1$	0.030*** (0.009)	0.013* (0.007)	0.020* (0.011)	0.042*** (0.013)
Market-to-book ratio $t-1$	-0.005*** (0.001)	-0.003*** (0.001)	-0.004*** (0.001)	-0.005*** (0.001)
Return on assets $t-1$	-0.173*** (0.021)	-0.169*** (0.024)	-0.219*** (0.023)	-0.124*** (0.037)
Sales growth $t-1$	-0.014 (0.009)	-0.015* (0.008)	-0.014 (0.012)	-0.011 (0.013)
R&D intensity $t-1$	0.049 (0.040)	-0.015 (0.036)	0.084 (0.054)	-0.009 (0.048)
Missing R&D $t-1$	0.001 (0.004)	-0.003 (0.005)	0.003 (0.005)	0.000 (0.007)
Log (1 + CEO tenure $t-1$)	-0.024*** (0.004)	-0.026*** (0.005)	-0.027*** (0.006)	-0.022*** (0.005)
Cash $t-1$	0.025* (0.014)	0.024** (0.011)	0.044*** (0.017)	0.000 (0.019)
ROA volatility $t-1$	-0.038*** (0.005)	-0.037*** (0.004)	-0.032*** (0.007)	-0.047*** (0.006)
Log (2-Year stock returns $t-1$)	0.200*** (0.029)	0.155*** (0.024)	0.163*** (0.041)	0.253*** (0.038)
Investment $t-1$	-0.116*** (0.034)	-0.084** (0.034)	-0.168*** (0.044)	-0.101** (0.048)
Acquisitions $t-1$	-0.010*** (0.003)	-0.002 (0.003)	-0.008* (0.004)	-0.013*** (0.005)
SEO proceeds $t-1$	-0.059 (0.050)	-0.058 (0.049)	-0.165*** (0.063)	0.064 (0.089)
Observations	7,676	4,294	3,995	3,681
Chi-squared	1140.45	837.36	1272.06	452.65
Industry-year fixed effects	Yes	Yes	Yes	Yes
One-tailed Welch-Satterthwaite t -test on Confident CEO (options) coefficients:				
Columns 3 and 4:		p -value = 0.087, d.f. = 7,565		

Table 5

The effect of CEO overconfidence on employee holdings of company stock

This table presents results from one-sided and two-sided panel Tobit regressions. *Employee stock holdings (%)* is the total dollar value of employer stock held in the benefit plan scaled by the total assets of the benefit plan in the year. *Employee stock holdings (\$)* is the dollar value of employer stock held in the benefit plan per active plan participant in the year. The dollar value of employer securities held in employee benefit plans are from item (1d1) of the firm’s Form 5500 Schedule H. In specifications of *Employee stock holdings (%)*, we perform two-sided panel Tobit regressions in Columns 1 and 2. In specifications of *Employee Stock Holdings (\$)* in Columns 3 and 4, we perform one-sided panel Tobit regressions. *Confident CEO (options)* is a dummy variable that equals one if the CEO’s vested option holdings are at least 67% in-the-money on at least two instances (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013), and equals zero otherwise. *Confident CEO (options)* switches from zero to one from the first such instance. *Annual employer cash match* is a dummy variable that equals one if the firm offers matching cash contributions to its employees in the plan, and equals zero otherwise. *Annual employer stock match* is a dummy variable that equals one if the firm offers matching non-cash contributions to its employees in the plan, and equals zero otherwise. The dollar values of matching non-cash contributions and cash contributions made by the firm are extracted from item (2a2) and item (2a1a) of the firm’s IRS Form 5500 Schedule H. *Per-employee non-stock wealth* is the total assets of the benefit plan less the total dollar value of employer securities, scaled by the number of active plan participants. Detailed variable definitions are in Appendix A. Table 1 contains a description of the sample. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 5 (continued)

	(1) Employee stock holdings (%)	(2) Employee stock holdings (%)	(3) Employee stock holdings (\$)	(4) Employee stock holdings (\$)
Require positive employee stock holdings in previous year	No	Yes	No	Yes
Confident CEO (options) $t-1$	0.027*** (0.007)	0.012** (0.006)	2.123*** (0.425)	1.570*** (0.386)
Annual employer cash match $t-1$	0.072*** (0.013)	0.023** (0.012)	3.610*** (0.822)	1.040 (0.690)
Annual employer stock match $t-1$	0.145*** (0.015)	0.048*** (0.011)	9.964*** (1.291)	3.755*** (0.899)
Per-employee non-stock wealth $t-1$	-0.002*** (0.000)	-0.002*** (0.000)	0.045*** (0.008)	0.062*** (0.007)
Dividends-to-price ratio $t-1$	0.192 (0.261)	0.615** (0.246)	46.644*** (15.459)	73.470*** (15.598)
Price-to-earnings ratio $t-1$	-0.000 (0.000)	-0.000 (0.000)	-0.007* (0.004)	-0.002 (0.003)
Working capital $t-1$	-0.187*** (0.027)	-0.111*** (0.025)	-12.486*** (1.593)	-7.077*** (1.564)
Log (Firm age $t-1$)	0.068*** (0.007)	0.028*** (0.005)	3.323*** (0.444)	1.138*** (0.325)
CEO stock ownership $t-1$	-0.002*** (0.001)	-0.002*** (0.000)	-0.192*** (0.046)	-0.227*** (0.042)
Log (Total assets $t-1$)	0.004 (0.003)	0.006** (0.002)	0.011 (0.182)	0.168 (0.149)
Leverage $t-1$	0.013 (0.018)	-0.015 (0.015)	1.454 (1.116)	0.939 (1.012)
Market-to-book ratio $t-1$	0.007*** (0.002)	0.008*** (0.002)	0.400*** (0.113)	0.466*** (0.099)
Return on assets $t-1$	0.089* (0.048)	0.119*** (0.042)	4.487 (3.498)	6.696** (3.027)
Sales growth $t-1$	-0.021 (0.017)	-0.004 (0.017)	-1.052 (1.226)	-0.767 (1.120)
R&D intensity $t-1$	-0.203*** (0.066)	0.050 (0.064)	-15.035*** (4.724)	4.526 (4.927)
Missing R&D $t-1$	-0.047*** (0.010)	-0.048*** (0.008)	-2.975*** (0.684)	-2.620*** (0.543)
Log (1 + CEO tenure $t-1$)	0.003 (0.010)	0.006 (0.008)	0.646 (0.557)	0.818* (0.489)
Cash $t-1$	0.065* (0.035)	0.115*** (0.030)	3.802* (2.031)	5.737*** (1.719)
ROA volatility $t-1$	-0.181*** (0.061)	-0.090 (0.060)	-10.680*** (4.041)	-5.523 (3.946)
Log (2-Year stock returns $t-1$)	0.016*** (0.006)	0.010** (0.005)	1.033*** (0.352)	0.608** (0.301)
Investment $t-1$	0.132 (0.084)	0.385*** (0.089)	-2.044 (4.623)	8.700** (4.355)

Table 5. (Continued)

	(1)	(2)	(3)	(4)
Acquisitions $t-1$	-0.002 (0.007)	-0.006 (0.006)	-0.132 (0.431)	-0.478 (0.369)
SEO proceeds $t-1$	-0.074 (0.101)	0.028 (0.090)	-0.752 (5.616)	3.944 (5.071)
Observations	7,887	3,577	7,887	3,577
Chi-squared	422.18	336.94	702.06	632.41
Industry-year fixed effects	Yes	Yes	Yes	Yes

Table 6

The effect of CEO overconfidence on suppliers in industries with high relationship-specific investment

This table presents results from conditional logit regressions. *Suppliers increase* is a dummy variable that equals one if a firm has more dependent suppliers in year t than it has in $t-1$, and equals zero otherwise. *Start of supplier network* is a dummy variable that equals one when the firm has at least one dependent supplier in year t and none in $t-1$, and equals zero otherwise. The key independent variable in Panel A and Panel C is *Confident CEO (options)*. *Confident CEO (options)* is a dummy variable that equals one if the CEO's vested option holdings are at least 67% in-the-money on at least two instances (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013), and equals zero otherwise. *Confident CEO (options)* switches from zero to one from the first such instance. The key independent variable in Panel B and Panel D is *CEO overconfidence up*. *CEO overconfidence up* indicator is a dummy variable that equals one if the continuous *CEO overconfidence* measure in year t is in the top quartile of the sample and not in the top quartile of the sample in year $t-1$, and equals zero otherwise. Panel A and Panel B include manufacturing firms split into subsamples of durable industries (SIC 3400 – 3999) and non-durable industries (SIC 2000 – 3399). Panel C and Panel D include manufacturing firms split into subsamples based on industry-level relationship-specificity according to the year 1997 contract intensity variable from Nunn (2007). An industry is classified as High (Low) RSI if its industry contract intensity is above (below) the sample median. Detailed variable definitions are in Appendix A. Table 1 contains a description of the sample. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 6 (continued)*Panel A: Durable/Non-durable manufacturing industries subsample analysis*

	(1)	(2)	(3)	(4)
	LOGIT: Suppliers increase	LOGIT: Suppliers increase	LOGIT: Start supp. network	LOGIT: Start supp. network
Industry type	Durable	Non-durable	Durable	Non-durable
Confident CEO (options) $_{t-1}$	0.225* (0.136)	-0.011 (0.171)	0.308* (0.187)	-0.055 (0.224)
Control variables	Yes	Yes	Yes	Yes
Observations	4,946	4,037	4,946	4,037
Pseudo R -squared	0.219	0.241	0.069	0.073
Industry dummies	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes
Firm cluster	Yes	Yes	Yes	Yes

Panel B: Durable/Non-durable manufacturing industries subsample analysis

	(1)	(2)	(3)	(4)
	C-LOGIT: Suppliers increase	C-LOGIT: Suppliers increase	C-LOGIT: Start supp. network	C-LOGIT: Start supp. network
Industry type	Durable	Non-durable	Durable	Non-durable
CEO overconfidence up $_{t-1, t}$	0.559*** (0.210)	0.249 (0.271)	0.564*** (0.149)	-0.207 (0.249)
Control variables	Yes	Yes	Yes	Yes
Observations	1,771	1,411	1,372	1,032
Pseudo R -squared	0.042	0.070	0.026	0.028
Firm strata	Yes	Yes	Yes	Yes
Industry cluster	Yes	Yes	Yes	Yes

Table 6 (continued)*Panel C: High/Low RSI manufacturing industries subsample analysis*

	(1)	(2)	(3)	(4)
	LOGIT: Suppliers increase	LOGIT: Suppliers increase	LOGIT: Start supp. network	LOGIT: Start supp. network
Relationship specificity	High	Low	High	Low
Confident CEO (options) $_{t-1}$	0.299* (0.164)	0.060 (0.176)	0.458** (0.219)	0.046 (0.233)
Control variables	Yes	Yes	Yes	Yes
Observations	4,472	4,057	4,472	4,057
Pseudo R -squared	0.210	0.251	0.083	0.070
Industry dummies	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes
Firm cluster	Yes	Yes	Yes	Yes

Panel D: High/Low RSI manufacturing industries subsample analysis

	(1)	(2)	(3)	(4)
	C-LOGIT: Suppliers increase	C-LOGIT: Suppliers increase	C-LOGIT: Start supp. network	C-LOGIT: Start supp. network
Relationship specificity	High	Low	High	Low
CEO overconfidence up $_{t-1, t}$	0.340** (0.140)	0.382 (0.358)	0.464*** (0.096)	0.040 (0.334)
Control variables	Yes	Yes	Yes	Yes
Observations	1,583	1,221	1,231	939
Pseudo R -squared	0.045	0.073	0.027	0.028
Firm strata	Yes	Yes	Yes	Yes
Industry cluster	Yes	Yes	Yes	Yes

Table 7

The effect of CEO overconfidence on relationship-specific investment

This table presents results from OLS regressions. The unit of observation is a customer-supplier pair in a year. *Supplier R&D intensity* is calculated by first multiplying the R&D expenses of the supplier by the supplier-to-customer pair sales to total supplier sales ratio, before scaling by the total assets of the supplier. We then multiply the raw value by a factor of 1,000 and take its logarithmic transformation. *Confident CEO (options)* is a dummy variable that equals one if the CEO's vested option holdings are at least 67% in-the-money on at least two instances (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013), and equals zero otherwise. *Confident CEO (options)* switches from zero to one from the first such instance. In Columns 3 and 4, we split the sample on the two-digit SIC industry median value of *Supplier-customer size ratio*. For each customer-supplier pair in a year, *Supplier-customer size ratio* is computed as the ratio of the total assets of the supplier to the total assets of the customer. Where multiple fixed effects are deployed, singleton observations are eliminated from the sample. Detailed variable definitions are in Appendix A. Table 1 contains a description of the sample. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 7 (continued)

	(1) OLS: Supp. R&D intensity	(2) OLS: Supp. R&D intensity	(3) OLS: Supp. R&D intensity	(4) OLS: Supp. R&D intensity
Sample	All supp-cust pairs	All supp-cust pairs	Small supp-cust size ratio	Big supp-cust size ratio
Confident CEO (options) $t-1$	0.116*** (0.042)	0.072* (0.038)	0.149*** (0.046)	-0.006 (0.048)
<u>Customer control variables</u>				
Stock ownership $t-1$	0.023*** (0.007)	0.020*** (0.007)	0.034*** (0.011)	0.019** (0.008)
Log (Market cap $t-1$)	0.010 (0.018)	0.054*** (0.017)	0.057** (0.025)	-0.013 (0.023)
Leverage $t-1$	-0.545*** (0.134)	-0.421*** (0.123)	-0.501*** (0.160)	-0.547*** (0.164)
Return on assets $t-1$	0.231 (0.352)	0.207 (0.330)	0.556 (0.423)	0.247 (0.408)
Sales growth $t-1$	0.267** (0.116)	0.259** (0.106)	0.132 (0.139)	0.308** (0.122)
R&D intensity $t-1$	4.991*** (0.591)	4.086*** (0.524)	3.552*** (0.759)	4.255*** (0.631)
Missing R&D $t-1$	-0.122 (0.097)	-0.098 (0.089)	-0.073 (0.117)	-0.121 (0.105)
Log (1 + CEO tenure $t-1$)	-0.071 (0.055)	-0.061 (0.048)	-0.039 (0.059)	-0.074 (0.060)
Cash $t-1$	0.862*** (0.250)	0.788*** (0.242)	0.542* (0.321)	0.910*** (0.292)
ROA volatility $t-1$	-0.238 (0.593)	-0.262 (0.554)	-1.136 (0.801)	0.515 (0.663)
Num. suppliers $t-1$	0.005*** (0.001)	0.006*** (0.001)	0.004*** (0.001)	0.007*** (0.001)
Log (2-Year stock returns $t-1$)	-0.047 (0.041)	-0.007 (0.039)	-0.052 (0.044)	0.031 (0.053)
Investment $t-1$	-1.420*** (0.518)	-1.238*** (0.479)	-1.470** (0.645)	-0.110 (0.585)
Acquisitions $t-1$	0.082** (0.041)	0.044 (0.037)	0.031 (0.047)	0.015 (0.048)
SEO proceeds $t-1$	-2.040 (2.512)	-1.957 (2.438)	1.612 (2.457)	-3.586 (2.550)
<u>Supplier control variables</u>				
Leverage $t-1$		-0.934*** (0.065)	-0.598*** (0.086)	-1.002*** (0.094)
M/B ratio $t-1$		0.026*** (0.003)	0.014*** (0.004)	0.030*** (0.005)
Return on assets $t-1$		-0.932*** (0.065)	-0.895*** (0.071)	-1.131*** (0.128)

Table 7 (continued)

	(1)	(2)	(3)	(4)
Sales growth $t-1$		0.067*** (0.024)	0.075** (0.030)	0.025 (0.041)
Investment $t-1$		-0.001*** (0.000)	-0.009*** (0.001)	-0.001*** (0.000)
Cash $t-1$		0.001*** (0.000)	0.002*** (0.001)	0.000*** (0.000)
ROA volatility $t-1$		0.773*** (0.080)	0.706*** (0.089)	0.768*** (0.155)
Log (Market cap $t-1$)		-0.072*** (0.009)	0.003 (0.017)	0.004 (0.018)
Observations	12,074	12,074	6,196	5,871
<i>R</i> -squared	0.277	0.364	0.338	0.402
Industry fixed effects	Yes	Yes	Yes	Yes
Customer-year cluster	Yes	Yes	Yes	Yes

Table 8

The effect of CEO overconfidence on duration of customer-supplier relationships

This table presents results from Andersen-Gill Cox regressions. The start of a customer-supplier relationship is either the first time the relationship is documented in the Compustat business segment file or the earliest year in which the previously terminated relationship is restarted. A customer-supplier pair with one previous termination and one restart is represented as two relationships (observations) in the data set. Following Fee, Hadlock, and Thomas (2006), if a relationship ends because the supplier/customer becomes inactive in Compustat or if the relationship lasts up till the final year of the sample period (2011), we treat the length of relationship as being right-censored. A relationship is classified as an *Overconfident CEO relationship* if a *Confident CEO (options)* customer CEO is incumbent between the start (inclusive) and end (exclusive) of the relationship. Since an *Overconfident CEO relationship* may take at least two years to be established, we only include relationships that begin on or after 1995. *Supplier R&D intensity*, *Sales dependency* and *Supplier-customer size ratio* are the values in the first year of the relationship, respectively. In Columns 2 to 5, we also include all supplier and customer control variables used in Table 7. These control variables are matched to the first year of the relationship. We include interactions of covariates with time in Columns 4 and 5. We allow customer-supplier relationships to end and restart in Columns 1 to 4. In Column 5, we restrict the sample to customer-supplier relationships with no restarts. Standard errors are reported in parentheses. Estimated coefficients of supplier and customer control variables are not presented in the table to facilitate presentation. These control variables are customer CEO stock ownership, customer CEO tenure, market capitalization, leverage, market-to-book ratio, ROA, sales growth, investment, R&D expenditure, cash, ROA volatility, number of suppliers, 2-year stock returns, and M&A activities. Estimated coefficients are presented as hazard ratios. Detailed variable definitions are in Appendix A. Table 1 contains a description of the sample. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 8 (continued)

	(1) Andersen-Gill Cox model	(2) Andersen-Gill Cox model	(3) Andersen-Gill Cox model	(4) Andersen-Gill Cox model	(5) Andersen-Gill Cox model
Supplier-customer relationship restarts allowed	Yes	Yes	Yes	Yes	No
<u>Main covariates</u>					
Overconfident CEO relationship	0.892*** (0.036)	0.825*** (0.042)	0.803*** (0.044)	0.708*** (0.089)	0.660*** (0.098)
Supplier-customer size ratio			1.067 (0.084)	1.317 (0.239)	1.529** (0.325)
Supplier R&D intensity			1.019 (0.022)	1.031 (0.049)	1.042 (0.060)
Sales dependency			0.781*** (0.030)	0.743*** (0.067)	0.675*** (0.073)
Supplier and customer controls	No	Yes	Yes	Yes	Yes
<u>Interactions with time</u>					
Overconfident CEO relationship				1.010 (0.013)	1.012 (0.015)
Supplier R&D intensity				0.999 -0.005	0.998 (0.006)
Sales dependency				1.006 (0.009)	1.010 (0.011)
Supplier-customer size ratio				0.982 (0.018)	0.967 (0.021)
Supplier and customer controls				Yes	Yes
Observations	4,974	3,579	3,060	3,060	2,208
Unique customer-supplier pairs	4,427	3,181	2,736	2,736	2,208
Customer-supplier pair-years	14,502	10,604	8,958	8,958	6,185
Log likelihood	-19,341	-13,445	-11,167	-11,142	-7,029

Table 9

Do overconfident CEOs overpay suppliers?

This table presents results from OLS regressions. *Input costs* is the ratio of cost of goods sold (COGS) to total sales. *Markup percentage* is the ratio of the difference between total sales and COGS to COGS. *Confident CEO (options)* is a dummy variable which equals one if the CEO's vested option holdings are at least 67% in-the-money on at least two instances (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013), and equals zero otherwise. *Confident CEO (options)* switches from zero to one from the first such instance. We restrict the sample to include only firms with at least one supplier in the current year and at least one supplier in the preceding year. Detailed variable definitions are in Appendix A. Table 1 contains a description of the sample. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 9 (continued)

	(1) OLS: Input costs	(2) OLS: Markup percentage
Confident CEO (options) $t-1$	-0.014* (0.008)	0.164*** (0.050)
CEO stock ownership $t-1$	0.001 (0.001)	-0.004 (0.004)
Log (Total assets $t-1$)	0.003 (0.003)	0.008 (0.023)
Leverage $t-1$	-0.002 (0.024)	-0.112 (0.152)
Market-to-book ratio $t-1$	-0.008*** (0.001)	0.043*** (0.011)
Return on assets $t-1$	-0.506*** (0.056)	2.796*** (0.506)
Sales growth $t-1$	0.094*** (0.020)	-0.147 (0.136)
R&D intensity $t-1$	-1.526*** (0.096)	9.668*** (1.187)
Missing R&D $t-1$	-0.006 (0.022)	0.010 (0.145)
Log (1 + CEO tenure $t-1$)	-0.008 (0.009)	0.014 (0.067)
Cash $t-1$	0.149*** (0.053)	-0.997** (0.459)
ROA volatility $t-1$	-0.169** (0.075)	1.765*** (0.642)
Num. suppliers $t-1$	0.001*** (0.000)	-0.005*** (0.002)
Log (2-Year stock returns $t-1$)	0.030*** (0.009)	-0.219*** (0.070)
Investment $t-1$	-0.045 (0.104)	-0.310 (0.760)
Acquisitions $t-1$	-0.013* (0.008)	0.096 (0.075)
SEO proceeds $t-1$	-0.151 (0.233)	0.600 (1.659)
Observations	2,244	2,244
R -squared	0.654	0.535
Industry-year fixed effects	Yes	Yes
Industry-year cluster	Yes	Yes

Table 10

The effect of CEO overconfidence on suppliers: High/low cash flow subsamples

This table presents results from logit regressions. *Suppliers increase* is a dummy variable that equals one if a firm has more dependent suppliers in year t than it has in $t-1$, and equals zero otherwise. Dependent suppliers are defined according to SFAS 14. *Confident CEO (options)* is a dummy variable which equals one if the CEO's vested option holdings are at least 67% in-the-money on at least two instances (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013), and equals zero otherwise. *Confident CEO (options)* switches from zero to one from the first such instance. We split firms by their lagged levels of *Cash flow*. Following Chang et al. (2014), we define cash flow as the sum of income before extraordinary items (*ibc*), extraordinary items and discontinued operations (*xidoc*), depreciation and amortization (*dpc*), deferred taxes (*txdc*), equity in net loss (*esubc*), gains in sale of PP&E and investment (*sppiv*), other funds from operation (*fopo*), the exchange rate effect (*exre*), and less changes in working capital (*wcapc*). We scale cash flow by total assets (*at*) to compute *Cash flow*. We define a firm to have high (low) *Cash flow* if its *Cash flow* is higher (lower) than the median value in the year. In Columns 3 and 4, we exclude observations in the Great Recession years—2008 and 2009. Detailed variable definitions are in Appendix A. Table 1 contains a description of the sample. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 10 (continued)

	(1) LOGIT: Suppliers increase	(2) LOGIT: Suppliers increase	(3) LOGIT: Suppliers increase	(4) LOGIT: Suppliers increase
Cash flow	Low	High	Low	High
Exclude great recession years	No	No	Yes	Yes
Confident CEO (options) $t-1$	0.246** (0.120)	0.176* (0.106)	0.253** (0.122)	0.201* (0.115)
Financial constraints $t-1$	0.009 (0.018)	-0.019 (0.017)	0.021 (0.019)	-0.020 (0.018)
CEO stock ownership $t-1$	0.006 (0.014)	-0.016 (0.014)	0.005 (0.014)	-0.013 (0.014)
Log (Total assets $t-1$)	0.854*** (0.055)	0.876*** (0.055)	0.849*** (0.057)	0.869*** (0.058)
Leverage $t-1$	-0.546* (0.282)	-0.646 (0.395)	-0.464 (0.293)	-0.742* (0.407)
Market-to-book ratio $t-1$	0.086*** (0.023)	0.024 (0.019)	0.090*** (0.025)	0.028 (0.020)
Return on assets $t-1$	0.672 (0.748)	-1.083 (0.820)	0.591 (0.773)	-1.029 (0.888)
Sales growth $t-1$	0.433** (0.184)	0.282 (0.261)	0.335* (0.182)	0.335 (0.275)
R&D intensity $t-1$	3.340** (1.323)	6.711*** (0.983)	3.055** (1.386)	6.716*** (1.072)
Missing R&D $t-1$	0.230 (0.214)	0.226 (0.199)	0.242 (0.215)	0.217 (0.222)
Log (1 + CEO tenure $t-1$)	0.029 (0.164)	-0.120 (0.137)	0.053 (0.171)	-0.148 (0.147)
Cash $t-1$	0.787 (0.511)	-0.127 (0.521)	0.899* (0.528)	-0.204 (0.545)
ROA volatility $t-1$	0.450 (1.200)	1.633* (0.881)	0.752 (1.239)	1.388 (0.915)
Log (2-Year stock returns $t-1$)	-0.132 (0.117)	0.036 (0.115)	-0.126 (0.121)	0.011 (0.117)
Investment $t-1$	1.464 (1.301)	1.267 (0.895)	1.873 (1.331)	1.240 (0.939)
Acquisitions $t-1$	-0.068 (0.123)	0.057 (0.100)	-0.063 (0.130)	0.069 (0.107)
SEO proceeds $t-1$	-0.330 (1.603)	0.060 (1.599)	-0.956 (1.608)	0.312 (1.613)
Num. suppliers $t-1$	-0.033 (0.021)	-0.081*** (0.020)	-0.042** (0.020)	-0.083*** (0.021)
Observations	7,262	7,236	6,299	6,196
Pseudo R -squared	0.224	0.208	0.215	0.206
Industry and year dummies	Yes	Yes	Yes	Yes
Firm cluster	Yes	Yes	Yes	Yes
Two-tailed Welch-Satterthwaite t -test on Confident CEO (options) coefficients:				
Columns 1 and 2:	p -value = 0.663, d.f. = 14,285			
Columns 3 and 4:	p -value = 0.757, d.f. = 12,461			

Table 11

Does the manner of leadership matter?

This table presents results from logit and two-sided panel Tobit regressions. The dependent variables are *Suppliers increase* in Columns 1 and 2, *Employee turnover* in Columns 3 and 4, and *Employee stock holdings (%)* in Columns 5 and 6. *Suppliers increase* is a dummy variable that equals one if a firm has more dependent suppliers in year t than it has in $t-1$, and equals zero otherwise. Dependent suppliers are defined according to SFAS 14. Following Carter and Lynch (2004), *Employee turnover* is the ratio of the number of options that are cancelled, terminated, forfeited, expired, or lapsed in the year to the number of options outstanding at the beginning of the year. The construction of *Employee turnover* begins in 2004 when Compustat begins coverage of option cancellation data (*optca*). *Employee stock holdings (%)* is the total dollar value of employer stock held in the benefit plan scaled by the total assets of the benefit plan in the year. We require benefit plans to have positive employee stock holdings in the previous year. *CEO media positivity* is a dummy variable that equals one if the number of articles containing references to confidence is more than the number of articles containing references to non-confidence during the year, and equals zero otherwise (Banerjee, Humphery-Jenner, and Nanda, 2015). Missing values of *CEO media positivity* are set to zero. *Confident CEO (options)* is a dummy variable which equals one if the CEO's vested option holdings are at least 67% in-the-money on at least two instances (e.g., Campbell et al., 2011; Ahmed and Duellman, 2013), and equals zero otherwise. *Confident CEO (options)* switches from zero to one from the first such instance. Detailed variable definitions are in Appendix A. Table 1 contains a description of the sample. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 11 (continued)

	(1)	(2)	(3)	(4)	(5)	(6)
	LOGIT: Suppliers increase	LOGIT: Suppliers increase	P-TOBIT: Employee turnover	P-TOBIT: Employee turnover	P-TOBIT: Employee stock holdings (%)	P-TOBIT: Employee stock holdings (%)
CEO media positivity $t-1$	0.034 (0.152)	0.028 (0.152)	-0.002 (0.005)	-0.001 (0.005)	0.008 (0.012)	0.005 (0.012)
Confident CEO (options) $t-1$		0.227** (0.109)		-0.017*** (0.003)		0.026*** (0.008)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7,923	7,923	5,231	5,231	3,121	3,121
Pseudo R -squared/Chi-squared	0.214	0.215	893.04	934.84	651.52	699.91
Industry and year dummies	Yes	Yes	No	No	No	No
Industry-year fixed effects	No	No	Yes	Yes	Yes	Yes
Firm cluster	Yes	Yes	No	No	No	No

Do security analysts learn from their colleagues?

Kenny Phua, T. Mandy Tham, and Chishen Wei*

Abstract

We examine how learning from colleagues affects security analyst forecast outcomes. We represent the brokerage house as an information network of analysts connected through industry overlaps in their coverage portfolios. Analysts who are more centrally connected in their brokerage network produce more accurate forecast estimates and generate more influential forecast revisions. Consistent with learning, more central analysts tend to unwind their colleagues' recent forecast errors in their own forecast revisions. Benefits from inter-colleague information exchange accrue to all analysts at the same brokerage.

JEL Classification Code: D83, G17, G24

Keywords: Learning, Networks, Analyst Forecast Accuracy

* Corresponding author: Chishen Wei (cswei@smu.edu.sg) +65 6828-9650. Phua is at Nanyang Business School, Nanyang Technological University, Singapore. Tham is at Wealth Management Institute, Nanyang Technological University. Wei is at the Lee Kong Chian School of Business, Singapore Management University. We are grateful for helpful comments from Gennaro Bernile, Ekkehart Boehmer, Steve Dimmock, Sinan Gokkaya, Umit Gurun, Jarrad Harford, Harrison Hong, Byoung-Hyoun Hwang, Chuan-yang Hwang, Bryan Lim, Roger Loh, Abhiroop Mukherjee, Chayawat Ornthanalai, Rui Shen, Avaniidhar Subrahmanyam, Gary Twite, Rong Wang, Bohui Zhang, Zhuo Zhong, and seminar and conference participants at the Asian FA Meeting (Bangkok), CICF (Xiamen), Complexity Institute at NTU, Deakin University, EFA (Oslo), FIRCG (Melbourne), FMA Meeting (Boston), Fudan University, Hong Kong Baptist University, Hong Kong Polytechnic University, University of Melbourne, Nanyang Business School, Singapore Management University, and SFS Cavalcade (Nashville). All errors are our own. The authors acknowledge research support from Nanyang Technological University.

A large body of research shows that securities analysts play an important role as information intermediaries in financial markets. For example, security analysts supply valuable information (e.g., Womack, 1996; Chen, Francis, and Jiang, 2005; Bradley, Clarke, Lee, and Ornathanalai, 2014; Loh and Stulz, 2011, 2017) and shape the information environment (e.g., Balakrishnan, Billings, Kelly, and Ljungqvist, 2014; Lou, Cohen, and Malloy, 2014; Merkley, Michaely, and Pacelli, 2017).¹ This naturally raises the question of how forecast estimates are generated (e.g., Bradshaw, 2011). Extant studies find that forecast accuracy is a function of individual attributes including experience, personal characteristics, personal ties, and professional connections. However, even the best security analysts do not work in isolation. Their brokerage houses provide not only operational and back-office resources, but also a group of colleagues. These colleagues provide a potentially valuable network of knowledge and information.

For example, an analyst who covers Apple may receive useful industry insights from a colleague who covers Intel.² This type of information sharing is valuable because studies show that industry knowledge helps analysts to produce better forecast estimates.³ Indeed, our conversations with analysts suggest that colleagues often workshop ideas and solicit feedback from their colleagues. Lehman Brother’s research department in the early 1990s is an example of the importance of knowledge sharing. To foster communication and learning, Lehman Brothers instituted a policy that every analyst’s presentation must reference the work of at least two other colleagues. During that period, Lehman Brothers was regularly ranked among the top brokerage firms.⁴

We hypothesize that analysts produce higher-quality equity research if they are better able to tap into the knowledge base of their in-house colleagues. To test our hypothesis, we map the information network within a brokerage house using industry overlaps among analyst coverage portfolios. Brown et al. (2015) write, “*Industry knowledge is the single*

¹ Securities analyst also influence financial policy and valuation (e.g., Lang, Lins, and Miller, 2004; Chang, Dasgupta, and Hilary, 2006; Derrien and Kecskes, 2013).

² Throughout the paper, we use colleagues to refer to analysts that work at the same brokerage house.

³ See for example Boni and Womack (2006), Clement, Koonce, and Lopez (2007), Kadan, Madureira, Wang, and Zach (2012), Hilary and Shen (2013), Brown et al. (2015), Bradley, Gokkaya, and Liu (2017).

⁴ “The Risky Business of Hiring Stars,” Harvard Business Review, May 2004

most useful input to analysts' earnings forecasts and stock recommendations." Our premise is that information exchange is more likely to occur between a pair of analysts if there is industry sector overlap in their coverage portfolios.

Figure 1 provides an example of an actual brokerage network in our sample. Each node represents an analyst in the brokerage and a line between two nodes depicts a linkage between two analysts. To measure the prominence of each analyst in the brokerage network, we create an *Analyst Centrality* score that encapsulates four commonly used network measures, which capture both the importance of direct connections (colleagues covering the same industry sector) and the value of indirect information exchange beyond sector boundaries. For example, Apple's recent secretive foray into the electric car market suggests that information exchange with colleagues who cover the automobile sector may also help in the assessment of Apple's valuation. Therefore, an analyst located in a central nodal position has a high *Analyst Centrality* score and is at the epicenter of information exchange within the brokerage network. If analysts can learn from their colleagues, we hypothesize that central analysts are better positioned to benefit from such information exchange.

Our first set of tests indicate that analysts with higher *Analyst Centrality* scores produce better equity research. Central analysts exhibit higher forecast accuracy and their forecasts are more influential. Economically, an analyst who is at the 75th percentile of the *Analyst Centrality* distribution is 10.5% more accurate than one in the corresponding 25th percentile (relative to the median forecast error).⁵ Controlling for the confounding effects of 8-K issuances and earnings announcements, an interquartile increase in *Analyst Centrality* is associated with 0.27% higher market reaction within a three-day window around analyst forecasts. Our test specifications include firm-year fixed effects to ensure that our findings are not driven by underlying heterogeneity in the coverage firms. We also conduct tests with brokerage-year fixed effects; this specification is important because it rules out common alternative explanations relating to brokerage prestige, sector specialization, resources, and other important yet unobservable brokerage level

⁵ To put this into perspective, a similar inter-quartile increase in General Experience corresponds to a 14.8% improvement in forecast accuracy.

characteristics. Overall, our evidence supports the view that analysts with higher *Analyst Centrality* have better access to the information transmitted through their brokerage networks.

We expect that by leveraging the wide-ranging expertise of their colleagues, central analysts are better able to analyze hard-to-value firm. Consistent with this prediction, we find that central analysts more accurately forecast the earnings of firms with high R&D intensity and complex companies with multi-industry operating segments (e.g., Cohen and Lou, 2012). This suggests that one particular advantage of centrality in brokerage networks is the greater ease of piecing together value-relevant information across sectors.

To gain insights on how learning among colleagues might occur, we examine whether an analyst unwinds the *ex-post* forecast errors of her colleagues. Such information is public news, but likely to be more salient among colleagues. Therefore, an analyst may revise her own forecasts to incorporate the newly revealed information or to unravel the information previously learnt from her colleagues. We find evidence that analysts with higher *Analyst Centrality* are more likely to issue revisions that adjust for their colleagues' *ex-post* forecast errors. This evidence shows a potential channel of information propagation within brokerage networks and suggests that central analysts incorporate more colleagues' information into their own forecasts.

We hypothesize that the brokerage environment may play an important role in information sharing among colleagues. We find stronger effects of *Analyst Centrality* on forecast accuracy among larger brokerages—where the quality and volume of information exchange are likely higher—and denser networks—where the opportunities for information sharing are likely greater. Interestingly, the effect of *Analyst Centrality* on forecast accuracy is weaker at brokerage houses with either very low or very high analyst employment turnover. We propose two plausible explanations for this hump-shaped pattern. First, if analyst turnover proxies for competition at the brokerage, there may be an optimal level of competition that best promotes information exchange. On one hand, analysts may have weak tournament-like incentives to learn and outperform one another when competition is low. On the other hand, analysts may be disincentivized to share

value-relevant information and learning is thus inhibited. Second, the hump-shaped pattern perhaps reflects that an optimal level of analyst turnover is needed for new ideas to be generated and shared. Low turnover may lead to stale ideas and no new knowledge to be learned. However, new ideas may become lost quickly with rapid exits of colleagues.

We also find evidence that benefits of information exchange generally accrue to *all* analysts in brokerages—analysts in denser brokerage networks exhibit higher forecast accuracy, in aggregate. Therefore, while central analysts disproportionately benefit from colleague learning, the average analyst also benefits from brokerage networks that structurally facilitate information exchange.

An unaddressed issue is the matching process between the analyst and the brokerage house. Concerns relating to brokerage characteristics are partially addressed in our main tests with the inclusion of brokerage-year fixed effects in our baseline specification. However, an analyst may be more likely to match with a brokerage that shares her specialization in a particular sector(s).⁶ To explore the causal relation between *Analyst Centrality* and forecast accuracy, we exploit shocks to *Analyst Centrality* using a sample of brokerage mergers (e.g., Hong and Kacperczyk, 2010; Kelly and Ljungqvist, 2012) from years 2000 to 2007. These mergers introduce exogenous changes to the acquirer’s brokerage structure, shocking the *Analyst Centrality* scores of incumbent analysts. We find that analysts who experienced increases in their *Analyst Centrality* scores subsequently exhibited higher forecast accuracy relative to analysts whose *Analyst Centrality* scores declined. Results from a generalized difference-in-difference model (Autor, 2003) show that 1) increases/declines in *Analyst Centrality* are effectively random with respect to pre-treatment forecast accuracy, and 2) gaps in forecast accuracy between both groups exist only in the post-treatment periods.

One may wonder to what extent the *Analyst Centrality* measure reflects an analyst’s ability or effort. Of course, the development of higher centrality may be a consequence of greater analyst talent or skill. What distinguishes our learning hypothesis is that *Analyst*

⁶ All else being equal, such an analyst would have more colleagues who share sector overlaps with her, and thus be more connected in the general sense. However, as Section 1.2 and Appendix I make clear, *Analyst Centrality* is not solely driven by a count of an analyst’s direct linkages.

Centrality has explanatory power beyond measures of ability or effort found in the prior literature. Notably, our main specifications include proxies for ability (i.e., experience) and effort (i.e., revision frequency). Separately, we control for analyst ability by using membership in the Institutional Investor All-America Research Team (i.e. All-American status) as a proxy for analyst skill. The relations between *Analyst Centrality* and forecasting outcomes remain unchanged with the inclusion of this control in our main tests.

Our paper contributes to a growing literature that seeks to penetrate the ‘black box’ of information generation of sell-side financial analysts (e.g., Bradshaw, 2011; Brown et al., 2015). Extant studies find that forecast accuracy is a function of experience, location, personal characteristics, personal ties, and professional connections.⁷ Our paper emphasizes the value of information transfer between colleagues. Related to our findings is the evidence in Hwang, Liberti, and Sturgess (2016) of information transmission among in-house analysts in the mergers and acquisition setting. Equity analysts may also acquire information from their colleagues in debt/macro research (Hugon, Lin, and Markov, 2016; Hugon, Kumar, and Lin, 2016), on the asset management side (Irvine, Simko, and Nathan, 2004), and on the brokerage trading desk (Li, Mukherjee, and Sen, 2017). In contrast, we measure the general propensity of an analyst to receive information from her brokerage colleagues using a network-theoretic approach. This approach accounts for dependencies (including higher-order ones) between all analysts in a given brokerage, and captures potential information flow from one analyst to another across the entire network.

Our study also links to a literature that explores how analysts incorporate various signals into their forecasts. Chen and Jiang (2006) show how analysts weight public and private information in their forecasts. Clement, Hales, and Xu (2011) find that analysts revise forecast following the revisions of their direct competitors who cover the same company. Our evidence suggests that analysts learn more from their brokerage colleagues’

⁷ See for example Clement (1999), Jacob, Lys, and Neale (1999), Malloy (2005), Bae, Stulz, and Tan (2008), Bradley, Gokkaya, and Liu (2017), Chen and Jiang (2006), Chen and Matsumoto (2006), Clement, Koonce, and Lopez (2007), Cohen, Frazzini, and Malloy (2010), Kumar (2010), Clement, Hales, and Xue (2011), Hilary and Shen (2013), Law (2013), Green, Jame, Markov, and Subasi (2014a, 2014b), Malloy (2005), Soltes (2014), Jiang, Kumar, and Law (2016). For a recent review, see Bradshaw (2011).

mistakes than from mistakes of non-colleagues. This supports the notion that information exchange with brokerage colleagues is an important channel of information acquisition for analysts.

Our findings also contribute to the growing literature on information transmission through networks in financial markets. Venture capitalists and stock market investors who hold central positions in the network generate better investment returns (Hochberg, Ljungqvist, and Lu, 2007; Ozsoylev, Walden, Yavuz, and Bildik, 2013). Anjos and Fracassi (2015) find that internal information markets help conglomerates deliver more innovative and greater value. Studies also identify information transmission through workplace connections, personal networks, and social connections using educational backgrounds or social affiliations.⁸

1. Sample and network definitions

This section describes our methodology and variable construction. We discuss our data and present summary statistics on our sample.

1.1. *Defining Network Connections*

Studies show that industry-specific information is valuable in making accurate earnings forecast.⁹ Therefore, we presume that information exchange and dissemination is more likely to occur between analysts that share industry sector overlaps in their coverage portfolios. This may occur for direct economic reasons. First, given the significant benefits of industry knowledge, there is economic value to discuss and solicit feedback from

⁸ For example, social ties with management affect both analyst forecast behavior and coverage decisions (e.g., Westphal and Clement, 2008; Cohen, Frazzini, and Malloy, 2010; Brochet, Miller, and Srivinasan, 2014). Information also transmits through social networks among investors (e.g., Hong, Kubik, and Stein, 2004, 2005; Pool, Stoffman, and Yonker, 2015), between boards and managers (e.g., Kuhnen, 2009), between corporate and bank managers (Engelberg, Gao, and Parsons, 2012), across corporate boards (e.g., Shue, 2013; Larcker, So, and Wang, 2013; Fracassi, 2016), and between corporate boards and mutual fund managers (Cohen, Frazzini, and Malloy, 2008).

⁹ See: Boni and Womack (2006), Clement, Koonce, and Lopez (2007), Kadan et al. (2012), Hilary and Shen (2013), Brown et al. (2015), Bradley, Gokkaya, and Liu (2017). Brown, Call, Clement, and Sharp (2015) write, “industry knowledge is the single most useful input to analysts’ earnings forecasts and stock recommendations.”

colleagues who cover the same sectors. Second, given economies of scale in backend resources such as sharing support staff, research assistants, and data sources, we expect interactions to arise naturally among analysts with sector overlaps in their coverage portfolios. Our study focuses primarily on these channels of information exchange, but we discuss the possibility of other related channels in the next section.

We construct the connections within brokerage networks based on the FY1 forecast data from the Detailed History file of I/B/E/S. We use the Global Industrial Classification Standard (GICS) codes to classify industries because these are commonly used among practitioners (e.g., Bhojraj, Lee, and Oler, 2003). To ensure that sample brokerages are meaningfully large for network analysis, we require that brokerages have at least five analysts in a given year. The relaxation of this sample filter does not have a substantial impact on any findings.

We allow brokerage network structures to vary through time to accommodate changes in analysts' coverage portfolios. For example, consider two analysts, A and B, at the same brokerage. A connection exists in year t if A makes forecasts in GICS sectors 20 and 45, and B makes forecasts in GICS sector 45. By construction, the connection weight is independent of the number of firms covered in the overlapping GICS sectors by either analyst. While it is reasonable to expect that the intensity of information exchange increases with the number of firms covered in the overlapping sectors of connected analysts, our simple counting approach is more parsimonious and conservative.

[Insert Figure 1]

As an example, we illustrate the network structure of a brokerage (I/B/E/S identifier: 481) in the year 2005. Each circular node represents an analyst in the brokerage network. The numbers below each node identify the GICS sectors covered by the analyst. Most analysts in this brokerage network cover up to two sectors, and only two analysts span four sectors in their coverage portfolios. A line between a pair of nodes denotes a direct connection between a pair of analysts. Figure 1 shows that analysts can have different levels of connectedness even if they cover the same number of economic sectors. For example, an analyst who covers GICS sectors 25 and 45 has more direct connections than

her colleague who covers GICS sectors 45 and 50.¹⁰ This suggests that the relation between an analyst’s coverage portfolio complexity and potential for information exchange is non-linear and dependent on the compositions of her colleagues’ coverage portfolios.

1.2. *Defining Analyst Centrality*

There are many dimensions of the information exchange process. Aside from different types of information (e.g., firm-specific, industry news, and macroeconomic news), there is variation in the amount (i.e., volume) and speed of information transmission. Our approach uses the four most common measures of network centrality to describe information exchange (e.g., Larker, So, and Wang, 2013). Each measure of network centrality captures a distinct feature of connectivity and information exchange. As such, any single measure is unlikely to capture all facets of information exchange (Newman, 2003). Characterizing the information exchange process can be challenging (Borgatti, 2005), but we describe below the intuition behind the types of inter-colleague information exchange that each measure is likely to capture and reserve the details for Appendix I.

- *Degree Centrality* (Newman, 1979) is conventionally used to capture information flow in networks (Borgatti, 2005). *Degree Centrality* of an analyst counts the number of her directly connected colleagues. An analyst with more direct connections has more opportunities for interaction and is thus more likely to receive information. Since analysts are connected by sector overlaps in our networks, *Degree Centrality* likely captures information that is industry-specific or related to competition dynamics.
- *Eigenvector Centrality* (Bonacich, 1972) is a close cousin to *Degree Centrality* and is often used to capture information flow in networks. Whereas *Degree Centrality* assigns equal weight to all direct connections, *Eigenvector Centrality* places more weight on directly connected colleagues who are themselves more central. This measure not only captures more valuable industry-specific information but also information that is more macroeconomic in nature.

¹⁰ This difference arises because there are more analysts in the brokerage network who cover GICS sector 25 than the ones who cover GICS sector 50.

- *Closeness Centrality* (Freeman, 1979) measures the aggregate network distance between an analyst and all her colleagues, and thus captures the speed (time-to-arrival) of receiving information that is produced anywhere in the network (Borgatti, 2005). The value of information in financial markets is time-sensitive, and speedy arrival of colleagues' information at an analyst helps her update her forecasts in a timely fashion.
- *Betweenness Centrality* (Freeman, 1979) of an analyst captures her position between industry teams in the network. Therefore, it relates to cross-industry information exchange¹¹ and the complementarities in information production from colleagues who cover different industries. An analyst who is positioned between industry teams can better bridge otherwise isolated information silos to reap cross-industry synergies.¹²

A large literature has studied information transmission through informal or social connections. In contrast, we focus on professional linkages based on sector overlaps in analysts' coverage portfolios. Nevertheless, it is possible that analysts in our sample are connected socially too, beyond these professional relations (e.g., Hwang and Kim, 2011). While data limitations prevent us from capturing these non-professional ties, we argue that information exchange outside of our constructed networks is likely to bias against us finding empirical support for our hypothesis.

1.3. *Principal Component Analysis*

Given the complexity of information flow in a brokerage network, any single measure of network centrality cannot completely capture analyst connectedness (Newman, 2003; Valente, Coronges, Lakon, and Costenbader, 2008). Each measure of centrality may be better at capturing different types of information exchange, as we discussed above.

¹¹ Kini et al. (2009) find that cross-industry information is valuable to forecast accuracy.

¹² Borgatti (2005) makes clear that, due to the algorithmic implications of betweenness centrality, it is incompatible with the canonical type of untargeted information exchange (i.e., those captured by degree and eigenvector centralities). We also do not adopt the typical interpretation that betweenness centrality represents information brokering capacity. Instead, we stress that analysts with synergistic coverage portfolios will mechanically be situated between many colleagues. Where there is an opportunity for an analyst to exploit synergies between two industries, the information exchange should be targeted and not via a circuitous path

The four centrality measures have cross-correlations ranging from 30.2% to 89.1%. We show graphical evidence of such high cross-correlations in Figure 2. It contains four snapshots of the brokerage network in Figure 1; each snapshot corresponds to one of our four measures of centrality. In each snapshot, we weight the size of the nodes (i.e. representing the analysts) and their color intensities by the corresponding measure of centrality; larger and more deeply colored nodes are more central. A visual inspection shows that analysts who are central under one measure are also likely to be central under another measure. To sidestep multicollinearity concerns, we perform a principal component analysis (PCA) on degree, betweenness, closeness and eigenvector centrality. From the first principal component, we extract its standardized factor score and define it as *Analyst Centrality*. We report more details of the PCA in Section 1.6.

[Insert Figure 2]

1.4. Forecast Accuracy

Following Clement (1999), we construct *Normalized Forecast Error* from the latest firm-year forecast values by calculating the absolute difference between the analyst's firm-year forecast value and the corresponding earnings-per-share (EPS) of the firm-year, scaled by the firm-year mean absolute forecast error for comparability across observations. For analyst i covering firm f in year t , we define her *Normalized Forecast Error* as in (1).

$$\text{Normalized Forecast Error}_{i,f,t} = \frac{|EPS \text{ Forecast}_{i,f,t} - \text{Actual } EPS_{f,t}|}{\frac{1}{N_{f,t}} \sum_j^{N_{f,t}} |EPS \text{ Forecast}_{j,f,t} - \text{Actual } EPS_{f,t}|} \quad (1)$$

1.5. Control Variables

We include controls for the determinants of forecasting performance found in prior research.¹³ We provide a brief description of the key variables and provide full construction details in Appendix II. All continuous variables are winsorized at the 1st and 99th percentile values to limit the influence of outliers.

¹³ See for example, Mikhail, Walther, and Willis (1997); Clement (1999); Jacob, Lys, and Neale (1999); and Brown (2001), Ivković and Jegadeesh (2004).

To control for forecast characteristics associated with accuracy, we include proxies for revision activity (*Revision Frequency*), timeliness based on the closeness to the actual earnings announcement (*Horizon*), and deviation from the consensus forecast (*Boldness*). As lowballing behavior is associated with forecast accuracy, we also construct a lowball measure (*Lowball*) following Hilary and Hsu (2013).¹⁴ Since centrality may be associated with ability, we proxy for analyst ability using total work experience (*General experience*) and experience covering the firm (*Firm Experience*). To account for the analyst’s portfolio complexity, we measure number of unique firms (*Firm Breadth*) and the number of 2-digit GICS sectors (*Industry Breadth*) covered by the analyst during the year. We measure *Brokerage Size* as the logarithm of the number of analysts employed by the brokerage.

We also account for the influence of firm heterogeneity on forecasting outcomes by controlling for the number of analysts (*Analyst Coverage*), market capitalization, book-to-market ratio, leverage, ROA volatility, and negative earnings (*Loss*)¹⁵.

1.6. Descriptive Statistics

Our sample comprises 5,539 firms, 317 brokerages, 9,170 analysts, and 274,671 firm-year forecasts from the years 1996 to 2014 using the I/B/E/S vintage from WRDS in May 2015. Panel A presents summary statistics of unique analyst-year pairs. The median analyst has 16 (12) connections, covers 1.56 (1) GICS sectors, and covers 9.66 (9) firms. The median brokerage employs 42 analysts. By definition, the four centrality measures are non-negative. However, many analysts in the sample have negative values of *Analyst Centrality* because we used standardized variables in the PCA-extraction.¹⁶

[Insert Table 1]

Panel B of Table 1 presents Pearson correlations between the network centrality measures and analyst/brokerage characteristics. Naturally, *Analyst Centrality* is positively correlated with its component measures (54.6% to 94.9%). *Analyst Centrality* is also

¹⁴ Hilary and Hsu (2013) show that analysts strategically increase forecast error consistency via lowballing behavior for firms with high institutional ownership. The lowballing strategy inevitably decreases forecast accuracy. Appendix II contains full details of the construction of lowball.

¹⁵ See Hwang, Jan, and Basu (1996) and Brown (1997; 1998).

¹⁶ Standardization accounts for the different scales of the four network centrality variables.

positively correlated with *Industry Breadth* (57.1%), but weakly related to *Firm Breadth* (16.5%), as analysts generally tend to cover firms in the same industry. *Analyst Centrality* has low correlations with measures of experience including *General Experience*, *Firm Experience*, and *Brokerage Experience*. Interestingly, *Analyst Centrality* is negatively correlated with *All-American* status, which suggests that our findings are unlikely driven by All-American analysts.

Panel C reports a principal component analysis of the four network centrality measures. The first principal component captures approximately 66.5% of the variance in the four network centrality measures and is the only eigenvalue greater than one. The incremental variance explained by the next principal component is only about 20.9%. Other loadings on some centrality measures turn negative in principal components, thus making interpretation difficult. In light of these considerations, we construct *Analyst Centrality* based on the first principal component.

2. Do central analysts produce better forecasts?

Our learning hypothesis proposes that analysts generate better equity research if they can learn valuable information from their colleagues. The tests in this section focus on the relation between *Analyst Centrality* and forecast outcomes. Our main prediction is that central analysts have higher forecast accuracy because they are better positioned to benefit from inter-colleague information exchange. We test this along two dimensions: 1) forecast accuracy, and 2) market reactions around issuances of forecast revisions.

2.1. Forecast accuracy

We test whether central analyst produce more accurate forecasts by estimating specification (2).

$$\begin{aligned} \text{Normalized Forecast Error}_{i,f,t} = & \alpha + \beta_1 \cdot \text{Analyst Centrality}_{i,t} \\ & + \theta \cdot \text{Controls}_{i,f,t} + \varepsilon_{i,f,t} \end{aligned} \quad (2)$$

Our forecast accuracy measure is *Normalized Forecast Error*, which allows for comparability across firms. The control variables include forecast horizon, revision

frequency, forecast boldness, analyst experience, coverage experience, lowballing behavior, and complexity of coverage portfolio (i.e., firm breadth of coverage and industry breadth of coverage). Additionally, we control for the following firm characteristics: earnings volatility, loss, firm leverage, growth opportunities, firm size, and analyst coverage (e.g., Heflin, Subrahmanyam, and Zhang, 2003; Hilary and Hsu, 2013).

Our baseline specification includes brokerage-year fixed effects to capture unobserved brokerage characteristics such as prestige and resources (e.g., Stickel, 1995; Clement, 1999; Hugon, Kumar, and Lin, 2016). We also estimate a specification with firm-year fixed effects to absorb unobserved heterogeneity across firms that affects analysts' forecasting performance in the firm-year. Standard errors are clustered along two dimensions following Petersen (2009) and Gow, Ormazabal, and Taylor (2012). First, we cluster by analyst-firm because an analyst's forecast error on a particular firm may be correlated over time. The second clustering dimension is either 1) firm-year to capture dependence in forecast errors of competing analysts in a firm-year or 2) brokerage-year to capture correlations of analysts' forecasting performance within each brokerage-year.

[Insert Table 2]

The results in Table 2 show a positive relation between *Analyst Centrality* and forecast accuracy. The brokerage-year specification in Column (1) shows that central analysts are more accurate than their peripheral colleagues in the same brokerage. The results are similar with the inclusion of controls for firm characteristics in Column (2). Economically, an analyst in the 75th percentile of *Analyst Centrality* is 10.5% more accurate relative to the median forecast error than an analyst in the 25th percentile. For comparison, a 25th to 75th interquartile increase in *General Experience* corresponds to a 14.8% improvement in forecast accuracy.¹⁷ Column (3) shows that the results are similar with the inclusion of firm-year fixed effects, which necessarily omit controls for firm characteristics. Our results

¹⁷ From Table 1, the difference in *Analyst Centrality* between the 75th and 25th percentiles of its distribution is 1.337. This difference corresponds to a decrease in *Normalized Forecast Error* of $1.337 \times 0.014 = 0.019$. We de-normalize this value by the mean raw forecast error to obtain $0.019 \times \$0.222 = \0.0042 , or about 10.5% of the median raw forecast error (\$0.04). An interquartile increase in *General Experience* (from 23 months to 90 months) is a 290% increment. This has an impact of $-0.020 \times \log(1+2.9) = 0.027$ decrease in *Normalized Forecast Error*. On average, this translates to $0.027 \times \$0.222 = \0.0059 , or about 14.8% of the median raw forecast error (\$0.04).

suggest that central analysts are more accurate than peripheral competitors from other brokerages who cover the same firm in the year.

Consistent with prior literature, forecast accuracy is associated with shorter forecasting horizons, more firm-specific experience, bolder forecasts, and less low balling behavior (e.g., Clement and Tse, 2005; Hilary and Hsu, 2013). Greater effort, as measured by higher revision frequency, is also associated with greater accuracy. Higher analyst coverage also creates greater accuracy, consistent with greater competition spurring more effort. (e.g., Merkley, Michaely, and Pacelli, 2017; Loh and Stulz, 2017).

It is noteworthy that analysts who cover more sectors (*Industry Breadth*) tend to have greater forecast errors as their coverage portfolio becomes more complex, consistent with prior literature (e.g., Clement, 1999; Clement and Tse, 2005). This is interesting because *Analyst Centrality* measure is related to *Industry Breadth*, but the two measures have opposing effects on forecast accuracy. An interquartile increase in *Industry Breadth* reduces accuracy by 3.3% relative to the median forecast error.¹⁸ This implies a tradeoff between the benefits of inter-colleague information exchange and the costs associated with a more complex coverage portfolio.

2.2. Forecast revision informativeness: Market reactions

If central analysts benefit from information exchange within their brokerages, we expect their forecast revisions to contain more novel and value-relevant information, which attracts greater market reactions. We test this prediction by estimating specification (3).

$$Abs.CAR_{i,f,t} = \alpha + \beta_1 \cdot Analyst\ Centrality_{i,t} + \theta \cdot Controls_{i,f,t} + \varepsilon_{i,f,t} \quad (3)$$

The unit of analysis is a forecast revision on a firm by a particular analyst. Standard errors are clustered by week to capture common time-varying macroeconomic shocks and also by firm because market reactions to forecast revisions may be correlated over time

¹⁸ From Table 1, the difference in *Industry Breadth* between the 75th and 25th percentiles of its distribution is 1. This difference corresponds to an increase in *Normalized Forecast Error* of $1 \times 0.006 = 0.006$. We de-normalize this value by the mean raw forecast error to obtain $0.006 \times \$0.222 = \0.0013 , or about 3.3% of the median raw forecast error (\$0.04).

for a given firm. We also control for forecast boldness (Clement and Tse, 2005) and stock performance in the run-up to the forecast revision date.

[Insert Table 3]

Table 3 shows that central analysts command larger market reactions around their forecast revisions. Column (1) shows a statistically positive association between *Analyst Centrality* and 3-day market reactions around forecast revisions. A standard deviation increase in *Analyst Centrality* is associated with a 0.33% larger market reaction around forecast revisions. The results are similar in Column (2) with the inclusion of controls for analyst characteristics and stock performance during run-ups to the forecast revisions. Consistent with Bradley, Gokkaya and Liu (2017), the forecast revisions of analysts who have longer general experience and who work in larger brokerages attract stronger market reactions. Overall, our evidence is consistent with the view that central analysts possess more novel and value-relevant information.

A potential concern is that confounding events including material firm announcements may coincide with analysts' forecast revisions. To alleviate this concern, we collect all dates on which firms file SEC Form 8-Ks¹⁹ from the EDGAR database. Thereafter, we drop a forecast revision from our sample if a Form 8-K is filed in the (-1, +1) day window of the revision. Additionally, we exclude a forecast revision from our sample if the firm has earnings announcements in the (-1, +1) day window of the revision. These filters reduce our sample size by about 48%, suggesting that analysts' revisions are often motivated by the disclosures of new information by firms. We repeat our analysis in this reduced sample and continue to find that the forecast revisions of central analysts attract larger market reactions in Column (3). An analyst at the 75th percentile of the *Analyst Centrality* distribution commands 0.27% larger market reaction around her forecasts than an analyst at the 25th percentile of this distribution.

¹⁹ In addition to annual and quarterly reports, public companies are required to report certain material corporate events on a more current basis via the Form 8-K. A comprehensive list of the event types that trigger a firm's obligation to file a Form 8-K is available on the SEC website.

Finally, we repeat our analysis on a further-reduced sample of standalone forecast revisions (Gleason and Lee, 2003; Chen and Matsumoto, 2006) in Column (4).²⁰ This filter inevitably introduces a look-ahead bias, but we continue to find a positive relation between *Analyst Centrality* and market reactions around forecast revisions. Overall, these findings are consistent with the notion that analysts learn from their brokerage colleagues. In the next section, we examine the learning channel more carefully.

3. The learning channel

Thus far, our evidence is consistent with the view that inter-colleague information exchange affects forecast performance. In this section, we design a series of tests to explore how securities analysts might learn from their colleagues.

3.1. Learning information on hard-to-value stocks

Since central analysts enjoy better access to the wide-ranging expertise of their colleagues, we hypothesize that their informational advantage will be more pronounced in the coverage of complex, hard-to-value stocks. For example, firms with intense R&D activities have greater uncertainty about their true level of innovation and future growth prospects. Processing information of firms with highly dispersed operations across numerous industry segments (e.g., Cohen and Lou, 2012) is also more complicated.

We test this prediction by interacting *Analyst Centrality* with two measures of forecasting complexity: *High R&D* and *Conglomerate* to assess the incremental effects of colleague learning on stocks. *High R&D* is an indicator equal one if either the ratio of R&D expenditure to total assets is above the yearly median or the firm has missing R&D expenditure data in Compustat following Koh and Reeb (2015).²¹ *Firm Complexity* is defined as the Herfindahl index of a firm’s sales across its industry segments multiplied by minus one; higher values of *Firm Complexity* correspond to greater firm complexity.

²⁰ Standalone forecast revisions have no other revisions released in the (-1, +1) day window of the revision.

²¹ Koh and Reeb (2015) find that a substantial fraction of firms that report missing R&D expenditure have patent activities. Our results also hold if we require sample firms to have non-missing R&D expenditure data.

[Insert Table 4]

Table 4 shows that central analysts are better able to forecast earnings for hard-to-value stocks. Column (1) reports a significantly negative coefficient on the interaction term *Analyst Centrality* $\times$ *High R&D*. Column (2) yields similar results in the model specification with the *Firm Complexity* measure. These findings suggest that the effects of *Analyst Centrality* on forecast accuracy are particularly evident in the valuation of complex stocks. Our conclusions are unchanged in Column (3) when we include both *High R&D* and *Firm Complexity* (and their interactions with *Analyst Centrality*) in the same model specification. Interestingly, the statistically significant and negative loadings on both interaction terms suggest that *High R&D* and *Firm Complexity* proxy for different dimensions of valuation complexity.

3.2. Learning from our colleague's mistakes

To gain deeper insights on the mechanisms of colleague learning, we examine how central analysts respond to the revelations of their colleagues' *ex post* forecast errors. Specifically, a central analyst may observe their colleagues' ex-post forecast errors on related firms whose earnings are announced, and in response, revises her own forecasts to incorporate these signals. This test is similar in spirit to the framework developed in Clement, Hales, and Xue (2011) who study how analysts revise their forecasts in response to news produced on firms in their coverage portfolios.²²

We hypothesize that a central analyst is likely to learn from her colleagues' revealed optimism and revise her forecast downwards. Such ex-post forecast errors are public information, but there are also countless other public signals that an analyst could potentially analyze to inform their revisions (i.e., curse of dimensionality). Given the shared workplace, an analyst is more likely to pay attention to a colleague's ex-post forecast error and incorporate this signal into her forecast revision. To measure the re-adjustment of the analyst's prior forecast, we construct *Analyst Revision*, a signed measure defined as the difference between an analyst's forecast revision value and her

²² The difference here is that brokerage colleagues never cover the same firm that is already followed by the analyst, so the information source is not based on the same firm.

prior forecast value, deflated by the absolute value of the latter. Therefore, a positive (negative) value of *Analyst Revision* reflects an increase (a decrease) in the analyst's forecast value from her previous forecast.

For a forecast revision of a given analyst, we collect instances of her colleagues' realized forecast errors that occurred within the past 30 days. We only retain the realized forecast errors of 1) colleagues who are directly connected to the analyst²³ (i.e. who cover the same sectors), and 2) colleagues who cover either the major suppliers or major customers of the analyst's firms.²⁴ Each realized forecast error is classified as optimistic if the forecasted value is above the actual reported EPS. We define *Colleague Optimism* as the proportion of optimistic forecast errors in the 30-day window. Given an analyst i in brokerage G covering firm f , colleague j , and making a forecast revision on date d ,

$$Colleague\ Optimism_{i,f,d} = \sum_{j \neq i, j \in G} \frac{Optimistic\ Revisions_{j,d}}{Total\ Revisions_{j,d}} \quad (4)$$

To test the learning channel, we estimate specification (5).

$$\begin{aligned} Analyst\ Revision_{i,f,d} = & \alpha + \beta_1 \cdot Analyst\ Centrality_{i,t} \times Colleague\ Optimism_{i,d} \\ & + \beta_2 \cdot Colleague\ Optimism_{i,d} \\ & + \beta_3 \cdot Analyst\ Centrality_{i,t} + \theta \cdot Controls_{i,f,t} + \varepsilon_{i,f,t} \end{aligned} \quad (5)$$

To assess the conditional effect of *Analyst Centrality* on *Analyst Revision*, we include the interaction term *Analyst Centrality* $\times$ *Colleague Optimism* as the key variable. While analysts may unconditionally revise their forecasts downwards after their colleagues' forecasts are revealed to be optimistic, we expect this effect to be stronger for central

²³ Our results also remain unchanged if we do not apply this filter, and collect all instances of directly or indirectly linked colleagues' realized forecast errors.

²⁴ We identify customer and supplier pairs from the business segment files of Compustat. In accordance with SFAS 14, public firms are required to disclose sales to their principal customers, defined as customers that contribute to at least 10 percent of the total revenue of the firm or if sales to a customer are material to the business of the firm. Principal customer names are manually matched to Compustat GVKEYs following the approach in Fee, Hadlock, and Thomas (2006). For customer names that are abbreviated, we hand-match and use industry affiliations to determine whether the customer is in Compustat. For the remaining unmatched customers, we check their corporate websites in the Directory of Corporate Affiliation (DCA) database to determine if the customer is a subsidiary of a listed firm. If so, we assign the customer to its parent's GVKEY. To ensure accuracy, we discard any customer name that cannot be unambiguously matched to a GVKEY.

analysts because they participate in inter-colleague information exchange more intensely. We include controls for analyst and brokerage characteristics, and the stock performance of the firm leading up to the forecast revision. Standard errors are double-clustered on the week and firm dimensions to capture dependencies in analysts' revisions.

[Insert Table 5]

Table 5 shows that central analysts respond more strongly to the revelations of their colleagues' *ex-post* forecast errors. Column (1) shows that *Colleague Optimism* predicts more negative forecast revisions. This suggests that, unconditionally, analysts' forecast revisions tend to be more negative in response to a higher incidence of optimistic errors made by their colleagues. In Column (2), we find that, conditional on *Colleague Optimism*, central analysts issue revisions that are more negative. This suggests that central analysts may have previously incorporated more of their colleagues' information in their own forecasts and thus need to unravel more erroneous information upon the revelations of their colleagues' optimism.

An alternative interpretation is that when an economic shock occurs, central analysts are too busy to update their forecasts due to the complexity of their coverage portfolios. Their revisions would then appear to lag their colleagues' revisions. However, given that the revisions of central analysts generate larger market reactions (see Table 3), they are more likely to contain novel and timely information. If the revisions of central analysts are delayed after the news shock, it is unlikely that those revisions would generate significantly larger market reactions.

The results may also reflect a central analyst's superior ability to process public information because ex-post forecast errors are publicly available information. To rule out the information-processing hypothesis, we introduce *Global Optimism*, the global analog of *Colleague Optimism*. To construct *Global Optimism*, we collect all realized forecast errors 1) in the 30 days leading up to the analyst's revision, and 2) that are not made by the analyst's colleagues. *Global Optimism* is the proportion of optimistic forecast errors

made by non-colleagues in the 30-day window.²⁵ If central analysts are simply more proficient at processing information, they should also respond incrementally to the publicly available optimism of non-colleagues.

We add *Global Optimism* and its interaction with *Analyst Centrality* to our specification in Column (3). Unconditionally, analysts' revisions tend to be more negative as *Global Optimism* increases; realized forecast errors, even those made by non-colleagues, may be indicative of economic surprises on a larger scale. However, the interaction between *Analyst Centrality* and *Global Optimism* no longer predicts analysts' revision patterns. This shows that central analysts do not respond incrementally to forecast errors made by non-colleagues. This finding disagrees with the hypothesis that central analysts are simply more proficient at processing public information. In contrast, the interaction between *Analyst Centrality* and *Colleague Optimism* remains negative and statistically significant, supporting the learning hypothesis.

3.3. Do all analysts within a brokerage benefit from information exchange?

An interesting question is whether the benefits of a network structure that facilitates information exchange accrue to all analysts in a brokerage. Both central and peripheral analysts may collectively benefit via two-way feedback in a brokerage. To examine the collective benefit from learning, we create a brokerage-level measure based on the concept of transitivity (Newman, Watts, and Strogatz, 2002) – *Network Density*.²⁶ A network is dense if its nodes are tightly interconnected with one another, thus facilitating information exchange (Smith-Doerr and Powell, 2010; Gibbert and Durand, 2009). Following this, we expect *Network Density* to be positively related to forecast performance as analysts exchange more information with colleagues in dense brokerages.

Since *Network Density* is a brokerage-level measure, we construct an analyst-level measure of forecast performance. For each analyst-year, we construct *Outperformance (%)* as the proportion of her forecasts with realized forecast errors that are lower than their

²⁵ We provide additional evidence in Internet Appendix Table 1 that our results hold when we increase the window length from 30 days to 60 days in the constructions of *Colleague Optimism* and *Global Optimism*.

²⁶ Appendix I presents the technicalities of *Network Density* and working examples.

firm-year averages.²⁷ Hence, analysts with higher values of *Outperformance (%)* are more accurate than competing analysts from other brokerages. We estimate a Tobit model following equation (6) and include controls for analyst experience and portfolio complexity. By construction, *Network Density* is normalized by network size but we also explicitly control for brokerage size. We cluster standard errors at the brokerage-year level.

$$\begin{aligned} \text{Outperformance } (\%)_{i,t} = & \alpha + \beta_1 \cdot \text{Analyst Centrality}_{i,t} \\ & + \beta_2 \cdot \text{Network Density}_{i,t} + \theta \cdot \text{Controls}_{i,f,t} + \varepsilon_{i,f,t} \end{aligned} \quad (6)$$

[Insert Table 6]

The evidence in Table 6 suggests that both central and peripheral analysts benefit from inter-colleague information exchange. Column (1) verifies that our earlier baseline findings hold at the aggregate analyst-year level as the relation between *Analyst Centrality* and *Outperformance (%)* is positive and statistically significant. Column (2) shows that analysts who reside in denser brokerage networks (*Network Density*) have higher forecast accuracy than competing analysts in sparser brokerage networks. This suggests that analysts benefit collectively from brokerage structures that facilitate information exchange. Notably, central analysts continue to receive incremental informational advantages as the estimated loading on *Analyst Centrality* remains positive and statistically significant. Since brokerages that cover only few or only one two-digit GICS sectors will mechanically have high values of *Network Density*, we repeat our analysis in Column (3) on a reduced sample of brokerages that cover at least three 2-digit GICS sectors.²⁸ We find that our results in Column (3) remain quantitatively and qualitatively unchanged.

Overall, we show that the benefits of inter-colleague information exchange accrue to both central and peripheral analysts. However, central analysts benefit disproportionately more than peripheral colleagues from colleague learning in brokerage networks.

²⁷ This construction helps to account for firm-specific heterogeneity of forecast difficulty across analysts in a brokerage. In untabulated results, our findings are qualitatively and quantitatively similar when we measure an analyst's performance as her median *Normalized Forecast Error* in the year.

²⁸ Our results are quantitatively and qualitatively similar when we adopt different values for this arbitrary cutoff.

3.4. Does the brokerage environment impact learning?

We expect that the brokerage environment will likely have a strong effect on inter-colleague information exchange. One important dimension is brokerage size. Larger and more prestigious brokerages may have higher quality information sharing because they attract analysts of higher ability and have more support resources (Clement, 1999). However, inter-colleague competition for year-end bonuses may also be more intense (e.g. Groysberg, Healy, and Maber, 2011), which may disincentivize analysts from collaboration or providing feedback.

[Insert Table 7]

We empirically examine the net effect of each dimension of the brokerage environment on information exchange in Table 7. Specifically, we split our sample on quartiles of *Brokerage Size* in every year, and analyze the effect of *Analyst Centrality* on forecast accuracy in these subsamples. Controlling for brokerage-year fixed effects, Panel A shows that the effect of *Analyst Centrality* on forecast accuracy is statistically insignificant among small brokerages. Instead, central analysts seem to exhibit higher forecast accuracy only among larger brokerages. On net, Panel A suggests that the benefits of high-quality information exchange exceed the potential effect of competition in large brokerages.²⁹

Another dimension that affects information exchange is network density, as discussed in the prior subsection. While denser brokerages have network architectures that facilitate information exchange, higher *Network Density* imply more overlaps in analysts' coverage portfolios, and potentially more intense competition within brokerages.³⁰

In Panel B, we find that the effect of *Analyst Centrality* on forecast accuracy increases with *Network Density*. Particularly, the effect of *Analyst Centrality* on forecast accuracy in the fourth quartile is about thrice (twice) the size of that in the first (third) quartile. Therefore, the upside of greater information exchange appears to outweigh the effect of

²⁹ While a slightly smaller coefficient on *Analyst Centrality* in the largest quartile of brokerages is consistent with the salience of within-brokerage competition, the coefficient in the largest quartile is not statistically different from those in the second and third quartiles.

³⁰ Refer to Section 3.3 and Appendix I.

intense competition among analysts that may accompany denser brokerage network structures.

We expect *Analyst Turnover* at brokerages to affect the incentives for inter-colleague information exchange in two ways. First, analyst turnover may proxy for competition at brokerages. On one hand, analysts may have weak tournament-like incentives to learn and outperform one another when competition is low. On the other hand, analysts may be disinclined to share value-relevant information with colleagues when competition is high. Second, an optimal level of analyst turnover may be required for new ideas to be generated and shared—low turnover may lead to stale ideas and no new knowledge to be learned, while new ideas may become lost quickly with rapid exits of brokerage colleagues. Panel C shows that *Analyst Centrality* does not load significantly in the extreme quartiles of *Analyst Turnover* but significantly predicts higher forecast accuracy in the second and third quartiles. This hump-shape pattern suggests that benefits accruing to central analysts are conditional on an optimal level of turnover in brokerages. Consistent with our hypothesis, both extremely high and low levels of *Analyst Turnover* can reduce the propensity or effectiveness of information exchange within brokerages.

Our results suggest that the brokerage environment affects information exchange in brokerages. On net, the outperformance of central analysts is more pronounced in larger brokerages and in brokerages with denser network structures and moderate levels of analyst turnover.

4. Causal effects of analyst centrality on forecast accuracy

An unaddressed issue is the matching process between the analyst and the brokerage house. Concerns relating to brokerage characteristics discussed above are partially addressed in our main tests with the inclusion of brokerage-year fixed effects in our baseline specification. However, it remains that an analyst with expertise in a particular industry may be more likely matched with a brokerage that specializes in that industry. To establish a causal relation between *Analyst Centrality* and forecast accuracy, we exploit brokerage mergers (Hong and Kacperczyk, 2010; Kelly and Ljungqvist, 2012) from years

2000 to 2007 as exogenous shocks to *Analyst Centrality* scores.³¹ For each merger event, we track all analysts who work at the same acquirer brokerages pre-event and post-event. We further require that each analyst covers the same firm before and after the merger. Therefore, our unit of observation in this quasi-natural experiment is an analyst-firm. Derrien and Kecskes (2013) find that most brokerage mergers are usually motivated by business reasons, suggesting that analysts who leave post-merger, if any, are not systematically different in forecast accuracy. Further, we exclude closures because it is plausible that only high-ability analysts manage to find new employment post-closure. Restricting our analysis to analysts in brokerage mergers alleviates concerns of selection bias due to unobserved analyst ability. Essentially, the shocks to *Analyst Centrality* of incumbent analysts emanate from changes in network structures after the mergers.

The treatments in this test are *Analyst Centrality Up* and *Analyst Centrality Down*. *Analyst Centrality Up* is an indicator that equal to one if the analyst’s average post-merger *Analyst Centrality* is higher than her average pre-merger value, and to zero otherwise. *Analyst Centrality Down* is defined symmetrically. Since brokerage mergers are scattered temporally, we use a difference-in-difference model, generalized to accommodate multiple treatment groups, and multiple shocks across time, following Autor (2003). We estimate the specification in (7).

$$\begin{aligned} \text{Normalized Forecast Error}_{g,t} = & \gamma_g + \tau_t + \sum_{j=-m, j \neq 0}^{+n} \beta_j \cdot D_{g,t}(t = k + j) \\ & + \theta \cdot \text{Controls} + \varepsilon_{g,t} \end{aligned} \quad (7)$$

where γ_g represents the group (analyst-firm) fixed effects and τ_t represents the year fixed effects. k is the time at which the brokerage merger occurs, the term $D_{g,t}$ is an indicator which switches to one in year t if the group receives the treatment. Note that this generalized model allows k to vary in different g . This is important because brokerage mergers in our sample are staggered across time. We let $j \neq 0$ because we skip the year

³¹ In Appendix III, we document the list of 17 brokerage mergers that we are able to match to our final sample.

of the brokerage merger. Visual inspection of the parallel trend assumption is tenuous in a model with shocks spread across time. Therefore, we include temporal leads and lags of the treatment in the model to test the assumption econometrically. Building m leads and n lags of the treatment effect β_j into the model allows us to estimate the pre-treatment dynamics (m leads) and post-treatment dynamics (n lags). The parallel trend assumption is fulfilled if β_j is not statistically significant for $j < 0$ – this suggests the absence of anticipatory effects of the treatment.

We use a six-year window centered on the brokerage merger event. We choose *Normalized Forecast Error* as the dependent variable because it is a normalized measure that allows for comparison of forecast accuracy across different analyst-firms. The key independent variables are the three temporal leads (*Pre-Treatment*) and three temporal lags (*Post-Treatment*) of the treatment. The m^{th} temporal *Pre-Treatment* is an indicator that equates to unity only in the m^{th} year before the brokerage merger and only if the analyst-firm is treated, and zero otherwise. Similarly, the n^{th} temporal *Post-Treatment* is an indicator that equates to unity only in the n^{th} year and only if the analyst-firm is treated, and zero otherwise. Apart from year dummies and analyst-firm dummies, we also add analyst time trends to help control for confounding heterogeneity.

[Insert Table 8]

We present results from the generalized difference-in-difference model in Table 8. In Column (1), we find that the coefficients of the *Pre-Treatment* indicators are statistically insignificant. This implies that before employment shocks, analysts who experienced *Analyst Centrality Up* display no systematic differences in forecast accuracy compared to analysts whose centrality scores did not improve. Econometrically, anticipatory effects of the treatment are unlikely to be the spurious driver of our results – this helps us to validate the parallel trend assumption. Crucially, this suggests that the assignment of analysts, into more central or less central network positions subsequent to the brokerage mergers, is independent on their pre-shock performance. On the other hand, we find that the positive effects of higher *Analyst Centrality* on forecast accuracy occur only after the treatment has been administered. The weak statistical significance of the $t+1$ *Post-*

Treatment indicator suggests the presence of post-merger adjustment costs which delay the gains to colleague learning.

Switching the treatment to *Analyst Centrality Down* in Column (2) yields symmetrically similar conclusions. In summary, our findings suggest that analysts who experience an increment in *Analyst Centrality* display higher forecast accuracy than those who did not. Moreover, we show that the differences in forecast accuracy manifest largely in the post-treatment period, and not in the pre-treatment period.

5. Discussion and robustness tests

In this section, we discuss and analyze alternative explanations to the learning hypothesis. We also perform a series of tests to ensure that our results are robust.

5.1. Does individual ability or talent explain our findings?

Analyst Centrality is likely to capture a component of talent or ability as a talented analyst may develop high centrality after surviving intense labor market competition. High-ability analysts are also likely to acquire greater coverage responsibilities over time. Therefore, our main tests include control for measures of analyst ability (i.e., experience and lowballing frequency).

To further disentangle the learning hypothesis from the ability explanation, we proxy for another dimension of ability with membership in the All-American Research Team (e.g., Leone and Wu, 2007). Studies show that membership in the Institutional Investor All-American Research Team (All-American) reflects analyst ability.³² We re-estimate the forecast accuracy and market reaction tests presented in Tables 2 and 3. Due to our limited data on All-American Research Team membership, we end the sample period in 2008 for these additional tests.

[Insert Table 9]

³² All-American analysts produce more accurate forecasts (Stickel, 1992), elicit stronger market reactions around their forecasts (Gleason and Lee, 2003), exhibit performance persistence (Leone and Wu, 2007), and attract more investment banking deal flows (Clarke et al., 2007)

Table 9 shows that our main results are unchanged after controlling for All-American status. We first report that our main *Analyst Centrality* results hold in this subsample period in Column (1). In Column (2), we include the *All-American* indicator in the model specification and find that *Analyst Centrality* continues to predict higher forecast accuracy. Consistent with prior literature, we also find that All-American analysts are more accurate. The market reaction tests in Panel B yield similar inferences. We first show in Column (1) that in this subsample period, central analysts still command larger market reactions around their forecast revisions. Column (2) shows our main conclusions are unchanged with the inclusion of the *All-American* indicator.

Together, the results suggest that the superior performance of central analysts is unlikely to be completely driven by greater individual ability or talent alone. Of course, it is plausible that talented and highly skilled analysts acquire higher centrality in brokerage networks over time. For instance, managers may design brokerage networks around high-ability analysts to leverage on their expertise. What distinguishes our learning hypothesis is that *Analyst Centrality* has incremental explanatory power beyond measures of ability found in the prior literature.

5.2. Regulation Fair Disclosure

Regulation Fair Disclosure (henceforth, Reg. FD) eliminated the practice of selective disclosure of information. Mohanram and Sunder (2006) argue that analysts in the post-Reg. FD regime need to heighten information discovery efforts to offset the impact of reduced access to firm management. It is possible that inter-colleague information exchange helped to fill the information void that emerged after Reg. FD.

[Insert Table 10]

Table 10 shows how the effect of *Analyst Centrality* changes with the implementation of Reg. FD. Column (1) shows that central analysts do not exhibit higher forecast accuracy than their peripheral colleagues in the pre-Reg. FD period. In contrast, Column (2) shows that *Analyst Centrality* has a positive and statistically significant impact on

forecast accuracy after the implementation of Reg. FD. Economically, the effect of *Analyst Centrality* in Column (2) is about 17 times larger than the effect in Column (1).

We repeat the split-sample analysis on market reactions to analysts’ forecast revisions. Following our empirical setup in Table 3, we also exclude forecast revisions if they coincide with issuances of Form 8-Ks or earnings announcements for this test. Columns (3) and (4) show that the forecast revisions of central analysts command higher market reactions both before and after the implementation of Reg. FD. Furthermore, the effect of *Analyst Centrality* on market reactions, conditional on deviations from the consensus forecast is also higher in both regimes. However, we do not find statistically significant pre- and post-difference in the effect of *Analyst Centrality*.

5.3. *Are only direct connections important?*

Analysts typically issue forecast revisions upon receiving value-relevant information. If information sharing occurs within a brokerage, this information shock is likely to diffuse to other colleagues. Such information may even flow to colleagues who are more than ‘one-step’ away in the brokerage network. As information propagates through the network, revision-activity may diffuse like a wave through the network. The information propagation may weaken with distance because 1) the information shocks may be less value-relevant further away from the origin and 2) analysts have limited information-processing capacities and may not ‘pass on’ the information with fidelity.

In the spirit of Ahern and Harford (2014), we construct a measure called *Closeness-weighted Revision Activity* that captures these waves of revision activity. For each analyst-brokerage pair, we find the earliest and latest forecast announcement dates. In between these two dates, we divide the analyst’s tenure at the brokerage by months. We exclude the first and last months of her tenure at the brokerage from our sample to avoid truncation issues. In each month, we construct *Revision Count* as the total number of forecast revisions made by the analyst. Next, we compute *Closeness-weighted Revision Activity* of the analyst as the weighted sum of all colleagues’ *Revision Count* in the

previous month. We use the reciprocal of the shortest-path distance³³ between an analyst i and her colleague j in brokerage G as the weight.

$$\text{Closeness weighted Revision Activity}_{i,t} = \sum_{j \in G, j \neq i} \frac{1}{\text{Distance}_{i,j,t}} \times \sum_r \text{Revision}_{r,j,t-1} \quad (8)$$

Since the dependent variable, *Revision Count*, is a discrete count variable, we estimate negative binomial regressions following equation (9). We include controls for analyst characteristics and brokerage size. Additionally, as the number of revisions made by an analyst is expected to increase with the number of firms under her coverage, we include *Firm Breadth* as an exposure variable in the model.

$$\begin{aligned} \text{Revision Count}_{i,t} = & \alpha + \beta_1 \cdot \text{Closeness weighted Revision Activity}_{i,t} \\ & + \theta \cdot \text{Controls}_{i,t} + \varepsilon_{i,t} \end{aligned} \quad (9)$$

5.4. Do central or peripheral analysts drive our findings?

We hypothesize that central analysts are better positioned to benefit from information exchange in brokerages. However, it is plausible that the empirical findings are solely driven by the underperformance of the most peripheral analysts. To rule out this possibility, we re-estimate our main tests using a tiered measure of *Analyst Centrality*. The results are available in the Internet Appendix. Overall, our evidence does not support the alternative hypothesis that the effect of *Analyst Centrality* is exclusively driven by a relative underperformance of the most peripheral analysts in brokerages.

5.5. Robustness: Alternative industry classification schemes in network construction

The broadness of two-digit GICS industry classification scheme may imprecisely measure connections between analysts, creating noise in our measurement. Such noise, however, tends to reduce the fidelity of the network structure and is likely to bias against uncovering evidence of inter-colleague information exchange. Alternatively, we could use a more granular industry classification scheme in our network construction, but this approach

³³ Consider the simplistic triad of analysts A , B , and C in the above example; analyst A has shortest-path distances of one and two to analysts B and C respectively.

risks the possibility of failing to capture actual inter-colleague linkages. The optimal choice of industry classification schemes concerns a tradeoff between measurement noise and the failure to identify linkages between colleagues.

To show that our baseline findings are not due to our choice of industry classification scheme, we use two industry classification schemes that are more granular – the two-digit SIC and four-digit GICS industry classification schemes – in our network construction methodology. We show that our main findings are not sensitive to the granularity of industry classification schemes adopted in our network construction methodology. The results are available in the Internet Appendix.

6. Conclusion

In knowledge-based professions, the quality and ability of one’s co-workers can have a significant impact on an employee’s productivity. Our study focuses specifically on the effects of learning from one’s colleagues by examining financial analysts. This setting is particularly well suited because information is a key input to an analyst’s performance.

Using GICS sector overlaps in their coverage portfolios, we build an information network within a brokerage house to capture potential information flow among colleagues. Our evidence suggests that analysts who are more centrally positioned in their brokerage networks produce equity research of higher quality. We provide finer evidence that central analysts are more likely to learn from their colleagues. If central analysts are tapping into their colleagues’ knowledge and expertise, we expect that they would be more likely to incorporate such information in their forecast revisions. Consistent with this argument, we find that high centrality analysts are also more likely to revise their forecasts after their colleagues’ forecasting mistakes are known. Controlling for analyst membership in the Institutional Investor All-America Research Team, supplementary analyses suggest that the positive effect of *Analyst Centrality* on superior forecasting outcomes is incremental to that of innate analyst ability.

The formation of within-brokerage networks may be endogenous. Therefore, to understand the causal relation between within-brokerage network centrality and forecast

accuracy, we exploit exogenous brokerage mergers from 2000 to 2007. These mergers introduce exogenous changes to the brokerage structures of the acquirers, and hence impact the *Analyst Centrality* of their analysts. We find that analysts who experienced increases in within-brokerage network centrality subsequently exhibited higher forecast accuracy.

REFERENCES

- Ahern, K., Harford, J., 2014. The importance of industry links in merger waves. *Journal of Finance* 69, 527-576.
- Anjos, F., Fracassi, C., 2015. Shopping for information? Diversification and the network of industries. *Management Science* 61, 161-183.
- Autor, D., 2003. Outsourcing at Will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of Labor Economics* 21, 1-42.
- Bae, K-H, Stulz R., Tan, H., 2008. Do local analysts know more? A cross-country study of the performance of local analysts and foreign analysts. *Journal of Financial Economics* 88, 581-606.
- Balakrishnan, K., Billings, M., Kelly, B., Ljungqvist, A., 2014. Shaping liquidity: On the causal effects of voluntary disclosure. *Journal of Finance* 69, 2237-2278.
- Bhojraj, S., Lee, C., Oler, D., 2003. What's my line? A comparison of industry classification schemes for capital market research. *Journal of Accounting Research* 41, 745-774.
- Bonacich, P., 1972. Factoring and weighting approaches to status scores and clique identification. *Journal of Mathematical Sociology* 2, 113-120.
- Boni, L., Womack, K., 2006. Analysts, industries, and price momentum. *Journal of Financial and Quantitative Analysis* 41, 85-109.
- Borgatti, S., 2005. Centrality and network flow. *Social networks* 27, 55-71.
- Bradley, D., Clarke, J., Lee, S., Ornathanalai, C., 2014. Are analysts' recommendations informative? Intraday evidence on the impact of time stamp delays. *Journal of Finance* 69, 645-673.
- Bradley, D., Gokkaya, S., Liu, X., 2017. Before an analyst becomes an analysts: Does industry experience matter? *Journal of Finance* 72, 751-792.
- Bradshaw, M., 2011. Analysts' forecasts: What do we know after decades of work? Unpublished working paper, Boston College.
- Brochet, F., Miller, G., Srinivasan, S., 2013. Do analysts follow managers who switch companies? An analysis of relationships in the capital markets. *The Accounting Review* 89, 451-482.
- Brown, L., 1997. Analyst forecasting errors: Additional evidence. *Financial Analyst Journal* 53, 81-88.
- Brown, L., 1998. Managerial behavior and the bias in analysts' earnings forecasts. Unpublished working paper, Georgia State University.
- Brown, L., 2001. A temporal analysis of earnings surprises: Profits versus losses. *Journal of Accounting Research* 39, 221-241.
- Brown, L., Call, A., Clement, M., Sharp, N., 2015. Inside the "black box" of sell-side financial analysts. *Journal of Accounting Research* 53, 1-47.
- Chang, X., Dasgupta, S., Hilary, G., 2006. Analyst following and financing decisions. *Journal of Finance* 61, 3009-3048.
- Chen, Q., Francis, J., Jiang, W., 2005. Investor learning about analyst predictive ability. *Journal of Accounting and Economics* 39, 3-24.

- Chen, Q, Jiang W., 2006. Analysts' weighting of private and public information. *Review of Financial Studies* 19, 319-355.
- Chen, S., Matsumoto, D., 2006. Favorable versus unfavorable recommendations: The impact on analyst access to management - provided information. *Journal of Accounting Research* 44, 657-689.
- Clarke, J., Khorana, A., Patel, A., Rau, P., 2007. The impact of all-star analyst job changes on their coverage choices and investment banking deal flow. *Journal of Financial Economics* 84, 713-737.
- Clement, M., 1999. Analyst Forecast Accuracy: Do ability, resources, and portfolio complexity matter? *Journal of Accounting and Economics* 27, 285-303.
- Clement, M., Tse, S., 2003. Do investors respond to analysts' forecast revisions as if forecast accuracy is all that matters? *The Accounting Review* 78, 227-249.
- Clement, M., Tse, S., 2005. Financial analyst characteristics and herding behavior in forecasting. *Journal of Finance* 60, 307-341.
- Clement, M., Koonce, L., Lopez, T., 2007. The roles of task-specific forecasting experience and innate ability in understanding analyst forecasting performance. *Journal of Accounting and Economics* 44, 378-398.
- Clement, M., Hales, J., Xue, Y., 2011. Understanding analysts' use of stock returns and other analysts' revisions when forecasting earnings. *Journal of Accounting and Economics* 51, 279-299.
- Cohen, L., Frazzini, A., Malloy, C., 2008. The small world of investing: Board connections and mutual fund returns. *Journal of Political Economy* 116, 951-979.
- Cohen, L., Frazzini, A., Malloy, C., 2010. Sell-side school ties. *Journal of Finance* 65, 1409-1437.
- Das, S., Levine, C., Sivaramakrishnan, K., 1998. Earnings predictability and bias in analysts' earnings forecasts. *The Accounting Review* 2, 277-294.
- Derrien, F., Kecskés, A., 2013. The real effects of financial shocks: Evidence from exogenous changes in analyst coverage. *Journal of Finance* 68, 1407-1440.
- Engelberg, J., Gao, P., Parsons, C., 2012. Friends with money. *Journal of Financial Economics* 103, 169-188.
- Fee, C., Hadlock, C., Thomas, S., 2006. Corporate equity ownership and the governance of product market relationships. *Journal of Finance* 61, 1217-1251.
- Fracassi, C., 2016. Corporate finance policies and social networks. *Management Science* forthcoming.
- Francis, J., Philbrick, D., 1993. Analysts' decisions as products of a multi-task environment. *Journal of Accounting Research* 31, 216-230.
- Freeman, L., 1979. Centrality in networks: Conceptual clarification. *Social Networks* 1, 215-239.
- Gibbert, M., Durand, T., 2009. *Strategic networks: Learning to compete*. John Wiley & Sons.
- Gleason, C., Lee, C., 2003. Analyst forecast revisions and market price discovery. *The Accounting Review* 78, 193-225.

- Gow, I., Ormazabal, G., Taylor, D., 2010. Correcting for cross-sectional and time-series dependence in accounting research. *The Accounting Review* 85, 483-512.
- Green, T., Jame, R., Markov, S., Subasi, M., 2014a. Broker-hosted investor conferences. *Journal of Accounting and Economics* 58, 142-166.
- Green, T., Jame, R., Markov, S., Subasi, M., 2014b. Access to management and informativeness of analyst research. *Journal of Financial Economics* 114, 239-255.
- Groysberg, B., Healy, P., Maber, D., 2011. What drives sell-side analyst compensation at high-status investment banks? *Journal of Accounting Research* 49, 969-1000.
- Heflin, F., Subramanyam, K., Zhang, Y., 2003. Regulation FD and the financial information environment: Early evidence. *The Accounting Review* 78, 1-37.
- Hilary, G., Hsu, C., 2013. Analyst forecast consistency, *Journal of Finance* 68, 271-297.
- Hilary, G., Shen, R., 2013. The role of analysts in intra-industry information transfer. *The Accounting Review* 88, 1265-1287.
- Hochberg, Y., Ljungqvist, A., Lu, Y., 2007. Whom you know matters: Venture capital networks and investment performance. *Journal of Finance* 62, 251-301.
- Honoré, B., 1992. Trimmed LAD and least squares estimation of truncated and censored regression models with fixed effects. *Econometrica* 60, 533-565.
- Hong, H., Kubik, J., Solomon, A., 2000. Security analysts' career concerns and herding of earnings forecast. *RAND Journal of Economics* 31, 121-144.
- Hong, H., Kubik, J., 2003. Analyzing the analysts: Career concerns and biased forecasts. *Journal of Finance* 58, 313-351.
- Hong, H., Kubik, J., Stein, J., 2004. Social interaction and stock-market participation. *Journal of Finance* 59, 137-163.
- Hong, H., Kubik, J., Stein, J., 2005. Thy neighbor's portfolio: Word-of-mouth effects in the holdings and trades of money managers. *Journal of Finance* 60, 2801-2824.
- Hong, H., Kacperczyk, M., 2010. Competition and bias. *Quarterly Journal of Economics* 125, 1683-1725.
- Hugon, A., Lin, A-P, Markov, S., 2016. Do equity analysts benefit from access to high quality debt research? Unpublished working paper, Arizona State University.
- Hugon, A., Kumar, A., Lin, A-P., 2016. Analysts, macroeconomic news, and the benefit of active in-house economists. *The Accounting Review* 91, 513-534.
- Hwang, L., Jan, C., Basu, S., 1996. Loss firms and analysts' earnings forecast errors. *Journal of Financial Statement Analysis* 1, 18-30.
- Hwang, B-H, Kim, S., 2009. It pays to have friends. *Journal of Financial Economics* 93, 138-158.
- Hwang, B-H, Liberti, J., Sturgess, J., 2016. Information sharing and spillovers: Evidence from financial analysts. Unpublished working paper, Northwestern University.
- Irvine, P., Simko, P.J., Nathan, S., 2004. Asset management and affiliated analysts' forecast. *Financial Analyst Journal* May/June 67-78
- Ivković, Z., Jegadeesh, N., 2004. The timing and value of forecast and recommendation revisions. *Journal of Financial Economics* 73, 433-463.

- Jacob, J., Lys, T., Neale, M., 1999. Expertise in forecasting performance of security analysts. *Journal of Accounting and Economics* 28, 51-82.
- Jiang, D., Kumar, A., Law, K., 2016. Political contributions and analyst behavior. *Review of Accounting Studies* 21, 37-88.
- Kadan, O., Madureira, L., Wang, R., Zach, T., 2012. Analysts' industry expertise. *Journal of Accounting and Economics* 54, 95-120.
- Ke, B., Yu, Y., 2006. The effect of issuing biased earnings forecasts on analysts' access to management and survival. *Journal of Accounting Research* 44, 965-999.
- Kelly, B., Ljungqvist, A., 2012. Testing asymmetric-information asset pricing models. *Review of Financial Studies* 25, 1366-1413.
- Kini, O., Mian, S., Rebello, M., Venkateswaran, A., 2009. On the structure of analyst research portfolios and forecast accuracy. *Journal of Accounting Research* 47, 867-909.
- Kuhnen, C., 2009. Business networks, corporate governance, and contracting in the mutual fund industry. *Journal of Finance* 64, 2185-2220.
- Kumar, A., 2010. Self-selection and the forecasting abilities of female equity analysts. *Journal of Accounting Research*, 48, 393-435
- Lang, M., Lins, K., Miller, D., 2004. Concentrated control, analyst following, and valuation: Do analysts matter most when investors are protected least? *Journal of Accounting Research* 42, 589-623
- Larcker, D., So, E., Wang, C., 2013. Boardroom centrality and firm performance. *Journal of Accounting and Economics* 55, 225-250.
- Law, K., 2013. Career imprinting: The influence of coworkers in early career. Unpublished working paper, Nanyang Technological University.
- Leone, A., Wu, J., 2007. What does it take to become a superstar? Evidence from Institutional Investor rankings of financial analysts. Unpublished working paper, University of Rochester
- Li, F., Mukherjee, A., Sen, R., 2017. Inside brokers. Unpublished working paper, Hong Kong University of Science and Technology.
- Lim, T., 2001. Rationality and analysts' forecast bias. *Journal of Finance* 56, 369-385.
- Loh, R., Mian, G., 2006. Do accurate earnings forecasts facilitate superior investment recommendations? *Journal of Financial Economics* 80, 455-483.
- Loh, R., Stulz, R., 2011. When are analyst recommendation changes influential? *Review of Financial Studies* 24, 593-627.
- Loh, R., Stulz, R., 2017. Is sell-side research more valuable during bad times? *Journal of Finance*, forthcoming.
- Malloy, C., 2005. The geography of equity analysis. *Journal of Finance* 60, 719-755.
- Marsden, P., 1990. Network data and measurement. *Annual Review of Sociology*, 435-463.
- Matsumoto, D., 2002. Management's incentives to avoid negative earnings surprises. *The Accounting Review* 77, 483-514.

- Merkley, K., Michaely, R., Pacelli, J., 2017. Does the scope of the sell-side analyst industry matter? An examination of bias, accuracy, and information content of analyst reports. *Journal of Finance*, forthcoming.
- Mikhail, M., Walther, B., Willis, R., 1997. Do security analysts improve their performance with experience? *Journal of Accounting Research* 35, 131-157.
- Mikhail, M., Walther, B., Willis, R., 1999. Does forecast accuracy matter to security analysts? *The Accounting Review* 74, 185-200.
- Mohanram, P., Sunder, S., 2006. How has Regulation FD affected the operations of financial analysts? *Contemporary Accounting Research* 23, 491-525.
- Newman, M., 2003. The structure and function of complex networks. *SIAM Review* 45, 167-256.
- Newman, M., Watts, D., Strogatz, S., 2002. Random graph models of social networks. *Proceedings of the National Academy of Sciences* 99, 2566-2572.
- Ozsoylev, H., Walden, J., Yavuz, M., Bildik, R., 2014. Investor networks in the stock market. *Review of Financial Studies* 27, 1323-1366.
- Petersen, M., 2009. Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies* 22, 435-480.
- Pool, V., Stoffman, N., Yonker, S., 2015. The people in your neighborhood: Social interactions and mutual fund portfolios. *Journal of Finance* 70, 2679-2732.
- Richardson, S., Teoh, S., Wysocki, P., 2004. The walk-down to beatable analyst forecasts: The role of equity issuance and insider trading incentives. *Contemporary Accounting Research* 21, 885-924.
- Shue, K., 2013. Executive networks and firm policies: Evidence from the random assignment of MBA peers. *Review of Financial Studies* 26, 1401-1442.
- Smith-Doerr, L., Powell, W., 2010. *The Handbook of Economic Sociology*. Princeton University Press.
- Soltes, E., 2014. Private interaction between firm management and sell - side analysts. *Journal of Accounting Research* 52, 245-272.
- Stickel, S., 1992. Reputation and performance among security analysts. *Journal of Finance* 47, 1811-1836.
- Valente, T., Coronges, K., Lakon, C., Costenbader, E., 2008. How correlated are network centrality measures? *Connections (Tor)* 28, 16-26.
- Westphal, J., Clement, M., 2008. Sociopolitical dynamics in relations between top managers and security analysts: Favor rendering, reciprocity, and analyst stock recommendations. *Academy of Management Journal* 51, 873-897.
- Womack, K., 1996. Do brokerage analysts' recommendations have investment value? *Journal of Finance* 51, 137-167.
- Yin, H., Zhang, H., 2014. Tournaments of financial analysts. *Review of Accounting Studies* 19, 573-605.

Appendix I

I. Degree Centrality

Degree centrality counts the number of colleagues that an analyst is immediately connected to in a brokerage network. In Figure 2, analyst **A** – represented by the red node – is immediately connected to four colleagues. Since degree centrality is increasing in the number of analysts **N** in a brokerage **G**, we normalize degree centrality by $N_G - 1$, or the maximum possible number of direct connections an analyst can have in a network.

Figure 2

Consider the network in Figure 2. Analysts **A** and **Z** are represented as red and blue nodes respectively. If **Z** cannot provide resources to **A**, **A** has the opportunity to ask her other three neighbors – represented as green nodes and is less dependent on any single colleague. However, **Z** does not have an alternative direct connection if **A** is unable to provide resources. Being connected to more colleagues also increases the likelihood that **A** will receive any information being circulated in the network. The formal mathematical definition of degree centrality for a given analyst **i** is as follows.

$$\text{Centrality}_{\text{Degree}}(i, G) = \frac{\sum_{j \neq i, j \in G}^{N_G - 1} f(i, j)}{N_G - 1} \quad \begin{array}{l} N_G = \text{Number of analysts in network } G \\ f(i, j) = \begin{cases} +1, & i \text{ is connected to } j \\ 0, & i \text{ is not connected to } j \end{cases} \end{array}$$

Following the above discussion, we show that analyst **A** has a higher degree centrality than analyst **Z**.

$$\text{Centrality}_{\text{Degree}} \text{ of analyst } \mathbf{A} = \frac{(3 \times 1_{\text{RedGreen}}) + (1 \times 1_{\text{RedBlue}}) + (3 \times 0_{\text{RedYellow}})}{9 - 1} = 0.500$$

$$\text{Centrality}_{\text{Degree}} \text{ of analyst } \mathbf{Z} = \frac{(1 \times 1_{\text{BlueRed}}) + (0 \times 1_{\text{BlueGreen}}) + (0 \times 0_{\text{BlueYellow}})}{9 - 1} = 0.125$$

II. Closeness Centrality

Closeness centrality is related to the distances between an analyst and all her colleagues (both immediately or not immediately connected) in a brokerage network. For example, in Figure 3a, the closeness centrality of analyst **A** – represented by the red node – is the reciprocal of the sum of its shortest-path distances to all other brokerage colleagues. Since the sum of shortest-path distances is increasing in the number of analysts, closeness centrality is normalized by the minimum possible sum of shortest-path distances, $N_G - 1$. For an analyst whose normalized closeness centrality equates to unity, all her colleagues are immediately connected to her.

Figure 3a

Figure 3b

Closeness centrality favors analysts who can access their colleagues at relatively shorter path lengths. An advantage of closeness centrality over degree centrality is that the former can also account for indirect connections in the network. To illustrate the difference between closeness and degree centralities, consider a brokerage network G in which analysts **A** and **Z** are represented as a red node in Figure 3a and a green node in Figure 3b respectively. The shortest-path distance of each colleague from analysts **A** and **Z** is indicated by the numbers in the orange nodes. In the calculations below, we show that analyst **Z** has a higher closeness centrality than analyst **A** even though both analysts have the same values of degree centrality. The formal mathematical definition of closeness centrality for a given analyst i is as follows.

$$N_G = \text{Number of analysts in network } G$$

$$\text{Centrality}_{\text{Closeness}}(i, G) = \frac{N_G - 1}{\sum_{j \neq i, j \in G} d(i, j)} \quad d(i, j) = \text{Shortest path distance between } i \text{ and } j$$

We now show that analysts **A** and **Z** have different values of closeness centrality despite having the same values of degree centrality.

$$\begin{aligned} \text{Centrality}_{\text{Closeness}} \text{ of analyst } \mathbf{A} &= \frac{8 - 1}{1 + 1 + 1 + 2 + 3 + 3 + 4} = 0.467 \\ \text{Centrality}_{\text{Closeness}} \text{ of analyst } \mathbf{Z} &= \frac{8 - 1}{1 + 1 + 1 + 2 + 2 + 3 + 3} = 0.538 \end{aligned}$$

III. Betweenness Centrality

Betweenness centrality is related to the number of geodesics (shortest paths) in the brokerage network that pass through an analyst.

Figure 4a

We first elaborate on the definition of geodesics with the aid of Figure 4a. For a pair of analysts, a geodesic between them is a path of the shortest possible length. An analyst pair may have more than one geodesic. Consider analysts **E** and **F**. Analyst **E** may reach analyst **F** via two paths – **EZF** and **EAF**. Since both paths have lengths of two, and the shortest possible path length between the analyst pair is two, both **EZF** and **EAF** qualify as geodesics.

The betweenness centrality of analyst **A** is the sum of proportions of all geodesics (not involving **A**) which pass through **A**. Revisiting the example of analyst pair **E** and **F**, there are 2 geodesics **EZF** and **EAF** between them but only path **EAF** passes through analyst **A**, yielding a proportion of 0.5. If we repeat this computation for all possible analyst pairs with reference to analyst **A**, we will obtain her betweenness centrality. Since these sums of proportions are increasing in the number of analysts, betweenness centrality is normalized by $\frac{2}{(N_G-1)(N_G-2)}$ the number of unique analyst pairs not involving **A**.

Figure 4b

In the context of this paper, betweenness centrality favors analysts whose coverage portfolios bridge the structural holes (Burt, 1992) between sector teams. In Figure 4b, the tightly knitted clique of green nodes represent the consumer discretionary sector team and the blue nodes represent the information technology sector team. If not for the presence of the red node, there would be two isolated information silos (one for each sector team) in the brokerage. The red node

in Figure 4b serves as a bridge between the two sector teams, and concurrently benefit by receiving information from both sector teams.

The formal mathematical definition of betweenness centrality for a given analyst i is as follows.

$$\text{Centrality}_{\text{Betweenness}}(i, G) = \frac{2}{(N_G - 1)(N_G - 2)} \sum_{x \neq i, y \neq i, x, y \in G} \frac{s(x, y|i)}{s(x, y)}$$

N_G = Number of analysts in network G
 $s(x, y)$ = Number of geodesics between any pair x and y
 $s(x, y|i)$ = Number of geodesics between any pair x and y passing through i

We now show that analyst $\mathbf{A}$ has a higher betweenness centrality than analyst $\mathbf{Z}$. The identities of analyst pairs are denoted by subscripts in the denominators of the fractions.

$$\text{Centrality}_{\text{Betweenness}} \text{ of analyst } \mathbf{A} = \frac{1}{6} \left(\frac{1}{1_{DE}} + \frac{1}{1_{DF}} + \frac{2}{2_{DZ}} + \frac{1}{2_{EF}} + \frac{0}{1_{EZ}} + \frac{0}{1_{FZ}} \right) = 0.583$$

$$\text{Centrality}_{\text{Betweenness}} \text{ of analyst } \mathbf{Z} = \frac{1}{6} \left(\frac{0}{1_{DA}} + \frac{0}{1_{DE}} + \frac{0}{1_{DF}} + \frac{0}{1_{AE}} + \frac{0}{1_{AF}} + \frac{1}{2_{EF}} \right) = 0.083$$

IV. Eigenvector Centrality

Eigenvector centrality is related to the notion that the centrality of an analyst is high if her connected colleagues are highly central in the brokerage network. A notable variant of eigenvector centrality is the PageRank algorithm used by the Google search engine to determine the importance of websites on the Internet. The underlying logic of the algorithm is that important websites are more likely to receive more web-links from other important websites. Similarly, an analyst has a high eigenvector centrality if her connected colleagues also possess high eigenvector centralities. Under eigenvector centrality, not only does the *quantity* of connections determine one's prominence in the network, but the *quality* of those connections matters as well.

To motivate the mathematical intuition behind eigenvector centrality, consider a simple network structure in Figure 5 and its corresponding adjacency matrix $\mathbf{M}_G$ below. The row-wise and column-wise sequences of the elements follow P , Q , R , and S . Where there is a connection, the element values equate to unity, and equate to zero otherwise. For example, element $\mathbf{m}_{1,2}$ equates to unity because P and Q are connected. On the other hand, element $\mathbf{m}_{1,3}$ equates to zero because there is no connection between P and R .

Figure 5

$$\mathbf{M}_G = \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix}$$

Next, suppose there is a 4 by 1 vector $\mathbf{K}_G$ of centrality values. For the purpose of exposition, we begin by choosing $\mathbf{K}_G$ to be a vector of un-normalized degree centralities. We arbitrarily choose un-normalized degree centralities as a starting point for its simplicity. For all purposes of this exposition, we could have defined $\mathbf{K}_G$ to be a vector of any other centrality values.

Formally, we write $\mathbf{K}_G$ as follows.

$$\mathbf{K}_G = \begin{bmatrix} 2 \\ 3 \\ 1 \\ 2 \end{bmatrix}$$

where $\mathbf{k}_{1,1}$, $\mathbf{k}_{2,1}$, $\mathbf{k}_{3,1}$, and $\mathbf{k}_{4,1}$ are the unnormalized degree centralities of P , Q , R , and S respectively

Now we perform the below matrix multiplication, we obtain another 4 by 1 matrix.

$$\mathbf{M}_G \cdot \mathbf{K}_G = \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 2 \\ 3 \\ 1 \\ 2 \end{bmatrix} = \begin{bmatrix} 0 \times 2 + 1 \times 3 + 0 \times 1 + 1 \times 2 \\ 1 \times 2 + 0 \times 3 + 1 \times 1 + 1 \times 2 \\ 0 \times 2 + 1 \times 3 + 0 \times 1 + 0 \times 2 \\ 1 \times 2 + 1 \times 3 + 0 \times 1 + 0 \times 2 \end{bmatrix} = \begin{bmatrix} 5 \\ 5 \\ 3 \\ 5 \end{bmatrix}$$

For each analyst in the network, the matrix multiplication sums up only the centralities of colleagues whom he is directly connected to. Alternatively, this multiplication is not only causing each analyst to receive her connected colleagues' centralities, but also causing her to distribute her centrality to connected colleagues concurrently. From the above example, the element $[1, 1]$, corresponding to $\mathbf{P}$, of the resulting matrix carries a value of 5, the cumulative centrality of her connections. This value is derived from $\mathbf{P}$'s connections – $\mathbf{Q}$ and $\mathbf{S}$ – who have degree centralities of 3 and 2 respectively. Now, we can interpret this matrix multiplication as 'spreading' the initial vector $\mathbf{K}_G$ across the network.

Suppose we repeat the multiplication to spread the initial vector $\mathbf{K}_G$ further, we obtain more 4 by 1 matrices.

$$\mathbf{M}_G \cdot \mathbf{M}_G \cdot \mathbf{K}_G = \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 5 \\ 5 \\ 3 \\ 5 \end{bmatrix} = \begin{bmatrix} 10 \\ 13 \\ 5 \\ 10 \end{bmatrix} \quad \mathbf{M}_G \cdot \mathbf{M}_G \cdot \mathbf{M}_G \cdot \mathbf{K}_G = \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 10 \\ 13 \\ 5 \\ 10 \end{bmatrix} = \begin{bmatrix} 23 \\ 25 \\ 13 \\ 23 \end{bmatrix}$$

While we observe that the values of elements in the matrices continue to become larger, one may speculate that there may exist an equilibrium where the proportion of total centralities held by each analyst is constant remains constant through additional stages of multiplication. At such an equilibrium, the centrality value of each analyst fully reflect the centralities of all connected colleagues. We can search for this equilibrium by choosing the initial vector $\mathbf{K}_G$. Upon closer inspection, the search for this equilibrium solution is in fact a search for the eigenvector of the adjacency matrix $\mathbf{M}_G$.

If we had replaced each element of the centrality vector $\mathbf{K}_G$ with values of eigenvector centralities, the brokerage network can be described as follows.

$$\begin{aligned} \text{Centrality}_{\text{Eigenvector}}(i, G) &= \text{element } v_i \text{ of matrix } \mathbf{V}_G \\ \mathbf{V}_G &\text{ is an eigenvector of the network's adjacency matrix } \mathbf{M}_G \\ \text{or, } \quad \mathbf{M}_G \cdot \mathbf{V}_G &= \lambda \cdot \mathbf{V}_G \\ \text{where } \lambda &\text{ is a scalar} \end{aligned}$$

Suppose we perform a matrix multiplication between $\mathbf{M}_G$ and its eigenvector $\mathbf{V}_G$.

$$\mathbf{M}_G \cdot \mathbf{V}_G = \lambda \cdot \mathbf{V}_G$$

Multiply the resulting vector by $\mathbf{M}_G$.

$$\mathbf{M}_G \cdot \lambda \cdot \mathbf{V}_G = \lambda \cdot \mathbf{M}_G \cdot \mathbf{V}_G = \lambda \cdot \lambda \cdot \mathbf{V}_G = \text{constant} \cdot \mathbf{V}_G$$

We observe that even with incremental steps of matrix multiplication, the resulting vector is always a scalar inflation of the starting vector $\mathbf{V}_G$. Thus, we can say that the vector $\mathbf{V}_G$ fully

represents the cumulative centrality (or prominence) of analysts and their colleagues in the brokerage network.

V. Network Density

Network density is a measure that describes the overall interconnectedness of analysts in a brokerage network. Unlike centrality measures which are analyst-centric, network density is a brokerage-level measure that relates to the network structure of the brokerage. Analysts in a brokerage with high network density would be strongly interconnected to one another. To define interconnectedness among analysts, we rely on the concept of transitivity (Newman, Watts, and Strogatz, 2002). The concept of transitivity readily maps to our intuition of interconnectedness when we consider the colloquialism – “we are interconnected if the friend³⁴ of my friend also happens to be my friend”. Notice that the colloquialism elicits an image of a triangle with a node at each of its three vertices. This idea is at the heart of defining network transitivity³⁵ (or network density).

Consider a simple network P in Figure 6.

Figure 6

We begin by counting the number of paths of length = 2. In total, there are 12 paths of length = 2: one originating from A (ABC), two originating from B (BCE ; BCD), three originating from C (CBA ; CDE ; CED), three originating from D (DCB ; DCE ; DEC), and three originating from E (ECB ; ECD ; EDC). Notice that paths are sequences so that, for example, ECB is distinct from BCE .

From the paths of length = 2, a subset of them form closed triangles (recall the colloquialism above). In this network, there are 6 closed triangles: two originating from C ($CDEC$; $CEDC$), two originating from D ($DCED$; $DECD$), and two originating from E ($ECDE$; $EDCE$). As the case with paths of length = 2, closed triangles are also sequences so that, for example, $ECDE$ is distinct from $EDCE$.

Notice that closed triangles can only be formed from three nodes that sit on the same path of length = 2. Following this, the transitivity of the network in Figure 6 is the ratio of the number of closed triangles to the number of paths of length = 2.

³⁴ To be clear, a “friend” from the perspective of an analyst is a colleague who shares a connection with her.

³⁵ Network transitivity is also known as global clustering coefficient.

$$\text{Transitivity of } \mathbf{P} = \frac{6}{12} = 0.500$$

Consider another simple network $\mathbf{Q}$ in Figure 7.

Figure 7

Again, we begin by counting the number of paths of length = 2. In total, there are 28 paths of length = 2: six originating from A (ABC ; ACB ; ACD ; ACE ; ADC ; ADE), five originating from B (BAC ; BAD ; BCA ; BCD ; BCE), six originating from C (CAB ; CAD ; CBA ; CDA ; CDE ; CED), six originating from D (DAB ; DAC ; DCA ; DCB ; DCE ; DEC), and five originating from E (ECA ; ECB ; ECD ; EDA ; EDC).

Of the set of paths with length = 2, 18 of them form closed triangles: four originating from A ($ABCA$; $ACBA$; $ACDA$; $ADCA$), two originating from B ($BACB$; $BCAB$), six originating from C ($CABC$; $CBAC$; $CADC$; $CDAC$; $CDEC$; $CEDC$), four originating from D ($DACD$; $DCAD$; $DCED$; $DECD$), and two originating from E ($ECDE$; $EDCE$).

$$\text{Transitivity of } \mathbf{Q} = \frac{18}{28} = 0.643$$

From a quick visual inspection of Figures 6 and 7, we observe that the nodes in network $\mathbf{Q}$ are more interconnected than those in network $\mathbf{P}$. This observation agrees with our calculations that $\mathbf{Q}$ has a higher transitivity score than $\mathbf{P}$.

Appendix II. Variable Definitions

Analyst Centrality	Standardized PCA-extracted factor score of 4 network centrality measures – Degree Centrality, Closeness Centrality, Betweenness Centrality and Eigenvector Centrality.
Degree Centrality	Degree Centrality is related to the number of colleagues with whom the analyst is immediately connected in the brokerage network (see Appendix I for details on centrality measures).
Closeness Centrality	Closeness Centrality is related to the path distances between the analyst and all her colleagues in the brokerage network (see Appendix I for details on centrality measures).
Betweenness Centrality	Betweenness Centrality is related to the number of geodesics (shortest paths) in the brokerage network that pass through the analyst (see Appendix I for details on centrality measures).
Eigenvector Centrality	Eigenvector Centrality is related to the notion that an analyst is central if her connected colleagues are also central in the brokerage network (see Appendix I for details on centrality measures).
Revision Frequency	Total number of revisions made by the analyst in the firm-year.
Horizon	Number of days elapsed between the analyst's firm-year forecast and the actual earnings announcement. We exclude all forecasts that are more than 365 days or less than 30 days from the actual earnings announcement.
Boldness	The deviation of an analyst's forecast from the consensus forecast in the previous year.
General Experience	Logarithm of the number of months between an analyst's first appearance in I/B/E/S and her firm-year forecast.
Firm Experience	Logarithm of the number of months between an analyst's earliest forecast of the firm in I/B/E/S and her firm-year forecast.
Brokerage Experience	Number of months elapsed between an analyst's first appearance in the brokerage in I/B/E/S and her firm-year forecast.
Firm Breadth	Number of unique firms covered by the analyst in the year.
Industry Breadth	Number of unique 2-digit GICS (Global Industry Classification Standard) sectors covered by the analyst in the year.
Lowball	Number of times over the past 3 years that lowballing forecasts were issued for the firm by the analyst. For a forecast to be classified as a lowballing one, 3 conditions must be met. First, the forecast value must be below the actual EPS value. Second, the forecast error (absolute difference between forecast value and actual EPS value) must be either greater than 3 cents or higher than 5% of the actual EPS value. Third, to reduce the likelihood of mistaking forecasting difficulty for lowballing behavior, the difference between the forecast value and the consensus value must be greater than 3 cents or higher than 5% of the consensus value.
Brokerage Size	Logarithm of the number of analysts employed by the brokerage in the year.
Loss	Indicator that equates to unity if the actual EPS of the firm is negative, and equates to zero otherwise.
Forecast Dispersion	Standard deviation of firm-year forecast values among analysts.
Analyst Coverage	Number of analysts who have contributed at least 1 forecast in the firm-year.
Leverage	Sum of short-term debt and long-term borrowings, scaled by total assets of the firm.
Market-to-Book Ratio	Ratio of firm book value to its market capitalization in the year.
Total Assets	Total assets of the firm.
ROA Volatility	Volatility of the ratio of net income to total firm assets over the past 36 months.

Appendix II. (Continued)

Analyst Revision	Difference between the forecast revision value and the previous forecast value, scaled by the absolute value of the latter variable. <i>Analyst Revision</i> is a signed variable – positive (negative) values of <i>Analyst Revision</i> reflect increments (declines) in the analyst's forecasted value from her previous forecast.
------------------	---

Forecast Error	Absolute difference between the analyst's final firm-year forecast and the actual firm-year EPS value.
Normalized Forecast Error	<i>Forecast Error</i> scaled by the average <i>Forecast Error</i> across analysts in the firm-year.
Number of Forecasts	The number of competitor analysts who issue a forecast in the (-1, +1) window surrounding an analyst's forecast.
Bold Forecast	Following Clement and Tse (2005), an analyst's forecast revision is classified as bold when it is either above both her prior forecast and pre-revision consensus forecast or below both her prior forecast and pre-revision consensus forecast.
Revision Magnitude	Absolute difference between the analyst's revision value and her prior forecast value, scaled by the absolute value of the latter variable.
Number of Forecasts Colleague Optimism	Number of firm-year forecasts issued by all analysts in the week of the analyst's forecast revision.
	For a forecast revision of a given analyst, we collect instances of her colleagues' realized forecast errors that occurred within the past 30 days. We only retain the realized forecast errors of 1) colleagues who are directly connected to the analyst (i.e. who cover the same industries), and 2) colleagues who cover either the major suppliers or major customers of the analyst's firms. For each forecast error of the analyst's colleagues, we classify them as <i>optimistic</i> if the forecasted value is above the actual earnings. If the forecast error is zero, it is neither optimistic nor pessimistic but is still counted in the window. <i>Colleague Optimism</i> is the proportion of <i>optimistic</i> forecast errors issued by colleagues in the 30-day window.
Global Optimism	For a given forecast revision of an analyst, we collect all instances (including non-colleagues) realized forecast errors in the same 2-digit GICS sector within the past 30 days. Alternatively, the firms covered by all analysts (including non-colleagues) in the same 2-digit GICS sector must have announced their actual earnings in the same 30-day window. For each forecast error, we classify it as optimistic if the forecasted value is above the actual earnings. If the forecast error is zero, it is neither optimistic nor pessimistic but is still counted in the window. <i>Global Optimism</i> is the proportion of optimism forecast errors issued by all non-colleague analysts in the 30-day window.
All-American	Indicator that equates to unity if the analyst belongs to the Institutional Investor All-America Research Team in the year, and equates to zero otherwise.
Abs. CAR [-1, +1]	Market-adjusted cumulative abnormal returns, centered on the forecast revision date.
Network Density	This is a network-level measure. <i>Network Density</i> is the ratio of the number of loops of length three (closed triad) to the number of paths of length two (triads). By definition, a closed triad is also a triad. See Appendix I for a detailed working example.
Outperformance	A firm-year forecast is defined as an outperforming forecast if its forecast error is lower than the mean forecast error. <i>Outperformance (%)</i> of an analyst is the percentage of her forecasts in a year that are categorized as outperforming forecasts.

Appendix II. (Continued)

Firm Complexity	Using segment data from Compustat, we determine the number of 2-digit SIC industries that each firm operates in, following the procedure in Cohen and Lou (2012). For each firm, we then compute the Herfindahl-Hirschman index (HHI) of its industry segment sales. A high HHI indicates that firm sales are concentrated in one or few industries, and implies that the firm has low complexity. A symmetrical reasoning may be applied to a low HHI. To facilitate interpretation, we multiply the HHI by -1 and define this measure as <i>Firm Complexity</i> . A firm with a more positive <i>Firm Complexity</i> score indicates that firm sales are dispersed across multiple industry segments, and that the firm is complex.
High R&D	An indicator that equates to unity if either 1) R&D expenses are missing, or 2) R&D intensity (ratio of R&D to total assets) is above the year median, and equates to zero otherwise.
Revision Count	For each analyst-brokerage pair, we find her earliest and latest forecast announcement dates. In between these two dates, we divide the analyst's tenure at the brokerage by months. We exclude the first and last months of her tenure from our sample to avoid truncation issues. In each month, Revision Count is the total number of forecast revisions made by the analyst.
Closeness-weighted Revision Activity	For a given analyst, we find the shortest-path length between her and her brokerage colleagues. For example, if analyst A is directly connected to analyst B, and analyst C is in turn connected to analyst B, the shortest-path length between analysts A and C is two. Using the reciprocals of shortest-path lengths as weights, <i>Closeness-weighted Revision Activity</i> of an analyst is the weighted sum of her colleague's <i>Revision Count</i> in the previous month.
Distance1 Revision Activity	<i>Distance1 Revision Activity</i> of an analyst is the total <i>Revision Count</i> of her colleagues who are separated by a shortest-path length of one, normalized by the number of such colleagues.
Distance2 Revision Activity	<i>Distance2 Revision Activity</i> of an analyst is the total <i>Revision Count</i> of her colleagues who are separated by a shortest-path length of two, normalized by the number of such colleagues.

Appendix III. Brokerage Merger Events in Table 8

We compile the list of brokerage mergers from Hong and Kacperczyk (2010), and Kelly and Ljungqvist (2012). The table below documents the 17 mergers in our final sample.

Acquirer Brokerage	Target Brokerage	Date	
Salomon Smith Barney	Schroder Wertheim	Apr	2000
Soundview	Wit Capital	May	2000
Paine Webber	J.C. Bradford	Jun	2000
Credit Suisse First Boston	Donaldson Lufkin & Jenrette	Oct	2000
UBS Warburg Dillon Reed	Paine Webber	Dec	2000
J.P. Morgan	Chase Manhattan	Dec	2000
J.P. Morgan Chase	Hambrecht & Quist	Jan	2001
Dresdner Bank	Wasserstein Perella	Feb	2001
Fahnestock	Josephthal Lyon & Ross	Sep	2001
First Union	Wachovia Securities	Oct	2001
UBS	Schwab Soundview	Oct	2004
Janney Montgomery Scott	Parker/Hunter	Mar	2005
Merrill Lynch	Advest	Dec	2005
Merrill Lynch	Petrie Parkman	Dec	2006
Stifel Financial	Ryan Beck & co	Jan	2007
Fox-Pitt Kelton	Cochran, Caronia Securities	Sep	2007
Fahnestock	CIBC World Markets	Jan	2008

Figure 1. Example of a Within-brokerage Network

This figure maps the network structure of I/B/E/S brokerage identifier (481) in the year 2005. Each node represents an analyst in the brokerage. The numbers below each node identify the GICS sectors covered by the analyst.³⁶

³⁶ The 2-digit GICS sector codes map to the economic sectors as follows: 20 (industrials), 25 (consumer discretionary), 35 (healthcare), 45 (information technology), and 50 (telecommunication services). For a full map of all eleven 2-digit GICS sector codes, refer to <https://www.msci.com/gics>.

Figure 2. Distribution of GICS Sector Coverage

This figure maps the network structure of I/B/E/S brokerage identifier (481) in the year 2005. Each node represents an analyst in the brokerage. We weight the size of nodes by different measures of their centrality in the same network; larger nodes are more central. The color intensity of nodes also increases with centrality.

Table 1.
Descriptive Statistics

Panel A reports the summary statistics of unique analyst-year pairs. Since an analyst covers multiple firms in a year, we report summary statistics at the analyst-year level to show analyst characteristics that drive the formation of within-brokerage networks. Panel B reports Pearson correlations among selected variables of unique analyst-year pairs. Panel C reports results of principal component analysis on the four centrality measures. Panel D reports the summary statistics of the baseline sample used in Table 2 Column (1). Due to data limitations, descriptive statistics of *All-American* is based on a sample truncated in the year 2008. In Panel A and B, *Firm Experience* and *Brokerage Experience* are averaged at the analyst-year level.

Panel A. Summary Statistics (unique analyst-year pairs)

	N	Mean	S.D.	p10	p25	p50	p75	p90
Number of Connections	60694	16.023	14.565	3	6	12	22	35
Industry Breadth	60694	1.564	0.914	1	1	1	2	3
Firm Breadth	60694	9.660	7.188	1	4	9	14	19
General Experience (mth)	60694	52.286	51.445	2	12	37	78	129
Firm Experience (mth)	60694	22.671	26.588	0	2	14	34	59
Brokerage Experience (mth)	60694	2.470	2.981	0	0	12	48	72
Brokerage Size	60694	58.753	51.107	11	19	42	88	123
All-American	42558	0.102	0.303	0	0	0	0	1
Analyst Centrality	60694	0.007	0.984	-1.034	-0.720	-0.229	0.597	1.512
Degree Centrality	60694	0.371	0.257	0.103	0.169	0.295	0.511	0.774
Closeness Centrality	60694	0.553	0.193	0.345	0.452	0.536	0.646	0.807
Betweenness Centrality	60694	0.021	0.063	0	0	0	0.010	0.058
Eigenvector Centrality	60694	0.151	0.137	0.006	0.033	0.115	0.232	0.354

Table 1. (Continued)
Panel B. Pearson Correlations (unique analyst-year pairs)

		[A]	[B]	[C]	[D]	[E]	[F]	[G]	[H]	[I]	[J]	[K]
Analyst Centrality	[A]	1										
Degree Centrality	[B]	0.949	1									
Closeness Centrality	[C]	0.889	0.887	1								
Betweenness Centrality	[D]	0.518	0.324	0.277	1							
Eigenvector Centrality	[E]	0.848	0.746	0.592	0.378	1						
Industry Breadth	[F]	0.571	0.494	0.437	0.521	0.476	1					
Firm Breadth	[G]	0.165	0.158	0.18	0.128	0.078	0.307	1				
General Experience	[H]	0.008	-0.004	-0.009	0.048	0.011	0.061	0.344	1			
Firm Experience	[I]	-0.053	-0.066	-0.062	0.014	-0.037	0.013	0.287	0.801	1		
Brokerage Experience	[J]	0.021	0.015	0.029	0.045	-0.01	0.111	0.339	0.641	0.628	1	
Brokerage Size	[K]	-0.376	-0.409	-0.176	-0.144	-0.46	-0.119	-0.042	0.001	0.051	0.085	1
All-American	[L]	-0.121	-0.133	-0.07	-0.031	-0.141	0.01	0.203	0.192	0.245	0.208	0.261

Panel C. Principal Component Analysis

	Component (1)	Component (2)	Component (3)	Component (4)
Degree Centrality	0.948	-0.208	0.085	-0.225
Closeness Centrality	0.886	-0.258	0.345	0.171
Betweenness Centrality	0.512	0.851	0.118	-0.002
Eigenvector Centrality	0.846	-0.011	-0.528	0.074
Eigenvalue	2.661	0.834	0.419	0.085
Variance explained (%)	66.5	20.9	10.5	2.1
Cumulative variance (%)	66.5	87.4	97.9	100

Table 1. (Continued)
Panel D. Summary Statistics (baseline sample)

	N	Mean	S.D	p10	p25	p50	p75	p90
Analyst Centrality	274673	0.083	0.958	-0.983	-0.619	-0.107	0.671	1.459
Degree Centrality	274673	0.385	0.247	0.111	0.192	0.326	0.538	0.750
Closeness Centrality	274673	0.569	0.177	0.377	0.469	0.555	0.663	0.796
Betweenness Centrality	274673	0.028	0.068	0.000	0.000	0.000	0.023	0.082
Eigenvector Centrality	274673	0.155	0.130	0.008	0.045	0.126	0.237	0.346
Horizon	274673	116.442	90.613	21	55	98	125	273
Brokerage Size	274673	60.444	48.143	12	22	47	91	119
Revision Frequency	274673	3.664	2.868	1	2	3	5	7
Boldness	274673	-0.002	0.256	-0.110	-0.025	0.001	0.023	0.093
General Experience	274673	80.503	54.008	20	37	69	114	160
Firm Experience	274673	42.068	38.132	10	13	28	58	96
Firm Breadth	274673	15.451	7.737	7	11	15	19	24
Industry Breadth	274673	1.789	1.039	1	1	1	2	3
Lowball	274673	0.159	0.366	0	0	0	0	1
Loss	274673	0.125	0.330	0	0	0	0	1
Analyst Coverage	274673	16.945	9.702	6	9	15	23	31
Forecast Error	274673	0.222	33.947	0.005	0.013	0.040	0.110	0.290

Table 2.
The Effect of Analyst Centrality on Forecast Accuracy

This table presents results from OLS regressions on the effect of *Analyst Centrality* on forecast accuracy. The dependent variable *Normalized Forecast Error* is defined as the absolute difference between an analyst's last firm-year forecast and the actual value scaled by the average forecast error in the firm-year. *Analyst Centrality* is the standardized PCA-extracted factor score of 4 network centrality measures – *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality* and *Eigenvector Centrality*. Columns (2) and (4) exclude firm financial variables because they are invariant at the firm-year level. Detailed definitions of other variables are in Appendix II. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 2. (Continued)

	(1)	(2)	(3)
	OLS: Normalized Forecast Error	OLS: Normalized Forecast Error	OLS: Normalized Forecast Error
Analyst Centrality	-0.019*** (0.003)	-0.014*** (0.003)	-0.009*** (0.003)
Revision Frequency	-0.006*** (0.001)	-0.005*** (0.001)	-0.009*** (0.002)
Boldness	-0.019*** (0.006)	-0.019*** (0.006)	-0.020** (0.009)
Horizon	0.003*** (0.000)	0.003*** (0.000)	0.004*** (0.000)
General Experience	-0.019*** (0.002)	-0.020*** (0.002)	-0.019*** (0.003)
Firm Experience	0.003 (0.003)	0.004 (0.003)	-0.011*** (0.003)
Firm Breadth	-0.000 (0.000)	-0.000 (0.000)	-0.001*** (0.000)
Industry Breadth	0.014*** (0.003)	0.006** (0.003)	0.012*** (0.002)
Lowball	0.067*** (0.005)	0.066*** (0.005)	0.084*** (0.006)
Loss		-0.011** (0.005)	
Analyst Coverage		-0.003*** (0.000)	
Leverage		0.010 (0.007)	
Book-to-Market Ratio		0.002 (0.002)	
Total Assets		0.008*** (0.001)	
ROA Volatility		-0.000 (0.000)	
Brokerage Size			0.003 (0.002)
Observations	274,673	274,673	274,673
R-squared	0.176	0.177	0.277
Brokerage-Year fixed effects	Yes	Yes	No
Firm-Year fixed effects	No	No	Yes
Brokerage-Year cluster	Yes	Yes	No
Firm-Year cluster	No	No	Yes
Analyst-Firm cluster	Yes	Yes	Yes

Table 3.

The Effect of Analyst Centrality on Market Reactions to Forecast Revisions

This table presents results from OLS regressions on the effect of *Analyst Centrality* on market reactions to forecast revisions. The dependent variable is the absolute 3-day market-adjusted cumulative abnormal returns (*Abs. CAR* $[-1, +1]$), centered on the forecast revision date. *Analyst Centrality* is the standardized PCA-extracted factor score of 4 network centrality measures – *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality* and *Eigenvector Centrality*. *Consensus Deviation* is the absolute difference between an analyst’s revision value and the prevailing consensus, normalized by the absolute value of the latter variable. Stock return variables are unitized in percentage points. Detailed definitions of other variables are in Appendix II. In Column (3), we exclude a forecast revision if the firm files SEC 8-Ks or has an earnings announcement in the $(-1, 0)$ day window of the revision. In Column (4), we use a further restricted sample of revisions that 1) do not occur in the $(-1, 0)$ day window of the revision, and 2) are standalone. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 3. (Continued)

	(1)	(2)	(3)	(4)
	OLS: Abs. CAR [-1, +1]	OLS: Abs. CAR [-1, +1]	OLS: Abs. CAR [-1, +1]	OLS: Abs. CAR [-1, +1]
Exclude 8-K events	No	No	Yes	Yes
Exclude earnings announcements	No	No	Yes	Yes
Standalone forecasts only	No	No	No	Yes
Analyst Centrality	0.338*** (0.022)	0.289*** (0.023)	0.212*** (0.020)	0.156*** (0.023)
Bold Forecast	0.473*** (0.019)	0.315*** (0.016)	0.275*** (0.017)	0.185*** (0.026)
Revision Magnitude	0.000 (0.001)	0.000 (0.000)	0.001 (0.001)	0.001*** (0.000)
General Experience		0.117*** (0.017)	0.053*** (0.017)	0.031 (0.021)
Firm Experience		-0.063*** (0.009)	-0.050*** (0.009)	-0.017 (0.012)
Firm Breadth		-0.027*** (0.002)	-0.021*** (0.002)	-0.019*** (0.002)
Industry Breadth		-0.140*** (0.021)	-0.143*** (0.018)	-0.093*** (0.023)
Brokerage Size		0.221*** (0.020)	0.227*** (0.019)	0.174*** (0.022)
Analyst Coverage		-0.031*** (0.004)	-0.015*** (0.003)	-0.016*** (0.003)
Leverage		0.724*** (0.135)	0.565*** (0.121)	0.460*** (0.139)
Total Assets		-0.783*** (0.020)	-0.630*** (0.017)	-0.598*** (0.019)
Book-to-Market Ratio		0.389*** (0.091)	0.350*** (0.085)	0.305*** (0.109)
ROA Volatility		0.001 (0.002)	0.004* (0.002)	0.007 (0.005)
Abs. CAR [-5, -2]		-0.022*** (0.005)	-0.028*** (0.004)	-0.030*** (0.006)
Number of Forecasts		0.204*** (0.008)	0.306*** (0.014)	
Observations	1,211,241	1,211,241	637,008	130,626
R-squared	0.006	0.126	0.127	0.093
Week cluster	Yes	Yes	Yes	Yes
Firm cluster	Yes	Yes	Yes	Yes

Table 4.
Analyst Centrality and Hard-to-Value Firms

This table presents results from OLS regressions on the effect of *Analyst Centrality* on the accuracy of hard-to-value firms. The dependent variable is *Normalized Forecast Error*, defined as *Raw Forecast Error* scaled by the average forecast error in the firm-year. *Raw Forecast Error* is defined as the absolute difference between an analyst’s last firm-year forecast and the actual value. The key independent variable is *Analyst Centrality* and its interactions with *High R&D* and *Firm Complexity*. *Analyst Centrality* is the standardized PCA-extracted factor score of 4 network centrality measures – *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality* and *Eigenvector Centrality* (see Appendix I for details on centrality measures). *High R&D* is an indicator that equates to unity if either 1) R&D expenses are missing, or 2) R&D intensity (ratio of R&D to total assets) is above the year median, and equates to zero otherwise. To construct *Firm Complexity*, we compute the Herfindahl index of firm sales across the industry segments in which the firm operates (Cohen and Lou, 2012). *Firm Complexity* is this Herfindahl index multiplied by negative one. Therefore, a more positive value of *Firm Complexity* correspond to higher complexity in a firm’s operations. Control variables used in the Column (2) of Table 2 are also included in the model estimations. However, their estimated coefficients are suppressed to facilitate presentation. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 4. (Continued)

	(1)	(2)	(3)
	OLS: Normalized Forecast Error	OLS: Normalized Forecast Error	OLS: Normalized Forecast Error
Analyst Centrality [A]	-0.009** (0.004)	-0.013*** (0.003)	-0.008* (0.004)
[A] x High R&D	-0.007* (0.004)		-0.007** (0.004)
[A] x Firm Complexity		-0.020** (0.009)	-0.022** (0.010)
High R&D	-0.009** (0.004)		-0.009** (0.004)
Firm Complexity		-0.005 (0.009)	-0.007 (0.009)
Revision Frequency	-0.006*** (0.001)	-0.005*** (0.001)	-0.006*** (0.001)
Boldness	-0.017*** (0.006)	-0.017*** (0.006)	-0.017*** (0.006)
Horizon	0.003*** (0.000)	0.003*** (0.000)	0.003*** (0.000)
General Experience	-0.020*** (0.003)	-0.020*** (0.003)	-0.020*** (0.003)
Firm Experience	0.003 (0.003)	0.004 (0.003)	0.003 (0.003)
Firm Breadth	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Industry Breadth	0.005** (0.003)	0.006** (0.003)	0.005** (0.003)
Lowball	0.067*** (0.005)	0.067*** (0.005)	0.067*** (0.005)
Loss	-0.009* (0.005)	-0.010** (0.005)	-0.009* (0.005)
Analyst Coverage	-0.003*** (0.000)	-0.003*** (0.000)	-0.003*** (0.000)
Leverage	0.009 (0.007)	0.010 (0.007)	0.009 (0.007)
Book-to-Market Ratio	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)
Total Assets	0.008*** (0.001)	0.009*** (0.001)	0.008*** (0.001)
ROA Volatility	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Observations	274,673	274,673	274,673
R-squared	0.178	0.178	0.178
Brokerage-Year fixed effects	Yes	Yes	Yes
Brokerage-Year cluster	Yes	Yes	Yes
Analyst-Firm cluster	Yes	Yes	Yes

Table 5.
Learning from Colleagues' Ex-Post Forecast Errors

This table presents results from OLS regressions on the influence of colleague's forecast errors on analyst revisions. The dependent variable *Analyst Revision* is defined as a signed variable equal to the difference between the revised forecast value and the previous forecast value, scaled by the absolute value of the previous forecast value. A positive (negative) value of *Analyst Revision* reflects an increment (a decline) in the analyst's forecasted value from her previous forecast. The key independent variables are *Colleague Optimism*, *Global Optimism*, and their respective interactions with *Analyst Centrality*. For a forecast revision of a given analyst, we collect instances of her colleagues' realized forecast errors that occurred within the past 30 days. We only retain the realized forecast errors of 1) colleagues who are directly connected to the analyst (i.e. who cover the same industries), and 2) colleagues who cover either the major suppliers or major customers of the analyst's firms. In other words, the firms covered by the analyst's colleagues must have announced their actual earnings in the same 30-day window. For each forecast error of the analyst's colleagues, we classify them as *optimistic* if the forecasted value is above the actual earnings. *Colleague Optimism* is the proportion of *optimistic* forecast errors in the 30-day window. To construct *Global Optimism*, we collect all realized forecast errors outside of the analyst's brokerage in the same 30-day window. Following this, *Global Optimism* is the proportion of *optimistic* forecast errors made by non-colleagues in the 30-day window. *Analyst Centrality* is the standardized PCA-extracted factor score of 4 network centrality measures – *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality* and *Eigenvector Centrality*. Detailed definitions of other variables are in Appendix II. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 5. (Continued)

	(1)	(2)	(3)
	OLS: Analyst Revision	OLS: Analyst Revision	OLS: Analyst Revision
Colleague Optimism [A]	-0.848*** (0.207)	-0.819*** (0.207)	-0.462*** (0.173)
Global Optimism [B]			-7.289*** (1.727)
[A] x Analyst Centrality		-0.402** (0.159)	-0.356** (0.158)
[B] x Analyst Centrality			-0.055 (0.810)
Analyst Centrality	-0.314** (0.152)	-0.168 (0.154)	-0.179 (0.300)
Number of Colleague Forecasts	-0.004 (0.003)	-0.005* (0.003)	-0.005* (0.003)
Number of Global Forecasts			0.000 (0.000)
General Experience	0.128 (0.084)	0.123 (0.083)	0.143* (0.083)
Firm Experience	-0.064 (0.054)	-0.065 (0.054)	-0.059 (0.054)
Firm Breadth	-0.022** (0.009)	-0.022** (0.009)	-0.021** (0.009)
Industry Breadth	-0.062 (0.120)	-0.072 (0.119)	-0.063 (0.119)
Brokerage Size	-0.045 (0.126)	-0.002 (0.133)	-0.042 (0.130)
Analyst Coverage	0.098*** (0.009)	0.098*** (0.009)	0.099*** (0.009)
Abs. CAR [-5, -2]	-0.359*** (0.031)	-0.358*** (0.031)	-0.355*** (0.030)
Observations	698,697	698,697	698,697
R-squared	0.012	0.012	0.013
Week cluster	Yes	Yes	Yes
Firm cluster	Yes	Yes	Yes

Table 6.
Do Analysts Perform Better in Denser Networks?

This table presents results from Tobit regressions on the effect of network density on forecast accuracy. The dependent variable, *Outperformance (%)*, is the percentage of an analyst's forecasts in a year that exhibit forecast error that is lower than the mean forecast error. *Analyst Centrality* is the standardized PCA-extracted factor score of four network centrality measures – *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality* and *Eigenvector Centrality*. *Network Density* is a brokerage-level variable that measures the density of the within-brokerage network in a year. Refer to Appendix I for details on its construction and working examples. In column (4), we exclude brokerages that cover fewer than three 2-digit GICS sectors. Detailed definitions of other variables are in Appendix II. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

	(1) TOBIT: Outperformance (%)	(2) TOBIT: Outperformance (%)	(3) TOBIT: Outperformance (%)
Exclude brokerages that cover less than 3 sectors	No	No	Yes
Analyst Centrality	0.779*** (0.214)	0.619*** (0.225)	0.898*** (0.276)
Network Density		3.661* (1.615)	4.644** (2.307)
General Experience	0.795*** (0.183)	0.790*** (0.183)	0.819*** (0.261)
Firm Experience	0.375** (0.181)	0.376** (0.181)	0.374* (0.223)
Industry Breadth	-2.557*** (0.219)	-2.430*** (0.226)	-2.658*** (0.213)
Firm Breadth	0.241*** (0.025)	0.243*** (0.025)	0.243*** (0.028)
Brokerage Size	0.029*** (0.003)	0.030*** (0.003)	0.030*** (0.005)
Observations	60,694	60,694	58,826
Brokerage-year cluster	Yes	Yes	Yes

Table 7.
Does the Brokerage Environment Impact Learning?

This table presents results from OLS regressions on the effect of *Analyst Centrality* on forecast accuracy in different brokerage environments. The dependent variable *Normalized Forecast Error* is defined as the absolute difference between an analyst’s last firm-year forecast and the actual value scaled by the average forecast error in the firm-year. *Analyst Centrality* is the standardized PCA-extracted factor score of four network centrality measures – *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality* and *Eigenvector Centrality*. In Panel (A), Panel (B), and Panel (C), we split the sample by *Brokerage Size*, *Analyst Turnover*, and *Network Density*, respectively. *Brokerage Size* is the number of analysts employed by the brokerage in the year. *Analyst Turnover* of a brokerage is the total number of analysts who left and joined the brokerage in the preceding year, scaled by the brokerage size in the preceding year. *Network Density* is a brokerage-level variable that measures the density of the within-brokerage network in a year. Refer to Appendix I for details on its construction and working examples. Detailed definitions of other variables are in Appendix II. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 7. (Continued)

Panel A. Brokerage Size Subsample Analysis				
	(1)	(2)	(3)	(4)
	Normalized Forecast Error	Normalized Forecast Error	Normalized Forecast Error	Normalized Forecast Error
Brokerage size quartile	Small	Q2	Q3	Big
Analyst Centrality	0.005 (0.008)	-0.019*** (0.007)	-0.020*** (0.006)	-0.013*** (0.005)
Control variables	Yes	Yes	Yes	Yes
Observations	26,183	24,404	47,198	176,888
R-squared	0.138	0.206	0.192	0.173
Brokerage-Year fixed effects	Yes	Yes	Yes	Yes
Brokerage-Year cluster	Yes	Yes	Yes	Yes
Analyst-Firm cluster	Yes	Yes	Yes	Yes
Panel B. Brokerage Network Density Subsample Analysis				
	(1)	(2)	(3)	(4)
	Normalized Forecast Error	Normalized Forecast Error	Normalized Forecast Error	Normalized Forecast Error
Network density quartile	Low	Q2	Q3	High
Analyst Centrality	-0.010* (0.005)	-0.013** (0.005)	-0.016** (0.007)	-0.032** (0.016)
Control variables	Yes	Yes	Yes	Yes
Observations	102,827	112,388	50,058	9,400
R-squared	0.171	0.168	0.198	0.226
Brokerage-Year fixed effects	Yes	Yes	Yes	Yes
Brokerage-Year cluster	Yes	Yes	Yes	Yes
Analyst-Firm cluster	Yes	Yes	Yes	Yes
Panel C. Analyst Turnover Subsample Analysis				
	(1)	(2)	(3)	(4)
	Normalized Forecast Error	Normalized Forecast Error	Normalized Forecast Error	Normalized Forecast Error
Analyst turnover quartile	Low	Q2	Q3	High
Analyst Centrality	-0.009 (0.006)	-0.022*** (0.006)	-0.014** (0.006)	-0.012 (0.008)
Control variables	Yes	Yes	Yes	Yes
Observations	76,396	87,536	74,126	36,615
R-squared	0.144	0.172	0.182	0.230
Brokerage-Year fixed effects	Yes	Yes	Yes	Yes
Brokerage-Year cluster	Yes	Yes	Yes	Yes
Analyst-Firm cluster	Yes	Yes	Yes	Yes

Table 8.
The Causal Effect of Analyst Centrality on Forecast Accuracy:
Shocks to Analyst Centrality due to Brokerage Mergers

This table presents results from a difference-in-difference model with multiple groups, and multiple shocks across time. We define shocks as brokerage mergers (Hong and Kacperczyk, 2010; Kelly and Ljungqvist, 2012) from years 2000 to 2007. Appendix III contains the list of 17 brokerage mergers that we are able to match to our final sample. For each merger event, we track all analysts who work at the acquirer before and after the mergers. We further require that each analyst covers the same firm before and after the merger. Therefore, our unit of observation in this quasi-natural experiment is an analyst-firm. The treatment variables are *Increase (Decrease) in Analyst Centrality*, an indicator variable which equates to unity if the average post-merger *Analyst Centrality* is higher (lower) than the average pre-merger *Analyst Centrality*, and equates to zero otherwise. Since visual inspection to validate the parallel trend assumption is tenuous in a model with shocks spread across time, we include temporal leads and lags of the treatment in the model to test the assumption econometrically (e.g. Autor (2003)). We use a (-3, +3) year window centered on the merger event. The dependent variable in this table is *Normalized Forecast Error*, defined as *Forecast Error* scaled by the average forecast error in the firm-year. The key independent variables are the temporal leads (*Pre-Treatment*) and lags (*Post-Treatment*) of the treatment. We also include analyst time trends to help control for confounding trends. Analyst time trends are *General Experience*, *Brokerage Size*, *Firm Breadth*, and *Industry Breadth*. Detailed definitions of the variables are in Appendix II. To facilitate presentation, coefficient estimates of analyst time trends are not presented. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 8. (Continued)

	(1)	(2)
	OLS: Normalized Forecast Error	OLS: Normalized Forecast Error
Treatment indicator	Increase in Analyst Centrality	Decrease in Analyst Centrality
<u>Leads and Lags of Treatment:</u>		
Pre-Treatment (t-3)	0.110 (0.113)	-0.193 (0.120)
Pre-Treatment (t-2)	0.136 (0.089)	-0.117 (0.077)
Pre-Treatment (t-1)	-0.060 (0.072)	-0.078 (0.070)
Post-Treatment (t+1)	-0.111 (0.078)	0.024 (0.079)
Post-Treatment (t+2)	-0.161* (0.083)	0.191** (0.088)
Post-Treatment (t+3)	-0.209** (0.094)	0.419*** (0.095)
Analyst-Firm dummies	Yes	Yes
Year dummies	Yes	Yes
Analyst time trends	Yes	Yes
Analyst-Firm cluster	Yes	Yes
Observations	4,515	4,515
R-squared	0.280	0.285

Table 9.
Analyst Centrality and Analyst Ability

This table presents results from OLS regressions. Due to data limitations on the identities of All-American analysts, our sample period ends in 2008. The dependent variables are *Normalized Forecast Error* and *Abs. CAR*. *Normalized Forecast Error* is defined as *Forecast Error* scaled by the average forecast error in the firm-year. *Abs. CAR* is the absolute 3-day market-adjusted cumulative abnormal returns, centered on the forecast revision date. The key independent variables are *All-American*, *Analyst Centrality*, and their interactions with *Consensus Deviation*. *All-American* is an indicator variable that equates to unity if the analyst belongs to the Institutional Investor All-America Research Team in the year, and equates to zero otherwise. *Analyst Centrality* is the standardized PCA-extracted factor score of 4 network centrality measures – *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality* and *Eigenvector Centrality* (see Appendix I for details on centrality measures). *Consensus Deviation* is the absolute difference between an analyst’s revision value and the prevailing consensus, normalized by the absolute value of the latter variable. Stock return variables are unitized in percentage points. In Panel B, we exclude a forecast revision if the firm either files SEC 8-Ks or has an earnings announcement in the (-1, 0) day window of the revision. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 9. (Continued)
Panel A. Analyst Centrality, Ability, and Forecast Accuracy

	(1) OLS: Normalized Forecast Error	(2) OLS: Normalized Forecast Error
Analyst Centrality	-0.011** (0.005)	-0.010** (0.005)
All-American		-0.033*** (0.007)
Revision Frequency	-0.008*** (0.001)	-0.008*** (0.001)
Boldness	-0.031*** (0.010)	-0.031*** (0.010)
Horizon	0.004*** (0.000)	0.004*** (0.000)
General Experience	-0.018*** (0.004)	-0.017*** (0.004)
Firm Experience	0.006 (0.005)	0.008 (0.005)
Firm Breadth	-0.001 (0.000)	-0.000 (0.000)
Industry Breadth	0.004 (0.004)	0.004 (0.004)
Lowball	0.064*** (0.008)	0.064*** (0.008)
Loss	-0.036*** (0.007)	-0.036*** (0.007)
Analyst Coverage	-0.001** (0.000)	-0.001** (0.000)
Leverage	0.004 (0.012)	0.005 (0.012)
Book-to-Market Ratio	0.001 (0.006)	0.001 (0.006)
Total Assets	0.007*** (0.002)	0.007*** (0.002)
ROA Volatility	0.022 (0.016)	0.021 (0.016)
Observations	166,023	166,023
R-squared	0.224	0.224
Brokerage-Year fixed effects	Yes	Yes
Brokerage-Year cluster	Yes	Yes
Analyst-Firm cluster	Yes	Yes

Table 9. (Continued)
Panel B. Analyst Centrality, Ability, and Market Reactions

	(1) OLS: Abs. CAR [-1, +1]	(2) OLS: Abs. CAR [-1, +1]
Exclude 8-K events	Yes	Yes
Exclude earnings announcements	Yes	Yes
Analyst Centrality	0.191*** (0.020)	0.194*** (0.020)
All-American		-0.166*** (0.037)
Bold Forecast	0.120*** (0.016)	0.122*** (0.016)
Revision Magnitude	2.135*** (0.111)	2.130*** (0.111)
General Experience	0.183*** (0.018)	0.185*** (0.018)
Firm Experience	-0.042*** (0.008)	-0.040*** (0.008)
Firm Breadth	-0.015*** (0.002)	-0.015*** (0.002)
Industry Breadth	-0.156*** (0.017)	-0.158*** (0.017)
Brokerage Size	0.189*** (0.017)	0.214*** (0.017)
Analyst Coverage	-0.002 (0.002)	-0.002 (0.002)
Leverage	0.221*** (0.065)	0.236*** (0.065)
Total Assets	-0.441*** (0.010)	-0.440*** (0.010)
Book-to-Market Ratio	0.165*** (0.037)	0.164*** (0.037)
ROA Volatility	0.004** (0.002)	0.004** (0.002)
Number of Forecasts	0.203*** (0.008)	0.203*** (0.008)
Abs. CAR [-5, -2]	-0.016*** (0.003)	-0.016*** (0.003)
Observations	364,943	364,943
R-squared	0.093	0.094
Week cluster	Yes	Yes
Firm cluster	Yes	Yes

Table 10.
Analyst Centrality and Regulation Fair Disclosure

This table presents results from OLS regressions. The dependent variable in Columns (1) and (2) is *Normalized Forecast Error*, defined as *Raw Forecast Error* scaled by the average forecast error in the firm-year. *Raw Forecast Error* is defined as the absolute difference between an analyst's last firm-year forecast and the actual value. The dependent variable in Columns (3) and (4) is the 3-day market-adjusted cumulative abnormal returns, centered on the forecast revision date and unitized in percentage points. The key independent variable is *Analyst Centrality* – the standardized PCA-extracted factor score of 4 network centrality measures – *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality* and *Eigenvector Centrality* (see Appendix I for details on centrality measures). The pre-Regulation Fair Disclosure sample period is from 1995 to 2000 (based on the fiscal year ends of firms). The post-Regulation Fair Disclosure sample period begins from 2001. Control variables used in the Column (2) of Table 2 or Column (3) of Table 3 are also included in the model estimations. However, their estimated coefficients are suppressed to facilitate presentation. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

	(1) OLS: Normalized Forecast Error	(2) OLS: Normalized Forecast Error	(3) OLS: Abs. CAR [-1, +1]	(4) OLS: Abs. CAR [-1, +1]
Sample period	Pre-Reg. FD	Post-Reg. FD	Pre-Reg. FD	Post-Reg. FD
Exclude 8-K events	-	-	Yes	Yes
Exclude earnings announcements	-	-	Yes	Yes
Analyst Centrality	-0.001 (0.008)	-0.017*** (0.004)	0.213*** (0.033)	0.190*** (0.024)
Control variables	Yes	Yes	Yes	Yes
Observations	50,999	223,643	167,135	469,855
R-squared	0.206	0.172	0.111	0.140
Brokerage-Year fixed effects	Yes	Yes	No	No
Brokerage-Year cluster	Yes	Yes	No	No
Analyst-Firm cluster	Yes	Yes	No	No
Week cluster	No	No	Yes	Yes
Firm cluster	No	No	Yes	Yes

Internet Appendix for “Do security analysts learn from their colleagues?”

This is the Internet Appendix for “*Do security analysts learn from their colleagues?*” which contains additional tests, robustness checks, and discussions referenced in the main text.

1. *Diffusion of revision activity*

Table 1 of the Internet Appendix shows that our results in Table 5 hold when we increase the window length from 30 days to 60 days in the constructions of *Colleague Optimism* and *Global Optimism*.

2. *Diffusion of revision activity*

Table 2 of the Internet Appendix shows that revision activity diffuses within the brokerage network. Column (1) shows a significantly positive relation between *Closeness-weighted Revision Activity* and *Revision Count*. This indicates that the analyst’s own revision activity increases in response to colleagues’ intense revision activity. Since the sum of colleagues’ revision activity mechanically increases with the number of analysts employed at the brokerage, we adopt a size-normalized measure of *Closeness-weighted Revision Activity*. Specifically, we scale the revision activity at each distance tier by the number of colleagues residing in that tier before applying the above distance-weighted aggregation. Our conclusion remains unchanged when we adopt this alternative measure in Column (2). Economically, a unit increase in size-normalized *Closeness-weighted Revision Activity* is associated with a 1.5% increase in the analyst’s *Revision Count*.

In Columns (3) and (4), we show that analysts are influenced more strongly by the revision activity of proximate (shorter paths) colleagues in the network. We define *Distance1 (Distance2) Revision Activity* of an analyst as the total *Revision Count* of all colleagues who are at one path-length (two path-lengths) away in the network, scaled by the number of colleagues residing at one path-length (two path-lengths). Comparing their estimated coefficients in Columns (3) and (4), *Distance1 Revision Activity* has an economic magnitude approximately two times larger than that of *Distance2 Revision*

Activity.³⁷ The Welch-Satterthwaite t-test indicates that this difference in economic magnitudes is statistically significant at the 1% level. Column (3) shows that an analyst has a 3.2% increase in his revision activity for an additional forecast revision made by a colleague at one path-length away. In Column (4), this effect shrinks to 1.1% for forecast revisions made by colleagues at two path-lengths away. This suggests that the diffusion of revision activity degrades with distance.

3. High and low centrality

To test whether our findings apply generally to all analysts, we introduce a tiered-decomposition of *Analyst Centrality*. Specifically, we assign *High (Low) Centrality* indicators to analysts in the top (bottom) tercile of *Analyst Centrality* in every year. Thereafter, we repeat our baseline analysis with these indicators in lieu of *Analyst Centrality* as key variables. This empirical design allows us to benchmark the performance of the most peripheral and most central analysts with their moderately central colleagues.

The results in Table 3 of the Internet Appendix indicate that the relation between *Analyst Centrality* and forecast accuracy applies generally to both peripheral and central analysts. In Column (1), we estimate the effect of *Low Centrality* on forecast accuracy in the absence of *High Centrality*. Thus, the benchmark category in Column (1) comprises analysts who are either moderately central or highly central. We find that *Low Centrality* loads positively and significantly on forecast error. In line with our expectation, peripheral analysts exhibit lower accuracy than their more central colleagues do. Following this, we replace *Low Centrality* with *High Centrality* in Column (2). Consistent with our prior findings that central analysts have higher forecast accuracy, we find a negative and statistically significant loading on *High Centrality*. In Column (3), we estimate the effects of *Low Centrality* and *High Centrality* on forecast error simultaneously. With moderately central analysts as the benchmark category, we expect *High Centrality* to load insignificantly on forecast error if our prior findings are driven exclusively by the most

³⁷ We refrain from including both *Distance1 Revision Activity* and *Distance2 Revision Activity* in the same specification because they are highly collinear with each other. We only present results up to a path-length of two because analysts in many brokerages are connected to any colleague by two path-lengths.

peripheral analysts. Our results indicate otherwise; *Low Centrality* and *High Centrality* load positively and negatively on forecast error when they are jointly included in the regression specification. In Columns (4) through (6), we repeat the preceding analysis but *Low (High) Centrality* corresponds to the bottom (top) quartile of *Analyst Centrality* instead. Our conclusions remain unchanged from Columns (1) through (3).

4. Brokerage networks based on 2-digit SIC or 4-digit GICS overlaps

Table 4 of the Internet Appendix reports results on the relations between *Analyst Centrality (SIC-2D or G-GROUP)* and 1) forecast accuracy, 2) market reactions around forecast revisions. Control variables are included per the model specifications in Tables 2 and 3 but their estimated coefficients are not tabulated for brevity. The definition of *Analyst Centrality (SIC-2D)* remains unchanged from Sections 1.2 and 1.3, except that it is computed on within-brokerage networks in which two analysts are connected if they cover at least one 2-digit SIC industry in common. *Analyst Centrality (G-GROUP)* is constructed based on the 4-digit GICS industry group classification scheme. Our conclusions remain unchanged with the adoption of either *Analyst Centrality (SIC-2D)* or *Analyst Centrality (G-GROUP)* in our tests. We continue to find that, controlling for brokerage-year fixed effects, more central analysts exhibit higher forecast accuracy. Market reactions around forecast revisions are also increasing in *Analyst Centrality (SIC-2D)* and *Analyst Centrality (G-GROUP)*. The economic magnitudes on the effect of analyst centrality are also comparable to those in Tables 2 and 3. Overall, we show that our main findings are not sensitive to the granularity of industry classification schemes adopted in our network construction methodology.³⁸

³⁸ However, when compared to Tables 2 and 3, our sample sizes in Table 11 are marginally smaller with the use of 2-digit SIC in our network construction methodology. The increased granularity of the industry classification scheme causes some brokerages in our sample to have no within-brokerage connections among their analysts. Since *Analyst Centrality* is undefined for all analysts in these brokerages, they are omitted from our sample.

Internet Appendix Table 1.
Additional Results: Learning from Colleagues' Ex-Post Forecast Errors

This table supplements the main results in Table 4. The dependent variable is *Analyst Revision*. *Analyst Revision* is a signed variable which equates to the difference between the revised forecast value and the previous forecast value, scaled by the absolute value of the latter. Hence, a positive (negative) value of *Analyst Revision* reflects an increment (a decline) in the analyst's forecasted value from her previous forecast. The key independent variables are *Colleague Optimism*, *Global Optimism*, and their respective interactions with *Analyst Centrality*. For a forecast revision of a given analyst, we collect instances of her colleagues' realized forecast errors that occurred within the past 60 days. We only retain the realized forecast errors of 1) colleagues who are directly connected to the analyst (i.e. who cover the same industries), and 2) colleagues who cover either the major suppliers or major customers of the analyst's firms. In other words, the firms covered by the analyst's colleagues must have announced their actual earnings in the same 60-day window. For each forecast error of the analyst's colleagues, we classify them as *optimistic* if the forecasted value is above the actual earnings. *Colleague Optimism* is the proportion of *optimistic* forecast errors in the 60-day window. To construct *Global Optimism*, we collect all realized forecast errors outside of the analyst's brokerage in the same 60-day window. Following this, *Global Optimism* is the proportion of *optimistic* forecast errors made by non-colleagues in the 60-day window. *Analyst Centrality* is the standardized PCA-extracted factor score of 4 network centrality measures – *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality* and *Eigenvector Centrality* (see Appendix I for details on centrality measures). Detailed definitions of other variables are in Appendix II. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Internet Appendix Table 1. (Continued)
Learning from Colleagues' Ex-Post Forecast Errors

	(1)	(2)	(3)
	OLS: Analyst Revision	OLS: Analyst Revision	OLS: Analyst Revision
Colleague Optimism [A]	-0.287*** (0.109)	-0.276** (0.109)	-0.155 (0.096)
Global Optimism [B]			-3.362** (1.351)
[A] x Analyst Centrality		-0.205* (0.109)	-0.191* (0.108)
[B] x Analyst Centrality			-0.072 (0.615)
Analyst Centrality	-0.319*** (0.111)	-0.249** (0.118)	-0.251 (0.246)
Number of Colleague Forecasts	-0.002 (0.003)	-0.002 (0.003)	-0.001 (0.002)
Number of Global Forecasts			-0.000 (0.000)
General Experience	0.283*** (0.063)	0.283*** (0.063)	0.293*** (0.064)
Firm Experience	-0.159*** (0.035)	-0.159*** (0.035)	-0.156*** (0.035)
Firm Breadth	-0.002 (0.006)	-0.002 (0.006)	-0.001 (0.006)
Industry Breadth	-0.056 (0.089)	-0.060 (0.089)	-0.051 (0.088)
Brokerage Size	-0.103 (0.096)	-0.102 (0.096)	-0.133 (0.092)
Analyst Coverage	0.087*** (0.009)	0.087*** (0.009)	0.088*** (0.009)
Abs. CAR [-5, -2]	-0.337*** (0.022)	-0.337*** (0.022)	-0.338*** (0.022)
Observations	698,697	698,697	698,697
R-squared	0.010	0.010	0.011
Week cluster	Yes	Yes	Yes
Firm cluster	Yes	Yes	Yes

Internet Appendix Table 2.
Diffusion of Revision Activity within Brokerage Networks

This table presents results from negative binomial regressions. The dependent variable is *Revision Count*. For each analyst-brokerage pair, we find her earliest and latest forecast announcement dates. In between these two dates, we divide the analyst's tenure at the brokerage by months. We exclude the first and last months of her tenure from our sample to avoid truncation issues. In each month, *Revision Count* is the total number of forecast revisions made by the analyst. The key independent variable is *Closeness-weighted Revision Activity*. For a given analyst, we find the shortest-path length between her and her brokerage colleagues. For example, if analyst A is directly connected to analyst B, and analyst C is in turn connected to analyst B, the shortest-path length between analysts A and C is two. Using the reciprocals of shortest-path lengths as weights, *Closeness-weighted Revision Activity* of an analyst is the weighted sum of her colleague's *Revision Count* in the previous month. In Column (2), we adopt a normalized measure of *Closeness-weighted Revision Activity* where the *Revision Count* in each of the analyst's distance tier is scaled by the number of colleagues. In Column (3), *Distance1 Revision Activity* of an analyst is the total *Revision Count* of her colleagues who are separated by a shortest-path length of one, normalized by the number of such colleagues. In Column (4), we define *Distance2 Revision Activity* analogously. Estimated coefficients are inflated by 100 to facilitate presentation. Detailed definitions of other variables are in Appendix II. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Internet Appendix Table 2. (Continued)

	(1) N-BINOM: Revision Count	(2) N-BINOM: Revision Count	(3) N-BINOM: Revision Count	(4) N-BINOM: Revision Count
Closeness-weighted Revision Activity [A]	0.014*** (0.003)			
Size-Normalized [A]		1.865*** (0.076)		
Distance1 Revision Activity [B]			3.235*** (0.116)	
Distance2 Revision Activity [C]				1.164*** (0.113)
General Experience	1.489*** (0.170)	1.042*** (0.169)	1.092*** (0.169)	1.339*** (0.171)
Industry Breadth	-12.582*** (0.169)	-11.897*** (0.167)	-12.274*** (0.166)	-12.348*** (0.167)
Brokerage Size	0.076*** (0.004)	0.090*** (0.003)	0.094*** (0.003)	0.091*** (0.003)
Global Revision Activity	24.140*** (0.170)	22.783*** (0.172)	22.825*** (0.169)	23.803*** (0.170)
Calendar month indicators	Yes	Yes	Yes	Yes
Exposure variable: Firm Breadth (coefficient constrained to 1)				
Welch-Satterthwaite t-test between estimated coefficients [B] and [C]:			t-stat: 12.79 df: 956,311	
Observations	478,485	478,485	478,485	478,485

Internet Appendix Table 3.
Do Central or Peripheral Analysts Drive the Findings?

This table presents results from OLS regressions. The dependent variable in Column (1) is *Normalized Forecast Error*, defined as *Raw Forecast Error* scaled by the average forecast error in the firm-year. *Raw Forecast Error* is defined as the absolute difference between an analyst's last firm-year forecast and the actual value. The key independent variables are indicators constructed based on *Analyst Centrality*. *Analyst Centrality* is the standardized PCA-extracted factor score of 4 network centrality measures – *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality* and *Eigenvector Centrality* (see Appendix I for details on centrality measures). In each year, we sort analysts by their *Analyst Centrality* scores into either quartiles or terciles. *Low Centrality* is an indicator that equates to unity if *Analyst Centrality* is less than the 25th (or 33rd) percentile value of the distribution, and equates to zero otherwise. *High Centrality* is an indicator that equates to unity if *Analyst Centrality* is more than the 75th (or 33rd) percentile value of the distribution, and equates to zero otherwise. In Columns (1) through (3), we sort *Analyst Centrality* into quartiles. In Columns (4) through (6), we sort *Analyst Centrality* into terciles. Detailed definitions of other variables are in Appendix II. In all specifications, we include control variables from Column (2) of Table 2. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Internet Appendix Table 3. (Continued)
Central and Peripheral Analysts

	(1) OLS: Normalized Forecast Error	(2) OLS: Normalized Forecast Error	(3) OLS: Normalized Forecast Error	(4) OLS: Normalized Forecast Error	(5) OLS: Normalized Forecast Error	(6) OLS: Normalized Forecast Error
Sort analyst centrality Benchmark category	Quartiles Above p25	Quartiles Below p75	Quartiles p25 to p75	Terciles Above p33	Terciles Below p66	Terciles p33 to p66
Low Centrality	0.011** (0.005)		0.011** (0.005)	0.012*** (0.004)		0.010** (0.004)
High Centrality		-0.016*** (0.006)	-0.016*** (0.006)		-0.019*** (0.005)	-0.017*** (0.005)
Control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	274,673	274,673	274,673	274,673	274,673	274,673
R-squared	0.177	0.177	0.177	0.177	0.177	0.177
Brokerage-Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Brokerage-Year cluster	Yes	Yes	Yes	Yes	Yes	Yes
Analyst-Firm cluster	Yes	Yes	Yes	Yes	Yes	Yes

Internet Appendix Table 4.
Networks Based on 2-digit SIC and 4-digit GICS Classification Schemes

This table presents results from OLS regressions. The dependent variable in Columns (1) and (2) is *Normalized Forecast Error*, defined as *Raw Forecast Error* scaled by the average forecast error in the firm-year. *Raw Forecast Error* is defined as the absolute difference between an analyst's last firm-year forecast and the actual value. The dependent variable in Columns (3) and (4) is the 3-day market-adjusted cumulative abnormal returns, centered on the forecast revision date and unitized in percentage points. The key independent variable is *Analyst Centrality (SIC-2d)* and *Analyst Centrality (G-Group)*. *Analyst Centrality* is the standardized PCA-extracted factor score of four network centrality measures – *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality* and *Eigenvector Centrality* (see Appendix I for details on centrality measures). Detailed definitions of other variables are in Appendix II. In Column (1), we include control variables from Column (2) of Table 2. In Column (2), we include control variables from Column (3) of Table 3. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

	(1) OLS: Normalized Forecast Error	(2) OLS: Abs. CAR [-1, +1]	(1) OLS: Normalized Forecast Error	(2) OLS: Abs. CAR [-1, +1]
Exclude 8-K events	-	Yes	-	Yes
Exclude earnings announcements	-	Yes	-	Yes
Analyst Centrality (SIC-2d)	-0.014*** (0.003)	0.227*** (0.019)		
Analyst Centrality (G-GROUP)			-0.011*** (0.003)	0.197*** (0.017)
Control variables	Yes	Yes	Yes	Yes
Observations	267,353	612,422	267,353	612,422
R-squared	0.177	0.131	0.177	0.130
Brokerage-Year fixed effects	Yes	No	Yes	No
Brokerage-Year cluster	Yes	No	Yes	No
Analyst-Firm cluster	Yes	No	Yes	No
Week cluster	No	Yes	No	Yes
Firm cluster	No	Yes	No	Yes

Commonality in analyst coverage and information diffusion in networks

Kenny PHUA[†]

Nanyang Business School, Nanyang Technological University

Abstract

I map an information-based network of the equity market by linking stocks that are jointly covered by sell-side equity analysts. I exploit a topological feature of networks—the friendship paradox—and show that complementarities in information production can generate information diffusion. Particularly, stock returns of firms are led by those of network neighbors. This effect is robust to adjustments for industry effects, and is not explained by firm characteristics. There are no return reversals in the longer term.

JEL classification:

G10, G12, G14, G17, L22

Keywords:

friendship paradox, information spillover, information production, networks

[†] Nanyang Technological University, Nanyang Business School, 50 Nanyang Avenue, Singapore 639798. E-mail: jphua005@e.ntu.edu.sg. I am grateful to Chishen Wei (my adviser), Jennifer Conrad, Steve Dimmock, Thuc Truc Do, Harrison Hong, Chuan-yang Hwang, Roger Loh, Angie Low, Rui Shen, Mathieu Taschereau-Dumouchel, Michael Weber, Russ Wermers, and participants at the Asian FA Meeting (Seoul), CIRANO-Walton Workshop (Montreal), FMA Doctoral Student Consortium (Boston), Nanyang Technological University Finance Conference, and VICIF (Hanoi) for constructive comments. I received research support from Nanyang Technological University. All errors are solely mine.

A large body of literature studies the transmission of firm information to markets, and how investors impound this information into prices. While traditional finance theories assume that information diffusion occurs instantaneously in markets, much empirical evidence suggests that stock prices can respond slowly to information. Particularly, some studies find that firms' stock prices respond slowly to price innovations of other firms (e.g., Lo and MacKinlay, 1999; Hou, 2007). Although this lead-lag effect can result from nonsynchronous trading or time-varying expected returns, several studies¹ show that these two mechanisms can only partially explain the lead-lag patterns.

In this paper, I focus on another group of explanations—some firms react more quickly to common information than others do.² Specifically, I examine the role of sell-side equity analysts in information diffusion between stocks. As information intermediaries between firms and investors, analysts not only disseminate and interpret major information releases³, but also produce information that is value-relevant across covered firms (Degeorge et al., 2013; Muslu et al., 2014; Gomes et al., 2016). Thus, the dynamics of analysts' information production can potentially affect how investors impound firm information into stock prices.

There are complementarities in analysts' information production because analysts have incentives to cover stocks with information commonalities (Veldkamp, 2006). For example, suppose Apple is jointly covered with Intel and Qualcomm. As more information is produced on Intel and Qualcomm, there are also more information spillovers unto Apple. These spillovers in turn facilitate information production on Apple. As a first step to examine these spillover dynamics, I map the information network of the market by linking stocks that are jointly covered by analysts. Since analysts process news from one firm into relevant information for other firms, the network linkages serve as a conduit for information diffusion.

To fix ideas, I adapt the theoretical framework of Jackson (2016) to analysts' information production characterized by complementarities. The model shows that equilibrium information production on a stock increases with its number of directly linked neighbors (i.e., its network degree). Next, I exploit a topological feature of

¹ Lo and MacKinlay, 1990; Mech, 1993; McQueen et al., 1996; Chordia and Swaminathan, 2000.

² E.g., Lo and MacKinlay, 1999; Brennan et al., 1993; Badrinath et al., 1995; Hou, 2007.

³ E.g., Brennan et al., 1993; Hong et al., 2000; Piotroski and Roulstone, 2004; Asquith et al., 2005.

networks—the friendship paradox—that states that the average neighbor has *strictly higher* degree than the typical firms do. Thus, it follows from the model that *strictly more* information is produced on neighbors than on firms, even on average. Following this, when a firm and its neighbor receive a common shock, investors should be able to impound news into the neighbor’s stock price more quickly. Therefore, I empirically examine whether firms’ returns are led by their neighbors’ returns.

My findings show that gradual information diffusion occurs in networks of stocks linked by analyst coverage. A dollar-neutral portfolio strategy that buys firms in the top quintile of lagged neighbors’ returns, and shorts firms in the corresponding bottom quintile yields 9.4 basis points (bps) per week. I find that the lead-lag effect is driven by diffusion of bad news—more than three quarters of the long-short differential comes from the short side. The profitability of the portfolio strategy is not driven by exposure to risk factors (Fama and French, 2015), and survives controls for a suite of firm characteristics. Additionally, my results continue to hold after controlling for the role of price momentum (Carhart, 1997). Notably, I find that even the smallest neighbors have information leadership over firms.⁴ This suggests that my results are not necessarily driven by the big-leads-small effect (e.g., Lo and MacKinlay, 1999; Hou, 2007). Results from vector autoregressions also support the information diffusion hypothesis.

Since analysts tend to be specialists in one or few industries, I examine the impact of industry effects on my findings. In supplementary tests I find that the lead-lag patterns hold when I 1) constrain firms to be matched with neighbors from different industries, 2) sort on neighbors’ past returns within the firms’ industries, or 3) use industry-adjusted firm returns to measure portfolio performance. From these tests, I estimate that industry effects explain only up to 40% of my baseline strategy profitability. Therefore, my proposed mechanism of information diffusion is compatible with, albeit distinct from, industry effects such as industry momentum (Moskowitz and Grinblatt, 1999).

⁴ Cohen and Lou (2012) and Scherbina and Schlusche (2015a; 2015b) find that returns of small firms can lead those of large firms.

Since the construction of network linkages are based on commonalities in analyst coverage, they likely support the diffusion of coverage-specific information—information that emphasizes commonalities among stocks in an analyst’s coverage (Muslu et al., 2014).⁵ To support this notion, I perform a falsification test by matching each firm with another randomly selected stock in the network, instead of its neighbor. These pseudo matches may capture diffusion of broad-based information but exclude transmission of coverage-specific information. These pseudo matches do not yield lead-lag patterns, thereby suggesting that meaningful information structures underlie the observed networks.

Additional evidence confirms that the lead-lag effect is driven by imbalances in information production between firms and their neighbors. I use analysts’ revision intensity for a stock as a proxy for information production (Jacob et al., 1999). When the difference in analysts’ revision intensities is greater (neighbors minus firms), the profitability of my portfolio strategy is also higher. Separately, if firms and neighbors share information commonalities, they may be exposed to common economic shocks and make correlated real decisions (Engelberg et al., 2017). To this end, I find a lead-lag pattern in acquisitiveness between firms and neighbors.

Investors appear to impound neighbors’ information into firms’ stock prices relatively quickly—the strategy ceases to be profitable beyond the third week after portfolio formation. More importantly, I do not find return reversals in the longer term. Per the reasoning in Cohen and Lou (2012), this suggests that my findings are driven by diffusion of fundamentally relevant information, and are unlikely explained by investors’ overreaction or underreaction.

The primary contribution of this paper is to map the information networks of the equity market, and show that lead-lag dynamics can emerge in these networks. In this vein, this paper links to the literature on cross-firm information flows. Menzly and Ozbas (2010) find that information flows between customer and supplier industries, and Hong et al. (2007) find that the returns of certain industries lead broader market returns. Additionally, Cohen and Frazzini (2008) find that information diffuses slowly

⁵ Coverage-specific information (Muslu et al., 2014) lies on the spectrum whose opposite extremes are marked by stock-specific and broad (marketwide or industry) information.

through the supply chain while Cohen and Lou (2012) show that single-segment firms lead multi-industry conglomerates. Scherbina and Schlusche (2015) also document lead-lag effects among stocks co-mentioned in the news media. These papers study information diffusion in the context of customer-supplier relationships, organizational structures, industry linkages, and media linkages. Instead, my paper shows that complementarities in analysts' information production can generate information diffusion. Furthermore, the network approach in this paper allows me to accommodate higher-order dependencies across stocks.⁶

This article also relates to recent works that apply network-theoretic techniques to explore complex interfirm dynamics. Particularly, Ahern (2013) constructs networks based on input-output trade flows and finds that stocks in central industries are riskier because they are more exposed to sectoral shocks that transmit through intersectoral trade. Separately, Ahern and Harford (2014) find that merger waves propagate across the network via customer-supplier linkages. While these studies emphasize the role of real economic channels as a conduit for interfirm interactions, my paper focuses on information diffusion between firms along 'soft' information-based network linkages. I also provide a novel application of the friendship paradox in financial networks to examine information production with complementarities.

Finally, my article also contributes to the literature on the informational roles of analyst research. Consistent with the theoretical model in Veldkamp (2006), Hameed et al. (2015) show that prices of stocks with greater analyst following are more informative and that investors use these prices to infer the values of less-followed stocks in the same industries. While their proposed mechanism is consistent with my findings, my network approach allows for information spillovers that transcend industry boundaries. Crucially, I also show that the lead-lag effect persists even when I control for industry effects. Hong et al. (2000) shows that analysts are instrumental in propagating firm-specific information. I extend their analysis by considering the diffusion of information that is not strictly firm-specific. Relatedly, Muslu et al. (2014)

⁶ In the presence of complementarities, information production on Apple is a function of that on its neighbors—Intel and Qualcomm. Information production on Intel and Qualcomm is, in turn, dependent on their own set of neighbors that not directly linked to Apple, so on and so forth. A network approach accommodates such higher-order dependencies between stocks. A similar point is emphasized in Ahern and Harford (2014).

find that analysts produce coverage-specific information that are value-relevant for stocks under shared analyst coverage. I extend this idea by showing that, in the presence of complementarities, information production may be preferentially tilted towards certain stocks in a complex interfirm network.

1. The friendship paradox and information production with complementarities

I adapt the theoretical framework of Jackson (2016) to information production characterized by complementarities. I begin by showing how the friendship paradox emerges in a network of stocks linked by commonalities in analyst coverage. Thereafter, I analyze the impact of the friendship paradox on analysts' information production in the presence of complementarities.

1.1. The friendship paradox in a random network

There is a finite set $N = \{1, 2, \dots, n\}$ of stocks, with generic indices i, j as nodes of a random network. If any two nodes are linked in the network, it follows that they are two different stocks jointly covered by at least one analyst. Stock $i \in N$ has $d_i \in \{1, 2, \dots\}$ linkages in the network. Equivalently, d_i is the number of direct neighbors that stock i has in the network.

For some firm i , let $P_i(d)$ be the degree distribution of the population of stock i 's potential neighbors under a random network where links between nodes are randomly generated. To be clear, $P_i(d)$ describes the degree distribution of any stock i 's neighbors, not conditioning on i being connected to the neighbor. I also assume that the marginal distribution $P_i(d)$ is same for every neighbor of i . Additionally, I let $E_i[\cdot]$ and $\text{Var}_i[\cdot]$ denote the expectation and variance of $P_i(d)$.

Now, conditional on stock i being linked to a neighbor (denoted by the tilde accent), the probability that the neighbor has degree d is given by the following equation (1).

$$\tilde{P}_i(d) = \frac{d \cdot P_i(d)}{\sum_d d \cdot P_i(d)} = \frac{d}{E_i[d]} \cdot P_i(d) \quad (1)$$

I reiterate the perspective here. An interfirm network has formed and stock i is linked to some neighbors. Conditional on these linkages, if I examine any one of stock i 's neighbors, $\tilde{P}_i(d)$ describes the distribution of degrees on that particular neighbor. Observe that $\tilde{P}_i(d)$ is increasing in d ; this relation drives the friendship paradox. Intuitively, stocks that have higher degrees must also be neighbors more frequently in the network. Appendix I contains a short working example to elucidate this intuition. This working example embodies the core of the friendship paradox – even on average, the set of neighbors linked to a stock is weighted not by the neighbors' proportions in the population, but by their network 'popularity' (i.e. degree).

To complete the formalization of the friendship paradox, I take expectations $\tilde{E}_i[\cdot]$ over $\tilde{P}_i$. In other words, conditional on linkages to neighbors, $\tilde{E}_i[\cdot]$ gives the expected number of linkages that each of stock i 's neighbors has. Together with equation (1), I express $\tilde{E}_i[d]$ in equation (2) below.

$$\tilde{E}_i[d] = \sum_d d \cdot \tilde{P}_i(d) = \sum_d d \cdot \frac{d}{E_i[d]} \cdot P_i(d) = \frac{E_i[d^2]}{E_i[d]} = E_i[d] + \frac{\text{Var}_i[d]}{E_i[d]} \quad (2)$$

Equation (2) states that the expected degree of a stock's neighbor is the unconditional population average $E_i[d]$ plus a bias factor $\frac{\text{Var}_i[d]}{E_i[d]}$. As long as there is some variation in the degrees of neighbors, such that $\text{Var}_i[d] > 0$, the friendship paradox holds.⁷ All stocks with degrees no higher than the population average are expected to have strictly less linkages than their neighbors do. Only stocks with substantially higher degrees than the population average (by at least the amount of the bias factor) are expected to have the same number of linkages as their neighbors do. Furthermore, if idiosyncrasies are not salient such that $\tilde{E}_i[\cdot]$ is common across stocks, I can drop the subscripts and write the expression as the following equation (3).

$$E[d] < \tilde{E}[d] = E[d] + \frac{\text{Var}[d]}{E[d]} \quad (3)$$

⁷ For this condition to be violated, every node in a network has to be linked to every other node in the network such that every node has the same number of connections. Such network structures are rare in reality. This condition is not violated in my interfirm networks.

1.2. Information production of sell-side equity analysts

In this section, I examine how complementarities generate imbalances in analysts' information production between a stock and its neighbors. Since stocks that share joint coverage are likely to share information commonalities (Veldkamp, 2006), information produced for a stock can be relevant for other stocks (Muslu et al., 2014), and vice versa. Such complementarities imply that information spillovers from a stock's neighbors may aid the valuation of the stock. Following this, as more information is produced on a stock's neighbors, information production on the stock itself should also increase.

To fix ideas, I adapt the theoretical framework of Jackson (2016) to the information production decisions of analysts facing an interfirm network per the setup in Section 1.1. An analyst who covers stock i has type $\theta_i \in \mathbb{R}_+$, faces a quadratic cost of information production c , and has complementarities captured by a . The analyst chooses to produce $x_i \in \mathbb{R}_+$ units of information for stock i , and gets utility of the following form in equation (4).

$$\theta_i x_i + a x_i \sum_{j \in \mathbf{N}_i} x_j - \frac{c x_i^2}{2} \quad (4)$$

The parameter θ_i captures any general preference of the analyst to produce information for stock i . The complementarities arising from joint information production in stock i and its neighbors j are controlled by $a > 0$. Intuitively, since a loads on the interaction between x_i and x_j , a positive a implies that the analyst choose a higher information production on stock i in response to more information produced on its neighbors. The quadratic cost $c > 0$ ensures that feedback effects between i and j do not drive optimal x_i to infinite. Stock i is connected to a set of neighbors $\mathbf{N}_i$. All analysts choose their quantities of information production for each stock under their coverage simultaneously. Knowing her own type and the degrees of the stock under her coverage, and expecting over the degrees of the stock's neighbors, the analyst maximizes the following equation (5).

$$\theta_i x_i + a x_i \cdot \tilde{\mathbb{E}} \left[\sum_{j \in \mathbb{N}_i} x_j \right] - \frac{c x_i^2}{2} \quad (5)$$

If the neighbors of stock i (i.e. any neighbor j) are drawn from the same distribution, I can rewrite the maximization problem as equation (6) below.

$$\theta_i x_i + a x_i \cdot d_i \tilde{\mathbb{E}}[x_j] - \frac{c x_i^2}{2} \quad (6)$$

1.3. Information production in equilibrium

To arrive at a Bayesian equilibrium, consider the setting in Section 1.1 where all stocks in the network face the same P_i . Let $\tilde{\mathbb{E}}[\cdot]$ denote $\tilde{\mathbb{E}}[\cdot]_j$ for any stock i . Analysts are allowed to differ in their types θ_i and in the degrees d_i of stocks under their coverage. Additionally, θ_i and d_i are assumed to be uncorrelated so that $\tilde{\mathbb{E}}[\theta_j] = \mathbb{E}[\theta_j]$. Maximizing x_i as a function of θ_i and d_i , the first order condition yields equation (7) below.

$$x_i = \frac{\theta_i}{c} + \frac{a \cdot d_i \mathbb{E}[x_j]}{c} \quad (7)$$

To obtain the expected information production of neighbors, I take expectations $\tilde{\mathbb{E}}$ of both sides of equation (7) yields the following.

$$\tilde{\mathbb{E}}[x] = \frac{\mathbb{E}[\theta]}{c} + \frac{a \cdot \tilde{\mathbb{E}}[d] \cdot \tilde{\mathbb{E}}[x]}{c} \quad (8)$$

Rearrangement of terms gives the below solution.

$$\tilde{\mathbb{E}}[x] = \frac{\mathbb{E}[\theta]}{c - a \cdot \tilde{\mathbb{E}}[d]}$$

Substituting the solution back into equation (7), the equilibrium quantity of information produced for stock i under friendship paradox is characterized by equation (9) below.

$$x_i^{\text{friendship paradox}} = \frac{\theta_i}{c} + \frac{a \cdot d_i \mathbb{E}[\theta]}{c(c - a \cdot \tilde{\mathbb{E}}[d])} \quad (9)$$

Equation (9) shows that increases in the complementarities of information production a , increases in the firm's degree d_i , decreases in the corresponding costs c , and increases in the variance of firms' degrees in the population (hence increasing $\tilde{E}[d]$), all heighten the quantity of information produced by analysts for stocks. The main effect is generated through d_i . As made clear in equation (3), neighbors have strictly higher degrees than the average stock in the network. Therefore, more information produced is produced for neighbors than for stocks in equilibrium. A positive feedback effect also operates through $\tilde{E}[d]$. As neighbors of a stock have higher degrees and thus higher information production, analysts respond to the greater information spillovers by producing more information on the stock too. This effect subsequently feeds back into neighbors, thus reinforcing the main effect.

To understand the impact of the friendship paradox on analysts' information production, we compare the equilibrium outcome of equation (9) to that in a benchmark setting without network effects. In the benchmark setting, stocks are randomly matched with the population for their interactions. Put differently, I ignore the information commonalities between stocks and their neighbors and form stock-pairs randomly from the population. Under the benchmark setting, the maximization problem appears similar to equation (7) save for a crucial difference – analysts take expectations over the whole population (E) rather than $\tilde{E}$ which conditions on the stocks' neighbors in the network. It is straightforward to show that the equilibrium choice of information production for stock i under the benchmark is given by equation (10) below, differing from equation (9) only in the denominator.

$$x_i^{benchmark} = \frac{\theta_i}{c} + \frac{a \cdot d_i E[\theta]}{c(c - a \cdot E[d])} \quad (10)$$

Given $a > 0, c > 0$ and the relation between $E[d]$ and $\tilde{E}[d]$ in equation (3), a comparison yields the following strict inequality.

$$x_i^{friendship\ paradox} > x_i^{benchmark}$$

This suggests that the equilibrium quantities of information produced on stocks are strictly higher under network effects, compared to a setting where information production on stocks are benchmarked against the population average. I examine this

contrast empirically by randomly rewiring the interfirm linkages in the network in later sections.

2. Data and methodology

In this section, I begin by outlining the multiple sources of data used in this article. Thereafter, I discuss how interfirm networks formed by commonalities in analyst coverage are constructed. Finally, I describe the portfolio strategies used to exploit imbalances in information production between firms and their neighbors.

2.1. Data

Data on analysts' 1-year-ahead earnings estimates is taken from the Institutional Brokers Estimate System (I/B/E/S) database. The universe of analysts' earnings estimates spans 11,968 unique firms over calendar years 1995 to 2014. Stock returns data is drawn from the Center for Research in Securities Prices (CRSP) database, and financial variables are obtained from Compustat. Following Jegadeesh and Titman (2001), stocks with share prices lower than five dollars at portfolio formations are omitted to reduce the likelihood that the results are driven by small, illiquid stocks or by bid-ask bounce. I do not require firms to be followed by a minimum number of analysts. My baseline sample is formed at the intersection of I/B/E/S, CRSP, and Compustat. For tests pertaining to the propagation of real effects, I draw data on corporate acquisitions from SDC Platinum. Following Ahern and Harford (2014), acquisitions must meet the following criteria: (1) both the target and the acquirer are U.S. firms; (2) the acquirer buys 20% or more of the target's shares; (3) the acquirer owns at least 51% of the target's shares after the deal; (4) the acquisition is completed; and (5) the transaction value is at least \$1 million.

2.2. Interfirm networks

My construction of interfirm networks centers on the notion that, due to costly information production (e.g. Veldkamp, 2006), analyst coverage decisions are

endogenous to the information commonalities between stocks. Following this, two stocks are linked in the network if they are jointly covered by at least one analyst in the calendar year. More specifically, the analyst must have issued at least one earnings forecast for each stock of the pair in the calendar year. Interfirm linkages are binary so that stocks are either linked or unlinked, and the weights allocated to linkages are invariant with the extent of shared analyst coverage.⁸

[Insert Figure 1]

Figure 1 contains a visual plot of the information network formed on shared analyst coverage between firms in 2005. Stocks are represented by circles and linkages between them are portrayed as lines. The plot is presented in the OpenOrd layout (Martin et al., 2011) which highlights stock clusters in the network to produce a cleaner plot of a large network. Included in Figure 1 are two inset magnifications to focus on local topological features of the network. As an illustration, the magnification at the heart of the network shows two clusters of stocks (colored red and yellow) that have many long-range linkages, direct or indirect, with other parts of the network. In contrast, the magnification at the periphery of the network shows a small cluster of four stocks (colored blue) that have linkages to one another but not to others outside this cluster. The magnification also picks up some isolated stocks without any neighbors; these firms are omitted from my networks because they play no part in the portfolio strategy.⁹ However, these stocks are in the extreme minority and constitute considerably less than 1% of the sample across all years, on average.

Using data on analysts' earnings estimates from January to December, I form interfirm networks at December end of every calendar year from 1995 to 2014. On average across all years, these interfirm networks contain 4,301 unique stocks, and each stock is connected to approximately 69 other stocks (i.e., average degree of 69). Network sizes peak in the late 1990s with more than 5,000 stocks but has generally declined to 3,613 in 2014. There is a U-shaped evolution of stocks' network degrees over the years, with $E[d]$ dipping to approximately 55 in the early 2000s and rising to

⁸ In other words, all interfirm linkages are equal in weight. I do not assign a high weight to a linkage even if there are many analysts who jointly cover the linked firm-pair.

⁹ For a firm to be isolated in the network, it must be covered by analysts who cover no other firms jointly. This is extremely rare in the I/B/E/S dataset in which the median analyst covers 16 firms.

about 80 on both temporal ends of the sample. I find that heterogeneity of stocks' degrees also declines over time as $\frac{\text{Var}[d]}{\text{E}[d]}$ decreases from about 42.3 in the first half of the sample period to about 26.6 in the latter half. While the friendship paradox is substantially stronger in the earlier years, my results are robust to subperiod analysis.

[Insert Table 1]

In Panel A of Table 1, I provide additional insights of the network structures by focusing on the largest components of the interfirm networks. A component is defined as the subset of the network where any stock is reachable by any other stock along the network paths. On average, more than 99.4% of stocks are connected in the largest component, and this ratio remains stable over the sample period. In this component, the typical stock is only about three linkages away from any other stock. In fact, the maximum shortest path length between any two stocks, also known as the component diameter, is 8.2 on average. The interfirm networks have a mean transitivity score of 0.39, indicating that 39% of the time two stocks linked to the same neighbor are also linked to each other. As a benchmark for this measure, I generate 100 random networks that preserve the total number of stocks and linkages in the interfirm networks¹⁰, and compute their transitivity scores. The average transitivity score from the random simulations is about 24 times smaller than the value in my observed networks. This supports the notion that shared analyst coverage between stocks does not emerge randomly but may be outcomes of strategic considerations. Taken together, the short separation lengths between stocks and high transivities point to a small-worlds (Watts and Strogatz, 1998) property of the interfirm networks. I also report stock-level statistics of various financial characteristics in Panel B. Apart from the period surrounding the 2008 financial crisis, there are no obvious trends in book-to-market ratios, past performance, idiosyncratic volatility, and Amihud (2002) illiquidity. Consistent with intuition, market capitalizations and past performance dipped during the crisis while book-to-market ratios, idiosyncratic volatility and illiquidity spiked concurrently.

[Insert Figure 2]

¹⁰ This procedure is commonly known as a degree-preserving rewiring of the original network.

Figure 2 reports the neighbor-to-firm ratios of financial characteristics across the sample period. Where the ratio of a particular characteristic is greater than unity (above the red line), neighbors have higher values of that characteristic than firms do. Across all years, neighbors have larger market capitalizations and idiosyncratic volatilities relative to firms. With the exceptions of crisis years around 2000 and 2008, relative idiosyncratic volatilities remain stable across years. However, the size dominance of neighbors over firms has declined steadily over time, falling from a peak ratio of 1.93 in 1996 to 1.40 in 2014. On average, neighbors also have poorer past performance, lower book-to-market ratios, and lower Amihud (2002) illiquidity. Interestingly, leading up to the 2008 financial crisis, neighbors have stronger past performance than firms do. Figure 2 also plots the proportions of firm-neighbor pairs that belong to identical Fama-French 49 or two-digit SIC industries in bars. Approximately 60% of firm-neighbor pairs do not reside in the same industries.

[Insert Figure 3]

Finally, I present the industry distribution of the networks in Figure 3. To facilitate presentation, I adopt the SIC divisions classification scheme which assigns stocks into 10 broad economic divisions.¹¹ In each calendar year, I compute the proportion of every SIC division in the interfirm network. Thereafter, I average these proportions across all years for every SIC division and report them in Figure 3. The majority of stocks in my sample (39.5%) belong to the manufacturing division. The finance, insurance, and real estate division and services division comprise, respectively, 20.2% and 16.3% of the sample. The remaining sample stocks are distributed across the other seven SIC divisions, with each of these SIC divisions not having a weight of more than 10%.

2.3. The portfolio strategy

Section 1 shows that, due to the friendship paradox in networks, neighbors have strictly higher degree than firms do, on average. Under information production with complementarities, this implies that neighbors receive relatively greater information spillovers, and hence more information is produced on neighbors. Therefore, when a

¹¹ There are, in fact, 11 SIC divisions under the classification scheme. However, no firms in my sample belong to the SIC division – public administration.

firm and its neighbor receive a common shock, investors should be able to impound news into the neighbor’s stock price more quickly. Following this, I design a portfolio strategy to examine whether the returns of firms are led by the returns of neighbors.

The portfolio strategy operates at a weekly frequency and uses only information known at the portfolio formation date. At the beginning of portfolio formation week w in year y ¹², I match each firm in the network with one of its randomly selected neighbor. To be clear, for a network of size N , I will obtain N firm-neighbor pairs. Thereafter, I track the stock returns of the neighbors over week w , and form quintile portfolios of firms sorted on their neighbors’ returns. To mitigate microstructure concerns, I skip a week after the formation week and hold the portfolios for one week (i.e., in week $w+2$). I calculate weekly portfolio returns as the equal-weighted average returns of constituent stocks. Finally, I execute a zero-cost portfolio strategy that is long in the quintile portfolio of firms whose neighbors have the highest lagged returns, and short in the corresponding bottom quintile portfolio of firms.

My portfolio strategy is related to cross-disciplinary applications of the friendship paradox in networks. For example, Cohen et al. (2003) show that instead of immunizing randomly-selected individuals in the population, targeting the randomly-chosen acquaintances of randomly-selected individuals is more effective. In the context of a flu outbreak at Harvard College, Christakis and Fowler (2010) find that the infection rate of randomly-chosen students lags that of the students’ reported friends by about 13.9 days. All told, there are some similarities between complex information flows among firms and epidemiological propagation among individuals. First, information shocks can spread through the network in an omnidirectional manner as viruses do among humans, such that it becomes difficult to determine the course of propagation *ex ante*. Second, like viruses, information shocks can spread via a parallel distributive mechanism (Borgatti, 2005) such that any single affected stock can transmit the shock to several others simultaneously.

3. Empirical results

¹² I operate on the network that is formed at the end of calendar year $y-1$ so I know the network structure at the point of portfolio formation.

In this section, I begin by showing baseline results from portfolio sorts on the lagged returns of network neighbors. Subsequently, I perform independent double-sorts on neighbor returns and various firm characteristics to ensure that my results are not driven by effects documented in the extant literature. Thereafter, I estimate multifactor regressions to test if the return patterns from my portfolio strategy are captured by the five factors of Fama and French (2015) and the momentum factor. I also estimate Fama and Macbeth (1973) cross-sectional regressions to show that neighbors' lagged returns can predict firms' returns. Following this, I show evidence that accords with the theoretical predictions in Section 1 – imbalances in information production between firms and neighbors are responsible for the gradual information diffusion in the interfirm networks. I also estimate vector autoregressions to show that information primarily diffuses from neighbors to firms. Finally, I show that corporate acquisitiveness can propagate along the interfirm linkages.

Since the portfolio strategy involves randomization, I perform all tests over 100 rounds of simulations and report average statistics in the tables.¹³ To perform statistical inference, I adopt the Fisher's method (Fisher, 1925) on the distribution of p -values across all rounds of simulations. Under the global null hypothesis, $-2 \sum_i^k \ln(p_i)$ should fit a chi-squared distribution with $2k$ degrees of freedom where k is the number of simulation rounds and p_i is the p -value from the i^{th} round of simulation.

3.1. Portfolio sorts on stock returns of network neighbors

If imbalances in information production between neighbors and firms cause information to diffuse gradually from neighbors to firms, neighbor returns should positively predict future returns of firms. I test this primarily through a dollar-neutral portfolio strategy that buys firms with good neighbors' past performance and sells firms with poor neighbors' past performance, as in Section 2.3.

[Insert Table 2]

¹³ My algorithm uses a pseudorandom generator to match each firm with one of its randomly selected neighbors. The pseudorandom generator guarantees the replicability of my results given the same initial conditions.

In the baseline specification in Panel A, I find that the long-short portfolio strategy yields an average of 9.4 bps per week (4.89% per annum). As expected of the friendship paradox, the neighbor-to-firm degree ratios are consistently well above unity (average=1.266), indicating that neighbors have higher degrees than firms do. Market capitalizations and four-week past performance of firms are similar across the quintile portfolios, suggesting that sorts on neighbors' returns are orthogonal to sorts on firms' sizes and past performance. However, I find a slightly hump-shaped pattern in firms' b/m ratios across the quintile portfolios. Of the high-minus-low spread, about 75% comes from the short side of the trade, that is, the difference between medium neighbor-returns stocks and low neighbor-returns stocks. This is consistent with the findings in Hou (2007) that lead-lag effects are more pronounced when bad news arrives (e.g., short sale constraints inhibit the incorporation of negative information into prices (Diamond and Verrecchia, 1987)). In Panel B, I rule out distortions due to the January effect but excluding portfolio months in December and January. Here, I find that the results remain both quantitatively and qualitatively similar.

The model in Section 1.3 shows that the imbalances in information production should occur exclusively between firms and their neighbors. Following this, I set up a falsification test by matching firms with pseudo-neighbors—firms are paired randomly with any other stocks in the networks. Essentially, this matching method disregards the information commonalities between a firm and its neighbors. If my baseline findings are not driven by network effects, I expect to find similar results when I match firms with pseudo-neighbors. However, I find no statistically significant evidence of a lead-lag effect between pseudo-neighbors and firms. Furthermore, the magnitude of the long-short return differential in Panel C is more than 90 times smaller than that in Panel A. This contrast suggests that an information-based mechanism drives gradual information diffusion in the network.

3.2. Portfolio double-sorts on neighbors' returns and firm characteristics

In this section, I perform independent portfolio double-sorts to ensure that I am not merely picking up returns patterns due to effects of firms' size, b/m, momentum, illiquidity, idiosyncratic volatility, and neighbors' size. Sizes of firms and neighbors are

defined as their market capitalization at the time of portfolio formation. Book equity (b) is defined as the Compustat book value of stockholders' equity, plus balance sheet deferred taxes and investment tax credit (where available), less the book value of preferred stock. Depending on their availability, redemption, liquidation, or par value (in this order of priority) is used to estimate the book value of preferred stock. To ensure that these variables are known at portfolio formation, I match the annual book equity value for fiscal years ending in calendar year $y-1$ with portfolio formation dates starting in calendar year y . Book equity is then scaled by market capitalization at portfolio formation to form b/m . Momentum is the firms' returns from $w-4$ to $w-1$. Illiquidity is the summation of the daily ratios of return magnitudes to trading volume over $w-4$ to $w-1$. For the illiquidity measure of a firm to be valid, the firm must have no less than 10 valid data points over $w-4$ to $w-1$. Following Campbell et al. (2001), I define idiosyncratic volatility of a firm as the summation of squared differences between daily returns and daily Fama-French 49 industry returns over $w-4$ to $w-1$. All summed values are then scaled to 21 days to ensure comparability (Brav et al. 2010).

[Insert Table 3]

The returns on the resulting 25 portfolios for each independent double-sort are presented in Table 3. Across all double-sorts, the average portfolio contains about 118 firms. I begin by double-sorting on neighbors' returns and firms' sizes in Panel A. The evidence suggests that the lead-lag effect is positive and statistically significant in all but the largest firm size quintile. Among the other four quintiles, the long-short returns are greatest among the smallest stocks (0.142% per week) and drop monotonically to 0.065% per week in the fourth firm size quintile. While the long-short return differential is positive among the largest firms, it is not statistically distinguishable from zero. However, I later show in Section 3.4 that, conditional on imbalances in information production, even the stock returns of the largest firms can be led. Next, the evidence in Panel B indicates that the lead-lag effect is present across b/m quintiles, with returns differentials ranging from 0.048% per week to 0.121% per week. Interestingly, the lead-lag effect is weaker among high-growth firms. A potential explanation is that high-growth firms attract considerable investor attention in themselves so their investors rely less on information produced on neighbors.

Neither is the friendship paradox effect driven by firms' momentum. In Panel C, the portfolio strategy yields statistically significant long-short profits across all quintiles of firms' momentum, ranging from 0.068% per week to 0.103% per week. Panel D shows that the lead-lag effect persists across all quintiles of idiosyncratic volatility. Notably, the long-short return differential among low-volatility firms is one-third the size of that among high-volatility firms. While I explicitly adjust for industry effects in later sections, this return pattern suggests that the lead-lag effect is unlikely driven by industry effects. If the information leadership of neighbors is purely industry-based, I should observe the strongest lead-lag effect among firms whose returns co-move more with industry returns (low idiosyncratic volatility).

Panel E indicates that the lead-lag effect is robust across illiquidity quintiles, except for the lowest one. The long-short return differential among the most liquid firms is positive but statistically insignificant. This pattern is consistent with the mechanism implied in the theoretical model. In the model, information diffusion from neighbors to firms is driven by imbalances in information production that, in turn, causes firms to have weaker informational efficiency on average. Since high firm liquidity is symptomatic of high informational efficiency (Admati and Pfleiderer, 1988; Chordia et al., 2008), much private information is likely already impounded into the stock prices of liquid firms, leaving little room for the information leadership of neighbors. Consistent with this explanation, I find that the lead-lag effect strengthens monotonically with firm illiquidity. Finally, I find that even small neighbors can have information leadership.¹⁴ Panel F shows that long-short returns are positive and statistically significant across all quintiles of neighbor sizes, and range from 0.350% to 0.476% per month. In the smallest quintile of neighbors, the average size of firms is about 21 times that of neighbors. Therefore, my findings are not a rehash of the big-leads-small effect in Lo and MacKinlay (1990).

Overall, I find that return patterns arising from the information leadership of neighbors are not subsumed by effects of various firm characteristics. The lead-lag

¹⁴ Cohen and Lou (2012) and Scherbina and Schlusche (2015a; 2015b) show that small stocks can have information leadership.

effect is stronger among firms that have smaller sizes, less growth opportunities, higher idiosyncratic volatility, and lower liquidity.

3.3. Multifactor regressions

In this section, I test whether the returns from portfolios sorted on neighbors' lagged returns can be explained by multifactor models. Table 4 presents results from the Fama and French (1993) three-factor model and the Fama and French (2015) five-factor model, both which I augment with the Carhart (1997) momentum factor. Reported coefficients are averages across 100 rounds of simulations and average t -statistics are reported in parentheses.

[Insert Table 4]

The estimated alphas indicate that the augmented three-factor model leaves substantial positive unexplained returns in all but the bottom quintile (4.9 bps – 9.2 bps per week). While the estimated alpha in the bottom quintile is negative, it is not statistically significant. The loadings suggest that firms across the quintile portfolios are discriminated neither on size nor on exposure to market risk. Firms in the extreme quintiles exhibit a slight tilt towards growth stock characteristics but behave like value stocks in the others. However, the portfolios' exposures to the value factor are weak with loadings ranging only between -0.051 and 0.171 . I also find that the sort on neighbors' returns produce little variation in past performance. The GRS (Gibbons et al., 1989) test rejects the hypothesis that the four-factor model explains the average returns on the five portfolios sorted on neighbors' lagged returns.

Table 4 also reports the estimates of the momentum-augmented Fama-French five factor model. I find that this model also leaves unexplained returns in the top four quintiles but not in the bottom one. The pattern of loadings on market risk and SMB is similar to the pattern found in the three-factor regressions. Here, I find that stocks in the bottom quintile exhibit a slight tilt towards weak profitability, loser, and aggressive investment behavior. However, the bottom quintile portfolio's loadings on these factors are small, ranging only between -0.052 and -0.127 . Again, the GRS test

rejects the hypothesis that the augmented five-factor model explains the average returns across the sorted portfolios.

On a weekly basis, the long-short portfolio strategy produces statistically significant alphas of 9.7 bps and 9.0 bps in the augmented three- and five-factor models, respectively. All told, the evidence suggests that sorts on lagged neighbors' returns do not produce portfolios that are sharply differentiated along conventional dimensions of factor risk. My results also indicate that the risk-adjusted profitability of my portfolio strategy appears to stem from the long side.

3.4. Relation between returns of firms and neighbors in cross-sections

This section employs Fama and Macbeth (1973) regressions at weekly cross-sections on my sample of stocks to illustrate the relation between firms' returns and lagged neighbors' returns. This analysis also acts as a robustness check since the regressions employ all stocks in my sample (i.e., no breakpoint specification is required), and control for stock characteristics that are known to predict stock returns. For each of 100 round of simulations, I estimate firm-level cross-sectional regressions of firm returns (in week $w+2$) on lagged neighbors' returns (in week w), firms' market capitalization, book-to-market, idiosyncratic volatility, and past performance measured at horizons of one week ($r_{-1,0}$) and 24 weeks ($r_{-24,0}$). Thereafter, I compute the time-series averages of their loadings with Newey-West adjustment for autocorrelations up to four lags.

[Insert Table 5]

Table 5 reports averages of estimates from Fama and Macbeth (1973) cross-sectional regressions over 100 rounds of simulations. Newey-West t -statistics are reported in the parentheses. The univariate specification in Column (1) shows that neighbors' returns predicts future firms' returns. The predictive power of neighbors' returns survives the addition of controls for various firm characteristics in Column (2). Moreover, the estimated coefficient on neighbors' returns remain similarly sized with the inclusion of the controls. While firms' returns are not predicted by short-term momentum, they are positively related to momentum over the prior 24 weeks.

Consistent with prior studies (e.g., Ang et al., 2009), I also find that higher idiosyncratic volatility predicts lower firm returns.

As a robustness check, I perform a subperiod analysis in Columns (3) and (4). I continue to find that firms' returns increase with neighbors' lagged returns in the earlier period (1996-2005) and the later period (2006-2015). Although Panel A of Table 1 shows that the network structures evolve over time such that the friendship paradox is weaker in the later period, I find similarly sized relations between the returns of firms and neighbors across both subperiods.

In untabulated tests, I re-estimate the Fama and Macbeth (1973) regressions but replace neighbors' returns with returns of pseudo-neighbors (see Section 3.1). In these specifications, I find that firms' returns cannot be predicted by pseudo-neighbors' returns in the cross-sections. Again, this supports the notion that the observed network structures capture meaningful information commonalities among stocks and their neighbors.

3.5. Mechanism – imbalances in information production

In this section, I reconcile my empirical findings with the theoretical predictions in my model. The lead-lag effect is generated by imbalances in information production between firms and neighbors. Therefore, the lead-lag effect should be stronger when these imbalances are more pronounced—more information is produced by analysts on neighbors than on firms.

I first operationalize information production on a stock as the frequency of revisions issued for the company per analyst, or *revision rate*. This is motivated by Jacob et al. (1999) who measure analysts' forecasting effort with their revision frequencies. Following this, I measure imbalances in information production within a firm-neighbor pair as the ratio of the neighbor's *revision rate* to the firm's *revision rate*. Higher values of *revision rate ratios* thus indicate that the information production is much higher on neighbors than on firms. To ensure that revision rates are knowable at the time of portfolio formation, the ratio is computed with revision rates of the preceding

calendar year. For analyst a covering firm i and analyst b covering neighbor j in calendar year y , I define revision rate ratio as follows.

$$revision\ rate\ ratio_{i,j,y} = \frac{\sum_{b \in N_j} |revision_{b,j,y-1}|}{\sum_{a \in N_i} |revision_{a,i,y-1}|} \frac{|N_i|}{|N_j|}$$

I find that larger stocks tend to have higher *revision rates* so I include firm size as a sorting dimension to isolate the variation in *revision rate ratios* more cleanly. To ensure that each portfolio is sufficiently populated, I form portfolios that are independently sorted on 1) firm size, 2) revision rate ratios, and 3) lagged neighbors' weekly returns. Each portfolio contains an average of 109 stocks.

[Insert Table 6]

Table 6 reports returns and characteristics of the $3 \times 3 \times 3$ portfolios.¹⁵ In the lowest terciles of *revision rate ratios*, the average revision frequency for neighbors is about 64% that of firms across size terciles. In contrast, neighbors receive 1.82 revisions for every revision issued for firms in the highest terciles of *revision rate ratios*. Among the smallest firms, I find that the lead-lag effect strengthens when imbalances in information production become more severe. I also find a similar returns pattern among medium-sized firms. Consistent with my previous results that the lead-lag effect weakens with firm size, the long-short return differentials for every tercile of *revision rate ratios* are all smaller among medium-sized firms than among the smallest firms (6.2 bps vs. 9.0 bps; 5.5 bps vs. 10.3 bps; 7.4 bps vs. 10.7 bps).

While Panel A of Table 3 shows that the lead-lag effect is absent among the largest firms, I show here that the returns of the largest firms can be led by neighbors' returns, conditional on high *revision rate ratios*. Among the largest firms, the long-short return differential is only 17 bps per week and not statistically significant when *revision rate ratios* are low. As information production on neighbors outstrips that on firms, the return differential becomes twice as large and statistically significant at the 5% level.

I acknowledge that revision frequencies are at best a proxy for the quantity of information produced by analysts. One may expect that other features such as the quality and novelty of information may be more salient in the lead-lag mechanism. To

¹⁵ I form tercile (instead of quintile) portfolios to ensure that every portfolio is sufficiently populated.

address this concern, I repeat the analysis with *revision rate ratios* constructed only with bold forecast revisions (Clement and Tse, 2005). As bold forecasts are deviations from both the analyst’s prior forecast and the consensus, the resulting *revision rate ratios* likely capture imbalances in production of novel information. The results remain qualitatively and quantitatively similar under this alternative definition of *revision rate ratios*.

3.6. Industry effects

Given that the interfirm networks are constructed on the foundation of analyst coverage, and that industry expertise is an essential and valued component of analyst skill (Boni and Womack, 2006; Kadan et al., 2012; Brown et al., 2015), I examine the role of industry effects in my findings. A candidate explanation of my results is industry momentum (Moskowitz and Grinblatt, 1999). If analysts tend to specialize in a narrow set of industries, it follows that my portfolio formation procedure will tend to produce firm-neighbor pairs that belong to the same industries. This implies that sorts on lagged returns of neighbors may be mechanically similar to sorts on lagged industry returns, and that my lead-lag effect is completely driven by industry momentum. This is a natural concern in light of the fact that around 40% of firm-neighbor pairs reside in the same industries.

In the tests to follow, I adopt three adjustment methods to disentangle the influence of industry effects from my findings. In the first method, I constrained my algorithm such that each firm is only matched with neighbors that do not reside in its industry. Hence, this directly removes the influence of industry effects in the information diffusion from neighbors to firms. Thereafter, I perform the baseline portfolio strategy by buying firms in the top quintile of lagged neighbors’ returns and selling firms in the corresponding bottom quintile.

In the following two methods, I constrain my algorithm to pair each firm with neighbors that are in its industry to maximize my strategy’s exposure to industry effects. In the second method, I sort firms within their industries on neighbors’ returns. Specifically, within each industry at a point in time, I buy firms in the top quintile of neighbors’ returns and short firms in the bottom quintile. Thereafter, I average the

long-short returns across all industries. This strategy is industry-neutral because the influence of industry effects is cancelled out by the long-short action within each industry. Third, I perform a global sort of firms on lagged neighbors' returns but I subtract contemporaneous industry returns from firm returns to correct for industry effects. If my findings are entirely driven by industry momentum, I expect these three alternative strategies to be unprofitable.

[Insert Table 7]

Table 7 shows that industry effects can only partially explain the lead-lag effect. The two-digit SIC industry classification scheme is used in Columns (1) through (3), following Moskowitz and Grinblatt (1999). Constraining firms and neighbors to be of different industries, the first adjusted strategy yields a long-short return of 6 bps per week in Column (1). The second and third industry-neutral strategies yield weekly returns of 5.7 bps and 9.1 bps in Columns (2) and (3), respectively. As a robustness check, I use the Fama-French 49 industries as an alternative industry classification scheme. Weekly long-short return differentials in Columns (4) through (6) are statistically significant and of similar magnitudes, ranging from 5.8 bps to 9.2 bps. Comparing the results in Table 7 with the baseline long-short returns of 9.4 bps in Table 2, a back-of-the-envelope calculation indicates that up to 40% of my baseline strategy profitability is attributable to industry effects.

Overall, I show that while the profitability of my strategy is not entirely driven by industry effects, they have a sizeable economic influence. Since companies within an industry face similar competitive pressures and regulatory environments, it is unsurprising that information commonalities between stocks (as represented in the interfirm networks) have non-trivial overlaps along industry lines. Notwithstanding the results in Table 7, I argue that this paper fundamentally contributes to our understanding of industry momentum by proposing and empirically testing a mechanism to explain it. Additionally, Moskowitz and Grinblatt (1999) and Hou (2007) also argue that industry momentum may be explained by information being received by industry leaders first before diffusing to other firms within the industry. In this vein, my paper examines *why* information production, in the presence of complementarities, is preferentially tilted towards certain firms in a network.

3.7. Portfolio performance before and after formation

To better understand the return patterns from my portfolio strategy, I track the performance of portfolios sorted on lagged neighbors' returns up to 10 weeks before and after the formation week. Particularly, I examine whether the return patterns are driven by some form of investors' overreaction, in which case I expect to observe a reversal in the longer term. In contrast, if the lead-lag effect is driven by diffusion of fundamentally relevant information, I should observe no reversals after firms' returns respond to neighbors' returns.

[Insert Figure 4]

Figure 4 shows that investors update firms' prices to information produced on neighbors relatively quickly. I find that the gulf in returns between the extreme quintile portfolios is largest in the first week ($w+1$) after portfolio formation; this gulf persists over the subsequent two weeks but shrinks dramatically from $w+4$ onwards. The long-short return differential is statistically distinguishable from zero in weeks $w+5$ and $w+6$ but its magnitude is only about one-thirds that in $w+2$ (my baseline strategy profitability). More importantly, I find no evidence of reversals in portfolio returns that would have been compatible with a behavioral explanation.

Although pre-formation portfolio returns are flat, I find that firms in the top quintile of lagged neighbors' returns exhibit higher returns up to three weeks before the formation week. Asymmetrically, I do not find that firms in the bottom quintile significantly underperforms the other firms. A potential explanation is that investors extrapolate information produced on neighbors to firms. While investors may bid up the prices of firms' stocks in response to positive news, short-sale constraints prevent them from impounding negative news into the firms' stock prices completely. After uncertainty about the shocks is resolved at week w , stock prices of firms, especially those in the bottom quintile, are then updated. To support the interpretation that

information is primarily transmitted from neighbors to firms¹⁶, I turn to a vector autoregressive analysis in the next section.

3.8. Vector autoregressions

In this section, I employ vector autoregressions (VARs) to gain sharper insights on the lead-lag dynamics between firms and neighbors. In the following tests, I distinguish between the information diffusion hypothesis and an alternative hypothesis based on nonsynchronous trading. Additionally, the VARs can also shed light on the primary direction of information diffusion.

In examining the information leadership of neighbors over firms, I first hold firm size constant. I do so because prior studies (e.g., Lo and MacKinlay, 1990; Brennan et al., 1993; Hou, 2007) already show that firm size is related to the speed of adjustment to information.¹⁷ To this end, my evidence in Panel A of Table 2 also shows that the lead-lag effect seems to be more pronounced in smaller firms. Therefore, controls for firm size helps me to filter out the influence of the size effect from my results. In each week, I form five portfolios of firms sorted on firm size. For each firm portfolio, I form a corresponding neighbor portfolio comprising of the firms' neighbors. I track the equal-weighted weekly returns for both types of portfolios.

On one hand, if the lead-lag effect is driven by firms' own autocorrelations coupled with high contemporaneous correlations between firms and neighbors, then cross-autoregressive terms in the VARs should have no predictive power once I control for lagged firms' returns ($\text{Ret}_{F,w-1}$). On the other hand, if the lead-lag effect is driven by quicker responses of neighbors' stock prices to common information, then lagged neighbors' returns ($\text{Ret}_{N,w-1}$) should continue to predict current firms' returns ($\text{Ret}_{F,w}$).

[Insert Table 8]

¹⁶ Under an alternative explanation, albeit one that is inconsistent with model predictions, investors may observe news on firms and use them to value neighbors; this mechanism can generate the return patterns in Figure 4.

¹⁷ Some possible drivers of this relation are investor attention (Merton, 1987), analyst coverage (Bhushan, 1989), and trading volume (Admati and Pfleiderer, 1988).

Panel A of Table 8 presents results from VARs of firms' and neighbors' returns, estimated for each firm size quintile separately. Across all firm size quintiles, I find that neighbors' returns Granger-cause firms' returns, that is— $\text{Ret}_{N,w-1}$ significantly predicts $\text{Ret}_{F,w}$ even after controlling for $\text{Ret}_{F,w-1}$. This result suggests that the lead-lag effect is unlikely driven by a nonsynchronous trading explanation. In addition, I find that the predictive power of $\text{Ret}_{N,w-1}$ on current firms' returns is generally greater than that of $\text{Ret}_{F,w-1}$ on current neighbors' returns and such cross-equation differences are statistically significant. This implies that information primarily diffuses from neighbors to firms, and not vice versa. Finally, in all firm size quintiles, returns of firm portfolios are more predictable than the returns of neighbor portfolios. This is expected if firms' stock prices respond more slowly to common information than neighbors' stock prices do.

In Panel B, I corroborate my prior findings that even small neighbors have information leadership by estimating VARs for each neighbor size separately. Even among the smallest neighbors, I continue to find that neighbors' returns Granger-cause firms' returns. In the other quintiles, I also find that the predictive power of lagged neighbors' returns on current firms' returns is incremental to controls for lagged firms' returns. Consistent with results in Panel A, cross-equation differences in estimated loadings also supports the notion that information diffusion occurs from neighbors to firms, and not in the opposite direction.

Overall, the evidence in Table 8 indicates that the information leadership of neighbors over firms persists across distributions of firm size and neighbor size. My results suggest that the lead-lag effect between firms and neighbors is incompatible with a nonsynchronous trading explanation. Instead, the results are consistent with the hypothesis that information primarily diffuses from neighbors to firms due to imbalances in their information environments.

3.9. Propagation of real effects

If firms and neighbors share information commonalities, they may be exposed to common economic shocks and make correlated corporate decisions. Since acquisitions

are salient corporate events and are readily observable, I use them to test the propagation of real effects from neighbors to firms.

I track the acquisition activity of firms and their neighbors in every month. In a given month, if a firm completes an acquisition, the firm acquisition indicator equals unity, and equals zero otherwise. The neighbor acquisition indicator is defined symmetrically. I estimate logistic and conditional-logistic regressions of firm acquisition indicator on neighbor acquisition indicator to test if acquisitiveness propagates from neighbors to firms. Regression specifications also include controls for firm's market capitalization, past performance, book-to-market and cash flow. In some specifications, I also include Fama-French 49 industries and calendar year indicators, industry-year fixed effects, or firm fixed effects.

[Insert Table 9]

Table 9 shows that the acquisition activity of neighbors heralds acquisitions by firms. Estimated coefficients are presented in odds ratios and are averaged over 100 rounds of simulations. Column (1) shows that the odds of a firm engaging in acquisitions increase by about 36% if its corresponding neighbor is involved in acquisitions in the preceding month. This implies that acquisitiveness propagates from neighbors to firms. I find similar results in Column (2) where I include Fama-French 49 and calendar year indicators. Next, I estimate fixed-effect logistic regressions in Columns (3) and (4). The link between the acquisitiveness of firms and neighbors survives the inclusion of industry-year and firm fixed effects, suggesting that my findings are not completely driven by unobserved heterogeneity at these dimensions. To test if these findings are driven by network effects, I replace the neighbor acquisition indicator with the pseudo-neighbor acquisition indicator and re-estimate conditional-logistic regressions in Column (5). My results indicate that the acquisition activities of pseudo-neighbors are unassociated with that of firms. These results suggest that the propagation of acquisitiveness from neighbors to firms is not completely driven by broader trends such as merger waves.

4. Conclusions

I show that information diffusion can emerge in information networks formed on shared analyst coverage. Due to complementarities in information production, my model shows that more information is produced on neighbors than on firms, even on average. Therefore, lagged neighbors' returns can predict future firms' returns. A strategy that buys firms in the top quintile of neighbors' returns and shorts firms in the corresponding bottom quintile yields a weekly return of 9.4 basis points. The profitability of this strategy is robust to controls for size, momentum, book-to-market, illiquidity, and idiosyncratic volatility. However, the effect is more pronounced among smaller and illiquid firms. While industry effects are compatible with my proposed mechanism, they explain only up to 40% of the lead-lag effect between firms and neighbors. Supplementary evidence from vector autoregressions suggests that the lead-lag effect is not explained by a nonsynchronous trading explanation. The results also indicate that information primarily diffuses from neighbors to firms, and not vice versa. Evidence from corporate acquisitiveness suggests that interfirm linkages are steeped in real economic roots.

References

- Acemoglu, D., Carvalho, V., Ozdaglar, A., Tahbaz-Salehi, A., 2012. The network origins of aggregate fluctuations. *Econometrica* 80, 1977-2016.
- Admati, A., Pfleiderer, P., 1988. A theory of intraday patterns: volume and price variability. *Review of Financial Studies* 1, 3-40.
- Ahern, K., 2013. Network centrality and the cross section of stock returns. Unpublished working paper, University of Southern California
- Ahern, K., Harford, J., 2014. The importance of industry links in merger waves. *Journal of Finance* 69, 527-576.
- Amihud, Y., 2002. Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets* 5, 31-56.
- Ang, A., Hodrick, R., Xing, Y., Zhang, X., 2009. High idiosyncratic volatility and low returns: International and further US evidence. *Journal of Financial Economics* 91, 1-23.
- Asquith, P., Mikhail, M., Au, A., 2005. Information content of equity analyst reports. *Journal of Financial Economics* 75, 245-282.
- Bhushan, R., 1989. Firm characteristics and analyst following. *Journal of Accounting and Economics* 11, 255-274.
- Boni, L., Womack, K., 2006. Analysts, industries, and price momentum. *Journal of Financial and Quantitative Analysis* 41, 85-109.
- Boone, A., White, J., 2015. The effect of institutional ownership on firm transparency and information production. *Journal of Financial Economics* 117, 508-533.
- Borgatti, S., 2005. Centrality and network flow. *Social networks* 27, 55-71.
- Brandt, M., Brav, A., Graham, J., Kumar, A., 2010. The idiosyncratic volatility puzzle: time trend or speculative episodes? *Review of Financial Studies* 23, 863-899.
- Brennan, M., Jegadeesh, N., Swaminathan, B., 1993. Investment analysis and the adjustment of stock prices to common information. *Review of Financial Studies* 6, 799-824.
- Brown, L., Call, A., Clement, M., Sharp, N., 2015. Inside the “black box” of sell-side financial analysts. *Journal of Accounting Research* 53, 1-47.
- Campbell, J., Lettau, M., Malkiel, B., Xu, Y., 2001. Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. *Journal of Finance* 56, 1-43.
- Carhart, M., 1997. On persistence in mutual fund performance. *Journal of Finance* 52, 57-82.
- Christakis, N., Fowler, J., 2010. Social network sensors for early detection of contagious outbreaks. *PloS One* 5, e12948.
- Chordia, T., Roll, R., Subrahmanyam, A., 2008. Liquidity and market efficiency. *Journal of Financial Economics* 87, 249-268.
- Clement, M., Tse, S., 2005. Financial analyst characteristics and herding behavior in forecasting. *Journal of finance* 60, 307-341.
- Cohen, R., Havlin, S., Ben-Avraham, D., 2003. Efficient immunization strategies for computer networks and populations. *Physical Review Letters* 91, 247901 - 247901.
- Cohen, L., Frazzini, A., 2008. Economic links and predictable returns. *Journal of Finance* 63, 1977-2011.
- Cohen, L., Lou, D., 2012. Complicated firms. *Journal of Financial Economics* 104, 383-400.

- Daniel, K., Grinblatt, M., Titman, S., Wermers, R., 1997. Measuring mutual fund performance with characteristic-based benchmarks. *Journal of Finance* 52, 1035-1058.
- Diamond, D., Verrecchia, R., 1987. Constraints on short-selling and asset price adjustment to private information. *Journal of Financial Economics* 18, 277-311.
- Degeorge, F., Derrien, F., Kecskes, A., Michenaud, S., 2013. Do analysts' preferences affect corporate policies? Unpublished working paper, European Corporate Governance Institute.
- Engelberg, J., Ozoguz, A., Wang, S., 2017. Know thy neighbor: industry clusters, information spillovers and market efficiency. *Journal of Financial and Quantitative Analysis*, forthcoming.
- Fama, E., MacBeth, J., 1973. Risk, return, and equilibrium: empirical tests. *Journal of Political Economy* 81, 607-636.
- Fama, E., French, K., 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33, 3-56.
- Fama, E., French, K., 2015. A five-factor pricing model. *Journal of Financial Economics* 116, 1-22.
- Feld, S., 1991. Why your friends have more friends than you do. *American Journal of Sociology* 96, 1464-1477.
- Fisher, R., 1925. *Statistical methods for research workers*. Genesis Publishing Pvt. Ltd.
- Gibbons, M., Ross, S., Shanken, J., 1989. A test of the efficiency of a given portfolio. *Econometrica* 57, 1121-1152.
- Giroud, X., Mueller, H., 2010. Does corporate governance matter in competitive industries? *Journal of Financial Economics* 95, 312-331.
- Gomes, A., Gopalan, R., Leary, M., Marcet, F., 2016. Analyst coverage network and corporate financial policies. Unpublished working paper, Washington University in St. Louis.
- Hameed, A., Morck, R., Shen, J., Yeung, B., 2015. Information, analysts, and stock return comovement. *Review of Financial Studies* 28, 3153-3187.
- Hoberg, G., Phillips, G., Prabhala, N., 2014. Product market threats, payouts, and financial flexibility. *Journal of Finance* 69, 293-324.
- Hong, H., Lim, T., Stein, J., 2000. Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies. *Journal of Finance* 55, 265-295.
- Hong, H., Torous, W., Valkanov, R., 2007. Do industries lead stock markets? *Journal of Financial Economics* 83, 367-396.
- Hou, K., 2007. Industry information diffusion and the lead-lag effect in stock returns. *Review of Financial Studies* 20, 1113-1138.
- Jackson, M., 2016. The friendship paradox and systematic biases in perceptions and social norms. Unpublished working paper, Stanford University.
- Jegadeesh, N., Titman, S., 2001. Profitability of momentum strategies: an evaluation of alternative explanations. *Journal of Finance* 56, 699-720.
- Kadan, O., Madureira, L., Wang, R., Zach, T., 2012. Analysts' industry expertise. *Journal of Accounting and Economics* 54, 95-120.
- Larcker, D., So, E., Wang, C., 2013. Boardroom centrality and firm performance. *Journal of Accounting and Economics* 55, 225-250.

- Lo, A., MacKinlay, A., 1990. When are contrarian profits due to stock market overreaction? *Review of Financial Studies* 3, 175-205.
- MacKay, P., Phillips, G., 2005. How does industry affect firm financial structure? *Review of Financial Studies* 18, 1433-1466.
- Menzly, L., Ozbas, O., 2010. Market segmentation and cross-predictability of returns. *Journal of Finance* 65, 1555-1580.
- Merton, R., 1987. A simple model of capital market equilibrium with incomplete information. *Journal of Finance* 42, 483-510.
- Moskowitz, T., Grinblatt, M., 1999. Do industries explain momentum? *Journal of Finance* 54, 1249-1290.
- Muslu, V., Rebello, M., Xu, Y., 2014. Sell-side analyst research and stock comovement. *Journal of Accounting Research* 52, 911-954.
- Piotroski, J., Roulstone, D., 2004. The influence of analysts, institutional investors, and insiders on the incorporation of market, industry, and firm-specific information into stock prices. *The Accounting Review* 79, 1119-1151.
- Scherbina, A., Schlusche, B., 2015a. Cross-firm information flows and the predictability of stock returns. Unpublished working paper, University of California, Davis.
- Scherbina, A., Schlusche, B., 2015b. Economic linkages inferred from news stories and the predictability of stock returns. Unpublished working paper, University of California, Davis.
- Veldkamp, L., 2006. Information markets and the comovement of asset prices. *Review of Economic Studies* 73, 823-845.
- Watts, D., Strogatz, S., 1998. Collective dynamics of 'small-world' networks. *Nature* 393, 440-442.
- Wermers, R., 2003. Is money really 'smart'? New evidence on the relation between mutual fund flows, manager behavior, and performance persistence. Unpublished working paper, University of Maryland.
- Xu, Y., Malkiel, B., 2003. Investigating the behavior of idiosyncratic volatility. *Journal of Business* 76, 613-645.

Appendix I.

Friendship paradox in networks

The friendship paradox states that neighbors have strictly higher degrees (i.e. number of neighbors) than firms in a network, on average. I motivate the intuition of the friendship paradox phenomenon in networks with a hypothetical network below. As will be made clear in Section 1, the intuition gleaned in this simple example can be generalized to more sophisticated networks.

There are 5 firms – **A**, **B**, **C**, **D**, and **E** – represented by spheres in this network. In each sphere, there is a number enclosed in parentheses. This number is the degree, or the number of direct neighbors, of the firm. For example, by being linked to firms **A**, **B**, **C**, and **D**, firm **E** has four neighbors and therefore a degree of 4. For generic firm i and j in set N , the average degree of firms is simply the arithmetic average of their degrees.

$$\begin{aligned} \text{Average degree of firms} &= \frac{1}{|N|} \sum_i \text{degree}(i) \\ &= \frac{1 + 1 + 2 + 2 + 4}{5} = 2 \end{aligned}$$

Next we compute the average degrees of neighbors. For each firm i , we sum up the degree of every neighbor j . Next, we cumulate these sums across all firms i and normalize by the total number of neighbors (which is simply the sum of firms' degrees) to obtain the average degrees of neighbors. The working example below clarifies this algorithm.

$$\begin{aligned} \text{Average degree of neighbors} &= \frac{1}{\sum_i \text{degree}(i)} \sum_i \sum_j \text{degree}(j) \\ &= \frac{\overset{\text{E}}{(\tilde{4})} + \overset{\text{E}}{(\tilde{4})} + \overset{\text{E}}{(\tilde{4} + \tilde{2})} + \overset{\text{D}}{(\tilde{4} + \tilde{2})} + \overset{\text{E}}{(\tilde{4} + \tilde{2})} + \overset{\text{C}}{(\tilde{1} + \tilde{1} + \tilde{2} + \tilde{2})}}{1 + 1 + 2 + 2 + 4} \\ &= 2.6 \end{aligned}$$

Unsurprisingly, the friendship paradox holds in this network – neighbors have strictly higher degrees than firms, on average. To understand what drives the friendship paradox, notice that firms are counted multiple times in proportion to their degrees (e.g. **A** is counted 1 time and **E** is counted 4 times). Intuitively, if a firm has a degree d , it must be counted as a neighbor by d other firms. Therefore, high-degree firms are over-sampled in the average firm's circle of neighbors. Equivalently, this bias means that the average firm's circle of neighbors tend to include firms with higher-than-average degrees, hence producing the friendship paradox.

Figure 1.
Visual plot of interfirm network in year 2005

Figure 1 contains a visual plot of the information network formed on shared analyst coverage between firms in calendar year 2005. Two firms are linked if they are jointly covered by at least one analyst in the calendar year. Nodes represent firms (refer to inset magnifications) and lines represent interfirm linkages. The visual plot of the network is presented in the OpenOrd layout (Martin et al., 2011) and generated on the software platform Gephi.

Figure 2.
Neighbor-to-firm comparisons of characteristics

Figure 2 tracks the neighbor-to-firm ratios of various financial characteristics across years in lines (items 1 through 6). Where the ratios of a particular characteristic are greater than unity (above the red line), neighbors have higher values of that characteristic than firms. Figure 2 also plots the proportions of firm-neighbor pairs across years that belong to identical FF-49 or SIC-2d industries (items 7 and 8) in bars. Definitions of the financial characteristics are contained in Section 3.2.

Figure 3.
Sample proportions by SIC divisions

Figure 3 tracks sample proportions by 10 SIC divisions. For each calendar year, I compute the proportion of SIC divisions in the sample. Thereafter, I average these proportions across all years for every SIC division and present them in Figure 3. The sample composition by SIC divisions does not vary meaningfully across years. The definitions of SIC divisions are available on the SEC portal.

Figure 4.
Portfolio performance before and after formation

Figure 4 tracks the performance of firm portfolios sorted on neighbors' returns. At the beginning of each formation week, for each firm in the network, I randomly select a neighbor and match it with the firm. Thereafter, I track neighbors' returns over the formation week w , and form quintile portfolios of firms sorted on their neighbors' returns. I then track the performance of these portfolios up to 10 weeks around formation (from $w-10$ to $w+10$).

Table 1.
Descriptive statistics

This table shows descriptive statistics of the network and sample by year. Panel A reports network topological statistics. The largest component of the interfirm network is the largest subset of firms that are connected to one another, either directly or indirectly. Put differently, any two firms in the largest component can reach each other along interfirm linkages. We report average shortest path lengths, diameter, and transitivity for the largest component in each year. Average shortest path length is defined as the average number of steps along the shortest paths between all possible pairs of firms. Diameter is the maximum length of shortest paths between all possible pairs of firms. Transitivity is the fraction of possible triangles in the network that are triangles. A possible triangle is identified by two firms which share a common neighbor; a possible triangle is a triangle if these two firms are also linked. I also describe the distributions of degrees and the friendship paradox bias term in Panel A. Panel B reports means and standard deviations of firm-level statistics in each year. Definitions of financial characteristics are contained in Section 3.2.

Panel A: Network topological statistics

Calendar year	Number of firms		Avg. shortest path length	Diameter	Transitivity	Degree distribution			
	Full network	Largest component				$E[d]$	$\min[d]$	$\max[d]$	$\frac{\text{Var}[d]}{E[d]}$
1995	4,528	4,528	2.742	7	0.372	80.071	1	407	53.985
1996	5,223	5,216	2.806	7	0.374	79.395	1	429	56.532
1997	5,447	5,440	2.878	8	0.374	74.449	1	427	52.330
1998	5,306	5,291	2.855	7	0.432	75.662	1	443	53.077
1999	5,004	4,993	2.923	9	0.365	66.822	1	371	44.696
2000	4,733	4,723	2.929	7	0.360	64.936	1	357	43.841
2001	4,117	4,091	3.021	7	0.373	58.237	1	270	33.286
2002	3,965	3,959	3.095	9	0.397	55.981	1	239	27.859
2003	3,905	3,854	3.070	9	0.396	56.513	1	266	29.277
2004	4,165	4,112	3.134	10	0.392	58.208	1	231	28.554
2005	4,306	4,242	3.180	11	0.399	60.161	1	227	26.260
2006	4,337	4,298	3.117	10	0.397	64.928	1	255	27.144
2007	4,328	4,300	3.063	9	0.391	68.992	1	289	29.170
2008	4,070	4,041	3.068	10	0.401	68.916	1	257	26.308
2009	3,947	3,915	3.001	7	0.400	68.758	1	305	28.357
2010	3,817	3,791	2.965	7	0.399	71.163	1	283	27.481
2011	3,711	3,700	2.941	7	0.400	73.934	1	300	25.903
2012	3,714	3,706	2.882	7	0.414	76.722	1	323	22.927
2013	3,798	3,793	2.887	8	0.405	80.085	1	294	26.837
2014	3,613	3,607	2.935	8	0.425	79.394	1	279	26.022

Panel B: Firm-level summary statistics

Calendar year	Mean					Standard deviation				
	Market-cap	b/m	(-12, -2) returns	I-volatility	Illiquidity	Market-cap	b/m	(-12, -2) returns	I-volatility	Illiquidity
1995	1532.7	2.012	1.550	2.288	6.014	5824.7	14.981	0.040	0.029	14.779
1996	1755.6	1.820	0.635	2.549	5.535	7164.9	14.503	0.040	0.032	13.789
1997	2094.9	1.718	1.529	2.664	4.518	9294.2	16.570	0.039	0.031	10.903
1998	2540.9	2.008	-2.076	3.886	3.997	13152.9	24.753	0.045	0.053	9.575
1999	3441.6	1.310	-0.050	4.208	3.486	19764.4	9.154	0.052	0.057	8.757
2000	3477.7	2.159	-0.557	5.875	3.457	17919.2	18.116	0.064	0.057	9.163
2001	3455.9	1.763	-1.098	4.756	3.095	17384.1	25.022	0.060	0.055	8.951
2002	2772.2	2.305	-2.492	3.743	2.509	13146.5	17.484	0.057	0.060	7.939
2003	3671.4	2.387	2.992	2.125	1.574	15922.0	25.062	0.040	0.088	5.838
2004	3995.8	2.877	0.341	1.486	1.213	16633.1	40.935	0.035	0.018	5.289
2005	4183.8	1.771	0.147	1.280	2.534	16858.1	14.870	0.033	0.018	18.996
2006	4603.4	1.375	0.814	1.260	2.028	18443.2	9.096	0.030	0.015	13.230
2007	4828.5	1.093	0.158	1.576	2.300	19394.8	4.967	0.038	0.028	19.127
2008	2859.1	2.134	-5.281	5.311	7.106	12834.0	13.657	0.050	0.134	59.174
2009	3808.3	1.383	2.670	4.323	11.221	14863.0	8.200	0.049	0.087	82.355
2010	4736.0	1.648	1.528	1.420	1.595	17011.0	12.207	0.031	0.018	13.575
2011	4563.2	2.368	-0.375	1.545	3.148	17183.4	24.213	0.033	0.019	28.096
2012	5361.5	2.255	0.777	1.374	3.120	20125.3	21.479	0.033	0.022	25.364
2013	6771.3	2.095	2.543	1.169	2.771	23469.3	19.845	0.033	0.021	31.303
2014	7337.3	1.959	0.200	1.250	0.842	25782.7	22.491	0.031	0.026	8.750

Table 2.
Returns to portfolios sorted on neighbors' returns

This table shows weekly equal-weighted average returns to portfolios sorted on lagged stock returns of neighbors in the network. At the beginning of each formation week, for each firm in the network, I randomly select a neighbor and match it with the firm. Thereafter, I track neighbors' returns over the formation week w , and form quintile portfolios of firms sorted on their neighbors' returns. I skip a week after the formation week and hold the portfolios for one week (in $w+2$). Apart from portfolio returns, I also report the average ratio of neighbors' degrees to firms' degrees, average number of firms, average firm size, average firms' book-to-market, and average firms' returns over the past four weeks in each portfolio. Panel A contains results from the baseline implementation of this portfolio strategy. Panel B contains results of the same baseline implementation but excludes portfolio-months in January and December. Panel C contains results from firms matched with pseudo-neighbors (i.e., randomly selected stocks that are not firms' neighbors in the network). Statistics presented in the panels are averages over 100 rounds of simulations.

Panel A: Baseline specification

Neighbor returns	Mean returns	Degree ratio	Number of firms	Firm size	Firm b/m	(-4, -1) returns
Low	0.205	1.209	587	4520.1	1.777	0.854
N2	0.262	1.291	588	4846.6	1.933	0.967
N3	0.276	1.309	588	4936.5	1.902	0.999
N4	0.283	1.304	588	4898.8	1.875	1.025
High	0.299	1.225	589	4592.2	1.770	1.060
High minus Low	0.094***					
Fisher $\chi^2 (\times 10^3)$	6.234					

Panel B: Exclude December/January observations

Neighbor returns	Mean returns	Degree ratio	Number of firms	Firm size	Firm b/m	(-4, -1) returns
Low	0.139	1.206	588	4512.3	1.775	0.603
N2	0.207	1.291	589	4840.9	1.939	0.749
N3	0.228	1.310	589	4931.7	1.900	0.794
N4	0.240	1.306	589	4893.6	1.869	0.820
High	0.253	1.227	590	4593.3	1.766	0.844
High minus Low	0.114***					
Fisher $\chi^2 (\times 10^3)$	7.355					

Panel C: Pseudo-neighbors benchmark

Neighbor returns	Mean returns	Degree ratio	Number of firms	Firm size	Firm b/m	(-4, -1) returns
Low	0.266	0.847	584	4751.5	1.835	0.985
N2	0.266	0.965	585	4750.8	1.836	0.983
N3	0.265	0.996	585	4749.3	1.835	0.984
N4	0.265	0.986	585	4750.0	1.834	0.984
High	0.265	0.873	585	4751.0	1.836	0.985
High minus Low	-0.001					
Fisher $\chi^2 (\times 10^3)$	0.958					

Table 3.
Independent portfolio double-sorts

This table shows weekly equal-weighted average returns to portfolios independently double-sorted on 1) lagged stock returns of neighbors, and 2) firm/neighbor characteristics. Sizes of firms and neighbors are defined as their market capitalization at the time of portfolio formation. Book equity (b) is defined as the Compustat book value of stockholders' equity, plus balance sheet deferred taxes and investment tax credit (where available), less the book value of preferred stock. Depending on their availability, redemption, liquidation, or par value (in this order of priority) is used to estimate the book value of preferred stock. To ensure that these variables are known at portfolio formation, I match the annual book equity value for fiscal years ending in calendar year $y-1$ with portfolio formation dates starting in calendar year y . Book equity is then scaled by market capitalization at portfolio formation to form b/m . Momentum is the firms' returns from $w-4$ to $w-1$. Illiquidity is the summation of the daily ratios of return magnitudes to trading volume over $w-4$ to $w-1$. For the illiquidity measure of a firm to be valid, the firm must have no less than 10 valid data points over $w-4$ to $w-1$. Following Campbell et al. (2001), I define idiosyncratic volatility of a firm as the summation of squared differences between daily returns and daily Fama-French 49 industry returns over $w-4$ to $w-1$. All summed values are then scaled to 21 days to ensure comparability (Brav et al. 2010). At the beginning of each formation week w , for each firm in the network, I randomly match each firm with one of its neighbors. Thereafter, I track neighbors' returns over the formation week, and form quintile portfolios of firms independently double-sorted on their neighbors' returns and firm/neighbor characteristics. I skip a week after the formation week and hold the portfolios for one week (in $w+2$). Panels A through F show details of portfolios independently double-sorted on neighbors' returns and firm size, b/m , momentum, idiosyncratic volatility, illiquidity, and neighbor size. Statistics presented in the panels are averages over 100 rounds of simulations.

Table 3. (Continued)

Panel A: Firm size and neighbors' returns

	Firm size quintiles				
	Small	S2	S3	S4	Big
Neighbor returns	Mean returns				
Low	0.184	0.211	0.222	0.212	0.213
N2	0.255	0.269	0.273	0.264	0.248
N3	0.278	0.294	0.283	0.269	0.253
N4	0.298	0.310	0.288	0.275	0.244
High	0.326	0.334	0.306	0.277	0.244
High minus Low	0.142 ^{***}	0.122 ^{***}	0.083 ^{***}	0.065 ^{***}	0.031
Fisher χ^2 ($\times 10^3$)	9.895	7.001	3.537	2.237	0.958
	Number of firms				
Low	125	122	118	112	110
N2	115	116	117	120	121
N3	113	114	117	121	123
N4	113	115	117	121	122
High	122	122	118	114	113
	Neighbors' returns				
Low	-0.073	-0.072	-0.071	-0.070	-0.069
N2	-0.020	-0.020	-0.020	-0.020	-0.020
N3	0.001	0.001	0.001	0.001	0.001
N4	0.022	0.022	0.022	0.022	0.022
High	0.084	0.083	0.081	0.080	0.079
	Firm size (\$'mil)				
Low	127.3	372.2	861.5	2148.5	20285.8
N2	128.5	373.5	865.9	2158.7	20229.4
N3	128.7	373.9	866.9	2161.4	20227.6
N4	128.8	373.7	866.4	2159.9	20219.6
High	127.8	372.4	862.4	2150.9	20304.4

Panel B: Firm b/m and neighbors' returns

	Firm b/m quintiles				
	Low	G2	G3	G4	High
Neighbor returns	Mean returns				
Low	0.205	0.209	0.231	0.240	0.193
N2	0.242	0.241	0.261	0.288	0.276
N3	0.249	0.255	0.266	0.310	0.290
N4	0.245	0.257	0.273	0.319	0.305
High	0.253	0.273	0.295	0.334	0.314
High minus Low	0.048 ^{***}	0.064 ^{***}	0.065 ^{***}	0.093 ^{***}	0.121 ^{***}
Fisher χ^2 ($\times 10^3$)	1.709	2.688	2.862	5.078	6.780
	Number of firms				
Low	132	121	114	109	111
N2	112	116	119	121	121
N3	107	115	120	124	122
N4	111	115	119	122	121
High	130	121	115	111	112
	Neighbors' returns				
Low	-0.074	-0.072	-0.070	-0.069	-0.069
N2	-0.020	-0.020	-0.020	-0.020	-0.020
N3	0.001	0.001	0.001	0.001	0.001
N4	0.022	0.022	0.022	0.022	0.022
High	0.086	0.082	0.080	0.079	0.078
	Firm b/m				
Low	0.099	0.337	0.519	0.758	7.436
N2	0.097	0.338	0.520	0.759	7.702
N3	0.098	0.339	0.520	0.760	7.461
N4	0.092	0.338	0.520	0.760	7.402
High	0.096	0.337	0.519	0.759	7.388

Table 3. (Continued)

Panel C: Firm momentum and neighbors' returns

	Firm momentum quintiles				
	Low	M2	M3	M4	High
Neighbor returns	Mean returns				
Low	0.241	0.227	0.207	0.178	0.190
N2	0.290	0.285	0.265	0.235	0.231
N3	0.298	0.309	0.282	0.244	0.245
N4	0.299	0.314	0.289	0.250	0.257
High	0.309	0.330	0.306	0.256	0.278
High minus Low	0.068***	0.103***	0.099***	0.078***	0.088***
Fisher $\chi^2 (\times 10^3)$	2.465	6.498	6.615	4.191	4.042
	Number of firms				
Low	129	113	109	111	125
N2	113	120	121	120	114
N3	109	122	125	122	111
N4	111	120	122	121	113
High	124	113	111	114	126
	Neighbors' returns				
Low	-0.073	-0.070	-0.069	-0.069	-0.072
N2	-0.020	-0.020	-0.020	-0.020	-0.020
N3	0.001	0.001	0.001	0.001	0.001
N4	0.022	0.022	0.022	0.022	0.022
High	0.084	0.080	0.079	0.080	0.084
	Firm momentum				
Low	-13.862	-3.559	0.775	5.215	16.338
N2	-13.192	-3.514	0.778	5.177	15.597
N3	-13.011	-3.501	0.779	5.166	15.432
N4	-13.127	-3.510	0.778	5.175	15.550
High	-13.711	-3.551	0.778	5.217	16.246

Panel D: Firm i-volatility and neighbors' returns

	Firm idiosyncratic volatility quintiles				
	Low	V2	V3	V4	High
Neighbor returns	Mean returns				
Low	0.241	0.241	0.227	0.202	0.178
N2	0.269	0.284	0.273	0.262	0.220
N3	0.279	0.291	0.283	0.275	0.234
N4	0.275	0.296	0.297	0.281	0.247
High	0.278	0.305	0.311	0.302	0.279
High minus Low	0.037***	0.064***	0.085***	0.100***	0.101***
Fisher $\chi^2 (\times 10^3)$	2.200	4.170	4.988	5.187	3.962
	Number of firms				
Low	96	106	116	128	142
N2	126	123	119	113	107
N3	135	126	118	109	99
N4	129	124	119	112	104
High	101	109	117	126	136
	Neighbors' returns				
Low	-0.066	-0.068	-0.070	-0.073	-0.076
N2	-0.019	-0.020	-0.020	-0.020	-0.020
N3	0.001	0.001	0.001	0.001	0.001
N4	0.022	0.022	0.022	0.022	0.022
High	0.074	0.077	0.080	0.084	0.089
	Firm idiosyncratic volatility				
Low	0.003	0.006	0.010	0.018	0.056
N2	0.003	0.006	0.010	0.018	0.054
N3	0.003	0.006	0.010	0.018	0.055
N4	0.003	0.006	0.010	0.018	0.054
High	0.003	0.006	0.010	0.018	0.056

Table 3. (Continued)

Panel E: Firm illiquidity and neighbors' returns

	Firm illiquidity quintiles				
	Low	I2	I3	I4	High
Neighbor returns	Mean returns				
Low	0.201	0.219	0.207	0.221	0.187
N2	0.241	0.265	0.268	0.278	0.253
N3	0.247	0.272	0.287	0.294	0.276
N4	0.240	0.277	0.299	0.305	0.292
High	0.234	0.290	0.311	0.333	0.323
High minus Low	0.033	0.071***	0.104***	0.111***	0.136***
Fisher $\chi^2 (\times 10^3)$	0.871	2.421	5.108	7.232	11.576
	Number of firms				
Low	115	118	120	119	114
N2	118	117	116	117	120
N3	119	117	116	116	121
N4	119	117	116	116	119
High	116	119	120	120	115
	Neighbors' returns				
Low	-0.070	-0.071	-0.072	-0.071	-0.070
N2	-0.020	-0.020	-0.020	-0.020	-0.020
N3	0.001	0.001	0.001	0.001	0.001
N4	0.022	0.022	0.022	0.022	0.022
High	0.080	0.082	0.082	0.082	0.081
	Firm illiquidity				
Low	0.294	1.204	3.352	10.731	279.310
N2	0.296	1.199	3.359	10.810	273.528
N3	0.296	1.198	3.362	10.833	276.378
N4	0.295	1.198	3.360	10.822	269.377
High	0.294	1.202	3.354	10.751	263.045

Panel F: Neighbor size and neighbors' returns

	Neighbor size quintiles				
	Small	S2	S3	S4	Big
Neighbor returns	Mean returns				
Low	0.222	0.208	0.200	0.191	0.206
N2	0.259	0.255	0.259	0.264	0.263
N3	0.281	0.274	0.274	0.277	0.277
N4	0.286	0.277	0.281	0.286	0.288
High	0.295	0.312	0.304	0.299	0.276
High minus Low	0.073***	0.105***	0.105***	0.108***	0.070***
Fisher $\chi^2 (\times 10^3)$	4.192	5.992	4.728	4.071	1.877
	Number of firms				
Low	159	132	113	97	87
N2	100	112	119	126	131
N3	90	105	119	132	142
N4	93	108	121	129	137
High	145	131	116	103	93
	Neighbors' returns				
Low	-0.082	-0.072	-0.068	-0.064	-0.059
N2	-0.020	-0.020	-0.020	-0.020	-0.019
N3	0.001	0.001	0.001	0.001	0.001
N4	0.022	0.022	0.022	0.022	0.022
High	0.104	0.083	0.075	0.070	0.063
	Neighbor Size (\$'mil)				
Low	123.5	458.9	1181.0	3187.4	26171.4
N2	134.7	468.8	1197.6	3222.1	30717.1
N3	136.1	470.3	1205.0	3235.8	31555.2
N4	135.9	469.4	1201.3	3232.7	30675.9
High	127.1	460.7	1180.9	3195.5	25677.9

Table 4.
Time-series multifactor tests for neighbor-returns quintiles

This table reports results of four-factor and six-factor regressions of weekly portfolio returns (sorted on neighbors' lagged stock returns) on Fama-French five factors and the momentum factor. At the beginning of each formation week, for each firm in the network, I randomly select a neighbor and match it with the firm. Thereafter, I track neighbors' returns over the formation week w , and form quintile portfolios of firms sorted on their neighbors' returns. I skip a week after the formation week and hold the portfolios for one week (in $w+2$). I also report the Gibbons, Ross, and Shanken (1989) tests for the four-factor and six-factor models. Time-series data of factors are drawn from Kenneth French's data library. Average Newey t-statistics are reported in parentheses. Coefficients and statistics presented in the table are averages over 100 rounds of simulations. ***, **, and * denote statistical significance of the Fisher's chi-squared test at the 1%, 5%, and 10% respectively.

Table 4. (Continued)

Portfolio	Factor sensitivities							Average R ²
	Alpha (%)	R _M - R _F	SMB	HML	UMD	RMW	CMA	
Low neighbor returns	-0.005 (0.35)	1.058*** (118.18)	0.767*** (44.44)	-0.051*** (3.55)	-0.139*** (14.62)			0.953
N2	0.049*** (3.40)	0.971*** (155.73)	0.654*** (54.43)	0.131*** (13.17)	-0.100*** (15.13)			0.970
N3	0.065*** (4.80)	0.935*** (159.58)	0.633*** (56.03)	0.171*** (18.25)	-0.089*** (14.21)			0.972
N4	0.073*** (5.30)	0.938*** (158.78)	0.663*** (58.21)	0.141*** (14.90)	-0.086*** (13.64)			0.972
High neighbor returns	0.092*** (5.21)	0.985*** (129.20)	0.778*** (52.99)	-0.026*** (2.17)	-0.103*** (12.69)			0.960
High minus Low	0.097*** (3.26)	-0.074*** (5.71)	0.011 (0.53)	0.024 (1.19)	0.036*** (2.66)			0.050
Low neighbor returns	0.009 (0.48)	1.032*** (102.97)	0.731*** (40.56)	0.010 (0.55)	-0.127*** (13.22)	-0.120*** (5.49)	-0.052*** (2.10)	0.951
N2	0.044*** (3.04)	0.981*** (138.71)	0.665*** (52.25)	0.103*** (7.19)	-0.105*** (15.49)	0.031*** (2.02)	0.035*** (1.99)	0.970
N3	0.055*** (4.12)	0.954*** (145.71)	0.654*** (55.57)	0.120*** (9.06)	-0.097*** (15.51)	0.067*** (4.70)	0.057*** (3.55)	0.971
N4	0.065*** (4.76)	0.953*** (143.23)	0.679*** (56.78)	0.100*** (7.44)	-0.093*** (14.52)	0.051*** (3.51)	0.047*** (2.86)	0.971
High neighbor returns	0.100*** (5.65)	0.973*** (113.05)	0.756*** (48.85)	-0.009 (0.59)	-0.097*** (11.77)	-0.080*** (4.26)	0.006*** (0.45)	0.959
High minus Low	0.090*** (3.01)	-0.058*** (3.99)	0.025 (0.97)	-0.019 (0.67)	0.030** (2.13)	0.040 (1.25)	0.058 (1.60)	0.055
Average GRS (4-factor) = 9.730								
Average GRS (6-factor) = 9.206								

Table 5.
Fama-Macbeth regressions

This table shows results from Fama-Macbeth regressions of weekly stock returns of firms on lagged stock returns of neighbors. At the beginning of each formation week w , for each firm in the network, I randomly select a neighbor and match it with the firm. Thereafter, I track neighbors' returns over the formation week w , skip a week, and track neighbors' returns over week $w+2$. Column (1) reports results from a univariate regression of firms' returns on neighbors' returns. I augment the regression model with control variables in Column (2). Column (3) contains results from the 1996 – 2005 subsample while Column (4) reports results from the 2006 – 2015 period. Firm size is defined as the market capitalization at the time of portfolio formation. Book equity (b) is defined as the Compustat book value of stockholders' equity, plus balance sheet deferred taxes and investment tax credit (where available), less the book value of preferred stock. Depending on their availability, redemption, liquidation, or par value (in this order of priority) is used to estimate the book value of preferred stock. To ensure that these variables are known at portfolio formation, I match the annual book equity value for fiscal years ending in calendar year $y-1$ with portfolio formation dates starting in calendar year y . Book equity is then scaled by market capitalization at portfolio formation to form b/m. Stock returns (-1, 0) is the firm's stock returns over the prior one week. Stock returns (-24, 0) is the firm's stock returns over the prior 24 weeks. Following Campbell et al. (2001), I define idiosyncratic volatility of a firm as the summation of squared differences between daily returns and daily Fama-French 49 industry returns over the prior four weeks. Coefficients presented in the table are averages over 100 rounds of simulations and are inflated by a factor of 10^3 . Average Newey t-statistics reported in parentheses. ***, **, and * denote statistical significance of the Fisher's chi-squared test at the 1%, 5%, and 10% respectively.

Dependent variable: Firm returns				
	(1)	(2)	(3)	(4)
Sample period	Full sample	Full sample	1996 – 2005	2006 – 2015
Neighbor returns	5.164*** (3.80)	4.681*** (4.27)	4.821*** (3.22)	4.541*** (2.85)
log (Firm size)		-0.136 (1.45)	-0.295*** (1.91)	0.022 (0.21)
Book-to-market		-0.004 (0.74)	-0.007 (1.07)	-0.001 (0.09)
Stock returns (-1, 0)		-0.222 (0.09)	-1.548 (0.49)	1.102 (0.26)
Stock returns (-24, 0)		2.240*** (2.52)	3.479*** (2.73)	1.004 (0.82)
Idiosyncratic volatility		-11.341*** (2.26)	-13.681 (1.59)	-9.004 (1.73)
Number of cross-sections	1,011	1,011	506	505

Table 6.
Imbalances in information production and gradual information diffusion

This table shows weekly equal-weighted average returns to portfolios independently triple-sorted on 1) lagged stock returns of neighbors, 2) firm size, and 3) the ratio of neighbors' revision rate to firms' revision rate (revision rate ratio). Sizes of firms and neighbors are defined as their market capitalization at the time of portfolio formation. Revision rate of a firm (neighbor) in year y is defined as the number of revisions per analyst covering the firm (neighbor) in the year $y-1$. At the beginning of each formation week w , for each firm in the network, I randomly match each firm with one of its neighbors. Thereafter, I track neighbors' returns over the formation week, and form quintile portfolios of firms independently triple-sorted on their neighbors' returns, firm size, and revision rate ratio. I skip a week after the formation week and hold the portfolios for one week. I skip a week after the formation week and hold the portfolios for one week (in $w+2$). Statistics presented in the table are averages over 100 rounds of simulations.

	Firm size Revision rate	Small			Med			Big		
		Low	Med	High	Low	Med	High	Low	Med	High
Neighbor- returns		Mean weekly returns								
	Low	0.230	0.225	0.214	0.247	0.245	0.214	0.244	0.223	0.215
	Medium	0.284	0.290	0.284	0.288	0.282	0.268	0.268	0.259	0.250
	High	0.320	0.328	0.321	0.309	0.300	0.288	0.261	0.256	0.247
	High minus Low	0.090 ^{***}	0.103 ^{***}	0.107 ^{***}	0.062 ^{***}	0.055 ^{***}	0.074 ^{***}	0.017	0.033 [*]	0.033 ^{**}
	Fisher $\chi^2 (\times 10^3)$	0.921	1.155	1.460	0.557	0.430	0.712	0.145	0.232	0.245
		Number of firms								
	Low	96	99	133	103	108	109	121	111	88
	Medium	89	93	125	102	108	108	129	120	94
	High	96	100	136	105	112	113	126	117	93
		Neighbors' returns								
	Low	-0.055	-0.054	-0.054	-0.053	-0.052	-0.052	-0.051	-0.050	-0.050
	Medium	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	High	0.062	0.061	0.060	0.060	0.058	0.057	0.057	0.056	0.055
		Revision rate ratio								
	Low	0.636	1.031	1.973	0.644	1.030	1.795	0.637	1.026	1.716
	Medium	0.641	1.031	1.962	0.648	1.030	1.789	0.640	1.026	1.708
	High	0.638	1.031	1.974	0.645	1.030	1.795	0.639	1.026	1.716
		Firm size								
	Low	203.4	209.1	193.1	900.6	894.1	875.2	14625.8	12472.9	11088.4
	Medium	205.7	210.7	195.5	911.1	902.9	881.0	14498.4	12467.1	11311.2
	High	204.6	209.8	194.1	902.3	896.4	876.5	14636.2	12471.5	11127.9

Table 7.
Adjustments for industry effects

This table shows weekly equal-weighted average returns to portfolio strategies that adjust for industry effects. At the beginning of each formation week, for each firm in the network, I randomly select a neighbor and match it with the firm. Thereafter, I track neighbors' returns over the formation week w , and form quintile portfolios of firms sorted on their neighbors' returns. I skip a week after the formation week and hold the portfolios for one week (in $w+2$). Columns (1) through (3) employ the two-digit SIC industry classification scheme. In Column (1), conditional on the industry of each firm, the algorithm is constrained to match the firm with a neighbor from a different industry. Thereafter, the neighbors' returns are sorted globally. In Column (2), conditional on the industry of each firm, the algorithm is constrained to match the firm with a neighbor from the same industry. Thereafter, the neighbors' returns are sorted within each industry and the portfolio returns are averaged across industries. In Column (3), conditional on the industry of each firm, the algorithm is constrained to match the firm with a neighbor from the same industry. Thereafter, the neighbors' returns are sorted globally; firm returns are industry-adjusted by subtracting the contemporaneous industry average returns from it. Columns (4) through (6) repeat the portfolio strategies of the preceding columns in the same sequence but employ the Fama-French 49 industry classification scheme. Statistics presented in the table are averages over 100 rounds of simulations.

	(1)	(2)	(3)	(4)	(5)	(6)
	2-digit SIC industries			FF49 industries		
Adjustment methods	(a)	(b)	(c)	(a)	(b)	(c)
Neighbor-returns	Mean weekly returns					
Low	0.225	0.230	-0.195	0.225	0.233	-0.191
N2	0.264	0.259	-0.130	0.265	0.258	-0.129
N3	0.274	0.269	-0.117	0.273	0.267	-0.119
N4	0.279	0.277	-0.115	0.283	0.277	-0.115
High	0.285	0.287	-0.104	0.287	0.291	-0.100
High minus Low	0.060***	0.057***	0.091***	0.062***	0.058***	0.092***
Fisher χ^2 ($\times 10^3$)	0.782	1.233	2.128	0.841	1.134	1.901

Adjustment methods

- (a) Firms are matched with neighbors in different industries; global sorts; raw returns
- (b) Firms are matched with neighbors in same industries; within-industry sorts; raw returns
- (c) Firms are matched with neighbors in same industries; global sorts; industry-adjusted returns

Table 8.
Vector autoregressions of firms' returns and neighbors' returns

This table shows results from vector autoregressions (VARs) of weekly stock returns of firms and their neighbors. At the beginning of each formation week, for each firm in the network, I randomly select a neighbor and match it with the firm. Firms' returns and neighbors' returns in week w are denoted by $\text{Ret}_{F,w}$ and $\text{Ret}_{N,w}$ respectively. To facilitate presentation, estimated loadings on Ret_F and Ret_N are denoted as a and b , respectively. The subscripts on loadings a and b represent the equations to which they belong (e.g., a_N is the loading of Ret_F in the equation with Ret_N on the LHS). $F(a_N = b_F)$ is the F-statistic of the cross-equation test of difference between a_N and b_F . Statistics presented in the table are averages over 100 rounds of simulations. In Panel A, I estimate VARs on subsamples split on firm size. In Panel B, I estimate VARs on subsamples split on neighbor size. ***, **, and * denote statistical significance of the Fisher's chi-squared test at the 1%, 5%, and 10% respectively.

Panel A: VARs on firm size quintiles

Firm size	LHS	$\text{Ret}_{F,w-1}$ (loading: a)	$\text{Ret}_{N,w-1}$ (loading: b)	Average R^2	$F(a_N = b_F)$
Small	$\text{Ret}_{F,w}$	0.083***	0.084***	1.72	1.09
	$\text{Ret}_{N,w}$	0.131***	0.009	1.31	
Q2	$\text{Ret}_{F,w}$	-0.019	0.090***	0.78	3.34***
	$\text{Ret}_{N,w}$	0.009	0.007	0.02	
Q3	$\text{Ret}_{F,w}$	-0.038	0.085***	0.81	5.33***
	$\text{Ret}_{N,w}$	-0.018	0.001	0.04	
Q4	$\text{Ret}_{F,w}$	-0.054	0.073***	0.85	6.99***
	$\text{Ret}_{N,w}$	-0.044	-0.006	0.19	
Big	$\text{Ret}_{F,w}$	-0.069***	0.053***	0.83	8.35***
	$\text{Ret}_{N,w}$	-0.076***	-0.013	0.50	

Panel B: VARs on neighbor size quintiles

Neighbor size	LHS	$\text{Ret}_{F,w-1}$ (loading: a)	$\text{Ret}_{N,w-1}$ (loading: b)	Average R^2	$F(a_N = b_F)$
Small	$\text{Ret}_{F,w}$	-0.002	0.074***	0.78	2.40**
	$\text{Ret}_{N,w}$	0.004	0.146***	2.15	
Q2	$\text{Ret}_{F,w}$	-0.019	0.063***	0.58	2.80***
	$\text{Ret}_{N,w}$	-0.013	-0.005	0.02	
Q3	$\text{Ret}_{F,w}$	-0.023	0.074***	0.72	3.12***
	$\text{Ret}_{N,w}$	-0.006	-0.033	0.12	
Q4	$\text{Ret}_{F,w}$	-0.026	0.075***	0.65	3.28***
	$\text{Ret}_{N,w}$	-0.006	-0.042	0.19	
Big	$\text{Ret}_{F,w}$	-0.026	0.072***	0.55	3.11***
	$\text{Ret}_{N,w}$	-0.006	-0.071***	0.51	

Table 9.
Propagation of acquisition activities in network

This table shows results from logit and conditional logit models. The dependent variable is *firm acquisition* which is an indicator that equals unity if the firm completes at least one acquisition in the current month, and equals zero otherwise. The key independent variables are *neighbor acquisition* and *pseudo-neighbor acquisition*. The indicator *neighbor acquisition* equals unity if the neighbor completes at least one acquisition in the previous month, and equals zero otherwise. The indicator *pseudo-neighbor acquisition* is similarly constructed for pseudo-neighbors that are randomly matched to the firm. Columns (1) and (2) reports results from estimation of logit specifications. Columns (3) to (5) reports results from estimation of conditional logit specifications. Full details on the construction of other variables are available in Section 3.2 and the appendix. Estimated coefficients are presented in the form of odds ratios and are presented as averages over 100 rounds of simulations. The Fisher's chi-squared values corresponding to the coefficients are deflated by a factor of 10^3 and reported in parentheses. ***, **, and * denote statistical significance of the Fisher's chi-squared test at the 1%, 5%, and 10% respectively.

Dependent variable: Firm acquisition indicator					
	(1) LOGIT	(2) LOGIT	(3) C-LOGIT	(4) C-LOGIT	(5) C-LOGIT
Neighbor acquisition	1.357*** (4.649)	1.155*** (1.255)	1.121*** (0.893)	1.108*** (0.724)	
Pseudo-neighbor acquisition					0.998 (0.181)
Firm market capitalization	1.000*** (28.286)	1.000*** (24.880)	1.000*** (20.164)	1.000 (0.200)	1.000*** (20.703)
Firm returns (-12, -2)	1.111*** (13.436)	1.096*** (9.250)	1.097*** (3.290)	1.097*** (3.121)	1.100*** (3.484)
Firm book-to-market	0.821*** (27.661)	0.850*** (4.335)	0.827*** (4.558)	0.536*** (1.451)	0.830*** (4.634)
Firm cash flow	1.000*** (4.801)	1.000*** (3.358)	1.000*** (2.303)	1.000*** (0.780)	1.000*** (2.413)
FF49 indicators	No	Yes	No	No	No
Year indicators	No	Yes	No	No	No
FF49-year fixed effects	No	No	Yes	No	Yes
Firm fixed effects	No	No	No	Yes	No
S.E. cluster	None	FF49-year	Year	FF49	Year