Do Equity Analysts Learn from Their Colleagues? Evidence Using an Information Network Centrality Measure*

J. KENNY PHUA *, T. MANDY THAM, and CHISHEN WEI

June 18, 2016

ABSTRACT

We examine how peer learning affects analyst forecast outcomes using network theory. We construct a measure of information network centrality among analysts within a brokerage using the degree of sector overlaps in their coverage portfolios. We find that analysts with higher centrality scores produce more accurate forecast estimates, exhibit less pronounced herding behavior, and experience stronger market reactions to their forecast revisions. Consistent with a peer-learning channel, high centrality analysts are also more likely to incorporate information from their colleagues' recent forecast errors into their forecast revisions. Overall, our evidence suggests that peer-learning within a brokerage is an important information acquisition channel for financial analysts.

JEL Classification Code: D83, G17, G24

Keywords: Analysts, Forecast Accuracy, Networks, Learning

* Corresponding author: J. Kenny Phua (JPHUA005@e.ntu.edu.sg). Phua and Wei are at Nanyang Business School, Nanyang Technological University, Singapore 639798. Tham is at Wealth Management Institute (Singapore). We are grateful for helpful comments from Rui Shen, Avanidhar Subrahmanyam, and seminar participants at Nanyang Business School. All errors are our own. We acknowledge research support from Nanyang Technological University.

Do Equity Analysts Learn from Their Colleagues? Evidence Using an Information Network Centrality Measure

ABSTRACT

We examine how peer learning affects analyst forecast outcomes using network theory. We construct a measure of information network centrality among analysts within a brokerage using the degree of sector overlaps in their coverage portfolios. We find that analysts with higher centrality scores produce more accurate forecast estimates, exhibit less pronounced herding behavior, and experience stronger market reactions to their forecast revisions. Consistent with a peer-learning channel, high centrality analysts are also more likely to incorporate information from their colleagues' recent forecast errors into their forecast revisions. Overall, our evidence suggests that peer-learning within a brokerage is an important information acquisition channel for financial analysts.

How sell-side equity research analysts generate their forecast estimates remains an open and important topic of research (e.g., Bradshaw, 2011). Extant studies find that an analyst's individual characteristics, such as work experience and ability, and personal and professional connections, such as access to management, determine her forecast accuracy.[1] However, even the best research analysts do not work in isolation. Their brokerage houses provide not only operational and back-office resources but also co-workers, who may provide a potentially valuable network of knowledge and information.

We propose that analysts who can better tap into the expertise of their in-house colleagues produce higher quality equity research. For example, an analyst covering Google may provide useful industry insights to a colleague covering Apple. This type of information sharing is potentially valuable because extant studies find that industry knowledge is an important factor in producing better forecast estimates (e.g., Clement, Koonce, and Lopez, 2007; Hilary and Shen, 2013). As conversations and workplace interactions facilitate the sharing of valuable knowledge, both analysts may be able to make better forecasts. Lehman Brother's research department in the early 1990s is an example of the importance of peer influence. To foster peer learning, Lehman Brothers instituted a policy that every analyst's presentation must refer to at least two colleagues. During that period, Lehman Brothers was regularly ranked among the top brokerage firms.[2]

While such anecdotes emphasize the importance of collaboration, measuring information flows between colleagues is a challenging task because workplace interactions are unobservable. To test our hypothesis, we adopt a network theory approach using analysts as nodes within the information network of their brokerage. 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. In contrast, an analyst in a peripheral nodal position is less well-positioned to benefit from in-house information exchange and has a low *Analyst Centrality* score. To construct network connections among analysts in a brokerage, we track the extent of sector overlaps among their coverage portfolios.[3] Our

[1] See for example Clement (1999), Jacob, Lys, and Neale (1999), Chen and Matsumoto (2006), Clement, Koonce, and Lopez (2007), Cohen, Frazzini, and Malloy (2010), Hilary and Shen (2013), and Soltes (2014).
[2] "The Risky Business of Hiring Stars," Harvard Business Review, May 2004
[3] Sectors are classified according to the 2-digit Global Industrial Classification Standard (GICS) codes.

premise is that information exchange is more likely to occur between a pair of analysts if there is sector overlap in their coverage portfolios. By this construction, an analyst with sector overlaps with a larger number of colleagues will shift towards the epicenter of the information network of her brokerage and will score higher on our *Analyst Centrality* measure.

Our tests indicate that analysts with higher *Analyst Centrality* scores produce more accurate forecasts. Economically, a standard deviation increase in *Analyst Centrality* is associated with a 7.5% reduction in absolute forecast error relative to the median. Our tests include firm-year fixed effects to ensure that our findings are not driven by unobserved firm heterogeneity. Our findings are also unlikely due to variation across brokerages as our tests also include brokerage-year fixed effects.

This evidence is consistent with the view that analysts with higher *Analyst Centrality* scores have better access to information. We also expect that better access to information translates into more informative forecast revisions. Consistent with this hypothesis, we find that analysts with higher *Analyst Centrality* scores are less likely to issue herding forecast revisions. Their forecast revisions command stronger market reactions, especially when the forecast revisions diverge from their prior forecast values or diverge from the prevailing forecast consensus. Holding consensus deviation at its sample mean (median), an increase in *Analyst Centrality* by one standard deviation is associated with a 4.86% (1.61%) increase in annualized market-adjusted abnormal returns.

While our findings indicate that analyst centrality influences forecast outcomes, they do not necessarily imply that analysts learn from their peers. High centrality may simply reflect an analyst's innate ability that is incremental to other factors discussed in the prior literature. Therefore, we design a specific test to distinguish the learning hypothesis from the innate ability hypothesis. We argue that learning is likely to occur when analysts observe the *ex-post* forecast errors of their colleagues. In response, they may revise their own forecasts to reverse the errors they observe in their colleagues' forecasts.

Our findings suggest that analysts with higher centrality scores are more likely to issue revisions that reverse their colleagues' forecast errors.[4] Moreover, we find that their adjustments only occur when their colleagues' forecast are revealed to be pessimistic. We find no evidence that analysts revise their forecasts when their colleagues' forecasts are revealed to be optimistic. This pattern of asymmetric learning is consistent with the management-relations cultivation hypothesis (e.g., Francis and Philbrick, 1993; Das et al., 1998; Lim, 2001; Matsumoto, 2002; Richardson et al., 2004; Ke and Yu, 2006). Together, this evidence supports the learning hypothesis. Using analyst membership in the Institutional Investor All-America Research Team, we show that the effect of *Analyst Centrality* on superior forecasting outcomes is incremental to innate analyst ability. However, we note that this is not necessarily inconsistent with the notion that high centrality analysts also possess higher abilities.

One concern with our *Analyst Centrality* measure is that it may be capturing unobservable brokerage effects. While our main tests include brokerage fixed effects, this does not directly address concerns on the matching process between analyst and brokerage house. Therefore, *Analyst Centrality* may be capturing unobservable brokerage characteristics or analyst ability. For example, an analyst with expertise in a particular industry may be more likely to work at a brokerage that specializes in the same industry. We address this particular concern by excluding brokerage houses that cover less than three industries and find that our results are unchanged.

To establish a causal relation between *Analyst Centrality* and forecast accuracy, we exploit employment shocks using a sample of brokerage closures (e.g., Kelly and Ljungqvist, 2012) from years 2000 to 2007. These closures exogenously forced analysts to join new brokerage houses.[5] We find that analysts who exogenously experienced improvements in their *Analyst Centrality* scores had superior forecast accuracy relative to those whose *Analyst Centrality* scores declined. Using a generalized difference-in-difference model that

---

[4] We expect that analysts who benefit more from inter-colleague information exchange also incorporate information, learnt from their colleagues, into their own forecasts to a greater extent.

[5] Derrien and Kecskes (2013) find that most brokerage closures stem from breakdowns in business strategy. This implies that the closures of brokerages are unlikely to be related to the forecasting abilities of the affected analysts.

accommodates multiple treatment groups and multiple treatments across time (Autor, 2003), we find that our results are not driven by anticipatory effects of the treatment. Instead, differences in forecast accuracy are found primarily in the post-treatment periods.

Our paper contributes to several strands of literature. First, we add to a growing literature that seeks to penetrate the 'black box' of information generation of sell-side financial analysts (e.g., Bradshaw, 2011; Brown, Call, Clement, and Sharp, 2015). Using tools from network theory, we develop a measure *Analyst Centrality* that captures the propensity of an analyst to participate in information exchange with her brokerage colleagues. We show that within-brokerage information exchange among colleagues is an important information acquisition channel.

Second, we contribute to a broader literature on learning in financial markets. Exploiting the educational links between analysts and senior corporate officers, Cohen, Frazzini, and Malloy (2010) find that analysts learn superior firm-specific information in social networks. To our knowledge, we are the first to examine the impact of inter-colleague learning within brokerage networks on the production of sell-side equity research.

Third, this study extends our understanding of the determinants of analyst forecast accuracy. Mikhail et al. (1997) and Clement (1999) show that firm-specific experience improves forecast accuracy. Clement et al. (2007) find that task-specific experience improves forecast accuracy, and that such experience can extend to other firms under the analyst's coverage (Hilary and Shen, 2013). The common thread of these studies is the focus on individual-level learning. We complement these earlier findings by showing that analyst-learning at the brokerage-level also has an impact on forecasting performance.

Finally, our study adds a new perspective of coverage portfolio complexity on analyst forecast accuracy. In a setting where analysts face resource constraints and diminishing returns to effort, Clement (1999) finds that forecast accuracy of an analyst declines as the numbers of industries and firms under her coverage increase. This view is echoed by Lees (1981) who concludes that there are economies of scale in acquiring information of other firms in the same industry. However, if we model brokerages as networks in which information exchange among analysts is possible, an analyst's coverage portfolio complexity may connect her to more colleagues and may heighten her propensity of benefitting from peer-learning. To be clear, we do not dispute the 'busy analyst' hypothesis.[6] Instead, we conclude that there may be a bright side to coverage portfolio complexity in a network where the information held by individual analysts is complementary to one another.

1. Sample and Data

In this section, we describe our methodology and construction of our variables. We also discuss our data and present summary statistics on our sample.

1.1. Defining Network Connections

Prior research has shown that analysts covering similar industries benefit from the firm-specific information of one another (Clement, Koonce, and Lopez, 2007; Hilary and Shen, 2013). Therefore, we hypothesize that information transfer is more likely to occur between a pair of analysts if there is sector overlap in their coverage portfolios. We construct within-brokerage network connections based on the forecast data for fiscal year one from the Detailed History file of I/B/E/S. We use the Global Industrial Classification Standard (GICS) codes to classify industries due to its popularity among financial practitioners (Bhojraj, Lee, and Oler, 2013). Brokerage-years with fewer than 5 analysts are dropped.

Our methodology allows the structures of within-brokerage networks to change with time. Let us illustrate with two analysts, A and B, who are colleagues in the same brokerage. A network connection exists between A and B if each analyst makes at least one forecast announcement in a common 2-digit GICS sector within the same year. For example, a connection is present in year t if A makes forecasts in GICS sectors 20 and 45, and B makes forecasts in GICS sector 45. Connections governed by a greater degree of industry overlaps between the coverage portfolios of the two analysts are given a larger weight. Suppose that B has instead made forecasts in GICS sector 20 and 45, the connection weight will then increase from one to two. Notably, by our construction, the connection weight is

[6] Our empirical results do not dispute the 'busy analyst' hypothesis. In our baseline regressions, the number of unique sectors covered by an analyst is negatively associated with her forecast accuracy.

independent of the number of unique firms covered in the overlapping GICS sectors by either analyst. While it is plausible that information exchange between the connected analysts intensifies with the number of firms in their overlapping sectors, our approach is arguably more conservative and parsimonious. In our sample, 93.46%, 5.81%, and 0.63% of connections have a weight of 1, 2, and 3 respectively. The maximum connection weight is 7.

We provide a graphical example of a within-brokerage network constructed with our methodology in Figure 1. Each red node represents an individual analyst. If there is a connection between a pair of analysts, a black line links the two corresponding red nodes. For ease of presentation, connection weights are not reflected in the diagram.[7] From Figure 1, we also observe that the numbers of connections possessed by analysts are highly heterogeneous, as is typical of the networks in our sample. We single out Analyst A (circled in blue) and Analyst B (circled in green) in the example. Analyst A is only connected to two other colleagues in the brokerage while Analyst B possesses 17 connections.

[Insert Figure 1]

We do not claim that within-brokerage information transfer among colleagues occurs exclusively along our defined connections. Rather, we are merely proposing that there is a higher likelihood of information exchange between two analysts if the two are at least connected in the manner as defined by our methodology.

1.2. Defining Analyst Network Centrality

In this section, we construct our *Analyst Network Centrality* measure to proxy for the propensity of information exchange between analysts. First, we treat analysts as nodes in their brokerage's information network. Under network theory, each nodal position reflects a unique set of benefit-constraint tradeoff with social power of the node strengthening with benefits and weakening with constraints. In other words, the benefit-constraint tradeoff at one node would be different from that at a second node and the two nodes would command different social power. A node with higher social power is more prominent in the

[7] It is common to represent the connection weight between a pair of nodes by the thickness of the line connecting them. Thicker lines can be used to represent a heavier connection weight.

network. Broadly, social power is synonymous with the connectedness of the node or in our context, the centrality of the node. An analyst at a node with high social power would have high connectivity and hence high centrality by our definition. By virtue of their prominence in the network, we hypothesize that high centrality analysts tend to be beneficiaries of information transfer and would therefore display superior forecasting performance.

To test our hypothesis, we employ four commonly used measures of network centrality – degree, closeness, betweenness, and eigenvector centrality (see for example, Larcker, So, and Wang (2013)). For each measure, we provide a mathematical and conceptual definition below. Details and examples of each centrality measure are provided in Appendix II.

1.2.1. Degree Centrality

First, an analyst may have more opportunities to participate in information transfer if she possesses more channels for inter-colleague interaction. This idea is captured by the *Degree Centrality* variable which measures the number of direct connections to other colleagues. Let $f(i,\ j)$ be an indicator which equates to unity if analysts i and j share a direct connection, and zero otherwise. We define *Degree Centrality* of analyst i in her brokerage network G as in (1).

$$Centrality_{Degree}(i,\ G) = \frac{\sum_{j \neq i, j \in G}^{N_G - 1} f(i,\ j)}{N_G - 1} \qquad (1)$$

where $N_G - 1$ is the maximum number of direct connections analyst i can have and is applied as a normalization factor in the denominator because *Degree Centrality* increases with the number of analysts in the brokerage.

1.2.2. Closeness Centrality

Second, an analyst may be more likely to benefit from information transfer if she can access her colleagues or is accessible by her colleagues at relatively shorter path lengths.[8] This idea is captured by the *Closeness Centrality* variable, which is the reciprocal of the sum of shortest-path distances to all other brokerage colleagues. Assuming that the path length between any two analysts in a brokerage is proportional to the cost of information exchange between them, analysts with higher values of *Closeness Centrality* would have lower costs of interaction with their colleagues. Let $d(i,\ j)$ be the shortest-path distance between analysts i and j. We define *Closeness Centrality* of analyst i in her brokerage network G following (2).

$$Centrality_{Closeness}(i,\ G) = \frac{N_G - 1}{\sum_{j \neq i, j \in G}^{N_G - 1} d(i,j)} \qquad (2)$$

where $N_G - 1$ is the minimum sum of shortest-path distances that analyst i can have and is applied as a scaling factor in the numerator because *Closeness Centrality* decreases with the number of analysts in the brokerage.

1.2.3. Betweenness Centrality

Third, information is more likely to pass through an analyst if she is positioned on many network paths between her brokerage colleagues. Such analysts are valuable as key brokers of information exchange and other inter-colleague interactions in the network. This idea is captured by the *Betweenness Centrality* variable, which is the summed proportions of all shortest-length paths (known as geodesics[9] in network theory) passing through the analyst. Since *Betweenness Centrality* and *Closeness Centrality* are both built on the notion of geodesics, analysts with higher values of *Betweenness Centrality* also tend to have lower costs of information exchange, by virtue of their brokering capacities. Let $s(x,\ y)$ be the number of geodesics between any analyst-pair x and y, and $s(x,\ y|i)$ be the number of geodesics between any analyst-pair x and y that passes through analyst i. We define *Betweenness Centrality* of analyst i in her brokerage network G following (3).

[8] Detailed explanations and examples of path lengths are contained in Appendix II.
[9] A working example of geodesics is available in Appendix II.

$$Centrality_{Betweenness}(i,\ G) = \frac{2}{(AP_G - 1)(AP_G - 2)} \sum_{x \neq i, y \neq i, x, y \in G} \frac{s(x, y|\ i)}{s(x, y)} \qquad (3)$$

where AP_G is the number of unique analyst-pairs not involving analyst i. A normalization factor $\frac{2}{(AP_G-1)(AP_G-2)}$ is applied in the denominator because *Betweenness Centrality* increases with the number of analysts in the brokerage.

1.2.4. Eigenvector Centrality

Lastly, not all connections are equal. An analyst may be prominent in her brokerage network only if she is connected to other well-connected analysts. This is encapsulated in the *Eigenvector Centrality* variable. While *Degree Centrality* incorporates the number of colleagues connected directly to an analyst, it does not consider the quality of these connections. Analysts with high values of *Eigenvector Centrality* tend to be connected to other colleagues who also possess high values of *Eigenvector Centrality.*[10] For her brokerage network G, letting $\mathbf{V_G}$ be the eigenvector of the network's adjacency matrix $\mathbf{M_G}$.[11], we define the *Eigenvector Centrality* of analyst i as in (4).

$$Centrality_{Eigenvector}(i,\ G) = element\ v_i\ of\ matrix\ \mathbf{V_G} \qquad (4)$$

A notable variant of *Eigenvector Centrality* is the PageRank algorithm used by Google to rank the importance of websites on the Internet. The underlying logic of the algorithm is that important websites are also likely to receive web-links from other important websites. In their study on boardroom centrality, Larcker, So, and Wang (2013) employ *Eigenvector Centrality* as a measure of prestige and power that confers special privileges to boards in

[10] Appendix II contains the mathematical intuition behind *Eigenvector Centrality* for the interested reader.
[11] A working example of an adjacency matrix is available in Appendix II. An adjacency matrix is a symmetric matrix that describes the connections among all nodes in a network. Consider a network of 3 analysts *X*, *Y*, and *Z*; connection weights are all set to unity for ease of discussion. The connections in the network are as follows. *X* is connected to *Y*, but is not connected to *Z*. *Y* is connected to *Z*. Since there are 3 nodes in the network, the adjacency matrix is a 3 x 3 matrix with all diagonal elements equating to zero (we assume no self-loops in our network) – $\begin{pmatrix} 0_{XX} & 1_{XY} & 0_{XZ} \\ 1_{YX} & 0_{YY} & 1_{YZ} \\ 0_{ZX} & 1_{ZY} & 0_{ZZ} \end{pmatrix}$. Each element in the adjacency matrix indicates whether a connection exists between a pair of analysts (denoted by the subscripts). If a connection is present between a pair of analysts, their corresponding element equates to unity, and zero otherwise.

obtaining information and favors. In the context of information transfer, analysts with high values of *Eigenvector Centrality* can be interpreted as better aggregator of widely-held information in the network. The intuition behind this interpretation is that the information arriving at such an analyst is likely to originate from colleagues who are in turn connected to many others in the network.

1.2.5. Principal Component Analysis

In spite of the four measures, there is no single measure that can describe the benefits and constraints of each node consummately because social power in a network is a multi-faceted construct (Newman, 2003). A particular measure may be optimal in describing the centrality of some nodes in the network but may be sub-optimal in describing the remaining nodes. This does not necessarily mean that the remaining nodes are unimportant in the network but rather that the measure is unable to capture their network prominence adequately. Generally, it is common for a given node in a network to be favored by one measure, and less so by another measure.

Since it is unclear, *ex-ante*, which measure can describe the network centralities of analysts optimally, we perform a principal component analysis (PCA) on the four network centrality measures. From the first principal component, we extract the standardized factor score and define it as our *Analyst Centrality* measure. The factor loadings on all four centrality measures are positive, supporting our use of *Analyst Centrality* as a measure of aggregate analyst connectedness within their brokerage networks. Consistent with prior literature, we also find that there is substantial correlation[12] among the four centrality measures – ranging from 28.2% to 89.1%. In view of this, our use of a PCA-extracted factor score to represent analyst connectedness helps us to avoid potential multicollinearity concerns in our empirical analyses.

1.3. Discussion of Analyst Network Centrality

[12] The substantial correlation among the four centrality measures is unsurprising because each centrality measure is merely capturing a distinct component of overall connectedness. A well-connected analyst would have relatively high scores across multiple centrality measures. For example, an analyst positioned at the center of a star-shaped network would be endowed with high values of degree, closeness, and betweenness centrality by definition.

Our modeling of information transfer within brokerages rests on a few salient assumptions. First, we assume that inter-colleague information transfer occurs primarily through sector overlaps in analysts' coverage portfolios. This is motivated by Brown, Call, Clement, and Sharp (2015) who wrote that "industry knowledge is the single most useful input to analysts' earnings forecasts and stock recommendations". It is plausible that a substantial portion of inter-colleague communication (e.g. informal interaction) occurs outside of our constructed networks. Furthermore, if analyst career progressions are structured as competitive tournaments (Yin and Zhang, 2015), it remains an open question whether analysts will cooperate with their colleagues and share information. Second, we assume that information tends to travel via the shortest possible paths within a brokerage network. This assumption is implicit in the *Closeness Centrality* and *Betweenness Centrality measures*, which are both built on the notion of geodesics. Where informal ties are dominant, geodesics may cease to be the least costly paths for information transfer and information may flow via longer paths in actuality. If the above considerations are true, the centrality measures may fail to describe the true network centralities of analysts. However, this mismeasurement of centrality should work against us in finding support for our hypothesis.

Third, there are alternative interpretations of a causal relation between *Analyst Centrality* and forecasting performance. Prior research finds that brokerage size, as a proxy for resources and prestige, can predict analyst forecasting performance (Clement, 1999). Assuming that portfolio coverage decisions are randomly made, the probability of an analyst gaining an incremental connection should increase with brokerage (network) size. Furthermore, if we view the analyst labor market as a tournament where performance determines continuation in the profession (Clement, 1999), we expect that higher-ability analysts tend to be employed longer. Consequently, the longer tenures confer higher-ability analysts with more opportunities to develop within-brokerage connections. This would then lead to higher *Analyst Centrality* values. In sum, the effects of *Analyst Centrality* on forecasting performance may be driven by unobserved heterogeneity and brokerage characteristics. We address these concerns as follows. First, we check that the univariate correlation between *Analyst Centrality* and *Brokerage Size* (*General Experience*) is low, at about 5.2% (2.0%). Next, we use normalized centrality measures that adjust for network size. Third, we address the influence of innate analyst ability on our findings by including an *All-Star* indicator alongside *Analyst Centrality* in horse-race regressions. Finally, we also include brokerage size, general experience, and brokerage fixed effects in our regression analyses.

1.4. Forecast Accuracy

We construct three measures of forecasting accuracy from the latest firm-year forecast values of analysts. First, *Forecast Error* is the absolute difference between the analyst's firm-year forecast value and the corresponding earnings-per-share (EPS) of the firm-year. Second, for comparability across observations, we construct *Normalized Forecast Error* in percentage points as *Forecast Error* scaled by the firm-year mean forecast error. Third, we use *Clement-Tse Accuracy* as an additional normalized measure of forecasting accuracy (Clement and Tse, 2005). In each firm-year, we first compute the maximum (*Max Forecast Error*) and minimum (*Min Forecast Error*) values of *Forecast Error*. Thereafter, the *Clement-Tse Accuracy* of an analyst in a particular firm-year is the ratio of the difference between *Max Forecast Error* and her *Forecast Error* to the difference between *Max Forecast Error* and *Min Forecast Error*. Hence, *Clement-Tse Accuracy* is bounded between zero and unity, with higher values reflecting a greater degree of forecasting accuracy. Formally, for analyst i covering firm f in year t, we define her *Clement-Tse Accuracy* as in (5).

$$Clement - Tse\ Accuracy_{i,f,t} = \frac{Max.Forecast\ Error_{f,t} - Forecast\ Error_{i,f,t}}{Max.Forecast\ Error_{f,t} - \ Min.Forecast\ Error_{f,t}} \qquad (5)$$

1.5. Herding Behavior

We adopt the methodology of Clement and Tse (2005) in defining herding behavior. For each firm-year forecast revision made by a given analyst, we compute the interim forecast consensus. This interim forecast consensus excludes the forecast contribution of the given analyst and only includes the most recent forecasts of other analysts. If the analyst's revision value is above her prior forecast and above the interim forecast consensus, we classify that forecast revision as *innovative*. A forecast revision is also classified as *innovative* if it is below the analyst's prior forecast and below the interim forecast consensus. If the prior forecast or interim forecast consensus is unavailable, we classify the forecast revision as *innovative*. All other forecast revisions are classified as *herding*. We also compute *Herding Rate*, defined as the number of herding forecast revisions made by an analyst in a firm-year, scaled by her total number of forecast revisions in the same firm-year. Hence, *Herding Rate* is bound between zero and unity, with higher values reflecting stronger herding behavior. Given an analyst i covering firm f and making p forecast revisions in year t, we let $h(i, f, t)$ equate to unity if the forecast revision is either above or below both the analyst's prior forecast value and the interim forecast consensus, and zero otherwise.

$$Herding\ Rate_{i,f,t} = \sum_{1}^{p} \frac{h(i, f, t)}{p} \qquad (6)$$

We also compute the *Clement-Tse Herding Rate* in a similar manner to the *Clement-Tse Accuracy* measure. In each firm-year, we first compute the maximum (Max HR) and minimum (Min HR) values of *Herding Rate*. Thereafter, the *Clement-Tse Herding Rate* of an analyst in a firm-year is defined as the ratio of the difference between Max HR and *Herding Rate* to the difference between Max HR and Min HR. Hence, *Clement-Tse Herding Rate* is bounded between zero and unity, with higher values reflecting stronger herding behavior.

1.6. Control Variables

We control for various analyst, forecast and firm characteristics that are found to affect forecasting performance in prior research.[13] *Revision Frequency* is the total number of forecast revisions made by the analyst in a particular firm-year. *Horizon* is the number of days that has elapsed between the analyst's firm-year forecast and the actual earnings announcement. Following standard practice in the literature, we remove forecasts issued either more than 365 days or less than 30 days from the actual earnings announcement

[13] See for example, Mikhail, Walther, and Willis (1997); Clement (1999); Jacob, Lys, and Neale (1999); and Brown (2001)

date. *Log* (*General Experience*) is the logarithm of the number of months elapsed between the analyst's first appearance in the I/B/E/S dataset and her firm-year forecast. *Log* (*Firm Experience*) is the logarithm of the number of months elapsed between the analyst's earliest forecast of the firm in the I/B/E/S dataset and her firm-year forecast. *Firm Breadth and Industry Breadth* are, respectively, the number of unique firms and the number of 2-digit GICS sectors covered by the analyst within the year. Lowballing behavior is defined to be present in any given firm-year if three conditions are met. First, the forecast value must be below the actual EPS. Second, the *Forecast Error* must be either greater than \$0.03 or higher than 5% of the actual EPS (i.e. non-trivial). Third, to reduce the likelihood of mistaking forecasting difficulty for lowballing behavior, the difference between the forecast and the consensus forecast must be greater than \$0.03 or higher than 5% of the consensus forecast. *Lowball* is the number of times over the past 3 years that lowballing forecasts were issued for the firm by the analyst. We measure *Brokerage Size* by the number of analysts employed by the brokerage. *Loss* is an indicator that equates to unity if the announced earnings of the firm are negative, and zero otherwise. *Forecast Dispersion* is the standard deviation of forecast values in a given firm-year. *Analyst Coverage* is the number of unique analysts who have contributed at least one forecast in that firm-year. *Leverage* is the sum of short-term debt and long-term borrowings, scaled by total assets. *Book-to-Market Ratio* is the ratio of firm book value to market capitalization in that year. *Log* (*Total Assets*) is the logarithm of total firm assets. *ROA Volatility* is the standard deviation of return on assets over the past 3 years. We also include lagged measures of forecasting performance and herding behavior. Lagged variables are indicated by the prefix "*Prev*". All continuous variables are winsorized at the 1st and 99th percentile values to limit the influence of outliers.

1.7. Descriptive Statistics

Our sample comprises 3462 firms, 282 brokerages, 5429 analysts, 123038 firm-year forecasts from the years 1996 to 2014. The median brokerage is substantially large with 56 analysts under its employment. Firms in our sample are covered by at least 5 analysts and the median *Analyst Coverage* is 18. The median firm and the smallest firm have *Total*

Assets of about $4300 million and $140 million, respectively. Most firm-years are profitable with only about 8.9% reporting losses. The median (mean) *Forecast Error* is 4 (14) cents while the median (mean) *Herding Rate* is 25.0% (28.4%). By definition, *Degree Centrality*, *Closeness Centrality*, *Betweenness Centrality*, and *Eigenvector Centrality* are non-negative. However, many analysts in the sample have negative values of *Analyst Centrality* because we used standardized[14] variables in the PCA-extraction. Most analysts seem to focus on a single industry with a median (mean) *Industry Breadth* of 1 (1.62). The average analyst has about 14 firms in her coverage portfolio.

[Insert Table 1]

Univariate correlations between the network centrality measures and analyst/brokerage characteristics are presented in Table 1 Panel B. Consistent with the notion that *Analyst Centrality* is capturing aggregate network connectedness, the correlations of *Analyst Centrality* with its component centrality measures are not only positive but large in magnitude. These correlations range from 47.5% to 75.9%. *Analyst Centrality* also shares a strong and positive correlation with *Industry Breadth.* This is unsurprising because an analyst's wide sector coverage will increase the probability that she will be connected to an additional colleague. In aggregate, more connections should generally increase one's centrality in the network. Comparatively, *Firm Breadth* is also positively correlated with *Analyst Centrality* but the correlation magnitude is low. This is expected because, at the extreme, all firms in an analyst's coverage portfolio may fall under the same GICS industry. In such a case, coverage of firms in one industry does not yield connectedness benefits beyond the first firm (refer to Section 1.1 for the definitions of network connections). *General Experience* is weakly correlated with *Analyst Centrality* (2.0%)*.* We exclude *Firm Experience* in correlation analysis because a given analyst-year has multiple values of it, making its correlations with *Analyst Centrality* uninformative for our purpose. *Brokerage Size* has a positive but weak correlation with *Analyst Centrality* (5.2%)*.*

2. Main Results

[14] Standardization accounts for the different scales of the four network centrality variables.

2.1. Analyst Centrality and Forecast Accuracy

To test the relation between *Analyst Centrality* and analyst forecast accuracy, we estimate regressions according to specification (7).

$$Forecast\ Accuracy_{i,f,t} = \alpha + \beta_1 \cdot Analyst\ Centrality_{i,t} + \theta \cdot Controls_{i,f,t} + \varepsilon_{i,f,t} \qquad (7)$$

[Insert Table 2]

Table 2 Panel A reports our first set of regression results. The dependent variable in Column (1) and (2) is the *Clement-Tse Accuracy* measure and we follow the regression specification in Clement and Tse (2005). The independent variables, including *Analyst Centrality*, are normalized using the normalization methodology of the *Clement-Tse Accuracy* measure. We first compute the maximum (*Max Variable*) and minimum (*Min Variable*) values of each independent variable in a given firm-year. Thereafter, each independent variable is normalized as the ratio of the difference between the *Raw Variable Value* and *Max Variable* to the difference between *Max Variable* and *Min Variable*. Hence, each normalized variable is bounded between zero and unity. Formally, each independent variable is normalized as follows.

$$Normalized\ Variable_{i,f,t} = \frac{Raw\ Variable\ Value_{i,f,t} - Min.Variable_{i,f,t}}{Max.Variable_{i,f,t} - Min.Variable_{i,f,t}} \qquad (8)$$

Variables that are invariant at the firm-year level are excluded from the specification in Column (1) due to the normalization procedure. From Column (1), *Analyst Centrality* positively predicts *Clement-Tse Accuracy,* implying that within-brokerage network centrality improves forecast accuracy. In Column (2), we replace the normalized independent variables with their raw variants while keeping *Clement-Tse Accuracy* as the dependent variable. In place of normalization, we include instead firm characteristics as additional controls. Furthermore, we include brokerage-year fixed effects to rule out concerns that *Analyst Centrality* is a spurious proxy of unobserved brokerage characteristics (e.g. prestige and resources). From Column (2), we obtain supporting evidence that *Analyst Centrality* positively predicts forecast accuracy.

Our findings are robust to alternative measures of forecast accuracy. The independent variables in Columns (3) and (4) are not normalized. In Column (3), we use *Forecast Error* as the dependent variable. The firm-year fixed effects absorb variables that have no variation within the firm-year. We find that *Analyst Centrality* is negatively associated with *Forecast Error*, suggesting that high centrality analysts tend to display better forecast accuracy. Specifically, one standard deviation increase in *Analyst Centrality* is associated with a decrease in forecast error of about 0.3 cents, or about 7.5% of the median forecast error. We arrive at similar conclusions when we employ brokerage-year fixed effects and use *Normalized Forecast Error* as the dependent variable in Column (4).

The relation between forecast accuracy and the control variables are generally consistent with the extant literature. Forecast revision frequency, general experience, firm experience, brokerage size, and lagged measures of forecast accuracy are significantly associated with higher forecast accuracy, while forecast horizon, industry breadth, ROA volatility, and historical lowballing behavior predict forecast accuracy negatively. In summary, the positive association between *Analyst Centrality* and various measures of forecast accuracy is incremental to various control variables that prior research finds to predict forecast accuracy.

To gain insights into the dimension of network centrality that is contributing to forecasting performance, we regress the four constituents of *Analyst Centrality* on *Normalized Forecast Error* in the same regression which also includes firm-year fixed effects. Panel B Column (1) suggests that our previous results are driven by *Betweenness Centrality*. In Columns (2) to (5), we rerun our tests and include each centrality measure by itself in order to avoid multicollinearity among the four measures. The results indicate that each of the four centrality measures is negatively associated with *Forecast Error* when employed in isolation. However, *Betweenness Centrality* dominates the other centrality measures, suggesting that the forecast accuracy of high centrality analysts can be attributed to lower information exchange costs.

2.2. Analyst Centrality and Herding Behavior

The sequence of our tests in Table 3 follows that in Table 2. To test the relation between *Analyst Centrality* and herding behavior, we estimate regressions following specification (9).

$$Herding\ Behavior_{i,f,t} = \alpha + \beta_1 \cdot Analyst\ Centrality_{i,t} + \theta \cdot Controls_{i,f,t} + \varepsilon_{i,f,t} \quad (9)$$

[Insert Table 3]

We present our results in Table 3 Panel A. In Column (1), we regress *Analyst Centrality* on *Clement-Tse Herding Rate* and include independent variables that are normalized according to Clement and Tse (2005). Next, we replace the normalized independent variables with their raw variants and include brokerage-year fixed effects in Column (2). Our results are robust to alternative measures of analyst herding behavior. We adopt *Herding Rate* as the dependent variable and include firm-year and brokerage-year fixed effects in Columns (3) and (4) respectively. In all these tests, we find that *Analyst Centrality* is statistically significant and is associated with a lower tendency to display herding behavior.

In Panel B, we regress the four constituents of *Analyst Centrality* on *Herding Rate*. Similar to our earlier analyses, we find that *Betweenness Centrality* primarily explains the negative effects of within-brokerage centrality on analyst herding tendencies. Additionally, *Degree Centrality* is negatively associated with *Herding Rate*, though the economic magnitude of this relation is considerably weak. Our results in Table 3 suggest that inter-colleague information exchange help analysts to issue innovative forecast revisions, and curtail their herding behavior through lower information exchange costs.

3. Analyst Centrality and Market Reactions to Revisions

If high centrality analysts are poised to benefit from information exchange within their brokerages, we hypothesize that their forecast revisions are also likely to contain more novel and value-relevant information. Since one cannot directly observe the information content of forecast revisions, we assume markets are efficient and test our hypothesis using market reactions to analysts' revisions.

[Insert Table 4]

The dependent variable in Table 4 is the absolute 3-day market-adjusted cumulative abnormal returns centered on the forecast revision date. We include *Forecast Revision* and *Consensus Deviation* as independent variables because we expect that market reactions are stronger when there is a greater divergence of the analyst's revision value from her prior forecast and the prevailing consensus forecast. *Forecast Revision* is the absolute difference between an analyst's revision value and her prior forecast value, scaled by the absolute value of her prior forecast value. *Consensus Deviation* is the absolute difference between an analyst's revision value and the prevailing forecast consensus, normalized by the absolute value of the forecast consensus. To uncover moderating effects of within-brokerage network centrality on characteristics of forecast revisions, we interact *Analyst Centrality* with *Forecast Revision* and *Consensus Deviation* separately. We also control for stock performance during the run-up to the forecast revision date and absorb some unobserved heterogeneity through firm and industry-week fixed effects.

In line with our expectations, Columns (1) and (2) show that forecast revisions that deviate more from the prevailing forecast consensus tend to elicit stronger market reactions. We also find that *Analyst Centrality* does not predict market reactions in isolation. We repeat the tests with *Forecast Revision* and find similar results in Columns (3) and (4). Controlling for the divergence from the analyst's prior forecast, market reactions are stronger for forecast revisions issued by high centrality analysts.

4. Analyst Centrality and Peer-Learning

Thus far, we attribute the relation between *Analyst Centrality* and forecasting outcomes to inter-colleague information exchange within brokerages. We examine this relation empirically using the notion of peer-learning. Suppose high centrality analysts tend to incorporate information or economic assumptions learnt from their colleagues into their forecasts. On average, we should also observe that high centrality analysts re-adjust the biased information in their forecasts when they realize that their colleagues were wrong. For example, if an analyst has incorporated her colleagues' aggregate optimism into her forecast, she ought to revise her forecast downwards, and vice-versa. Furthermore, we conjecture that this mechanism is more pronounced as *Analyst Centrality* increases since high centrality analysts are better positioned to benefit from information exchange within the brokerage.

[Insert Table 5]

In Table 5, we use *Analyst Revision* as the dependent variable, since it measures re-adjustment of the analyst's prior forecast in response to new information such as her colleagues' biases. *Analyst Revision* is a signed variable defined as the difference between an analyst's forecast revision value and her prior forecast value, deflated by the absolute value of her prior forecast value. Therefore, a positive (negative) value of *Analyst Revision* reflects an increment (a decline) in the analyst's forecast value from her previous forecast. Note that *Analyst Revision* is defined differently from *Forecast Revision* (Table 4) which only measures the magnitude of the forecast revision. To examine the effects of peer-learning, we focus on analysts' revision behavior and the ex-post forecasting performance of their brokerage colleagues. For a given forecast revision of an analyst, we collect all instances of her colleagues' *realized* forecast errors within the past 30 days. In other words, the firms covered by the analyst's colleagues must have announced their actual earnings in the abovementioned 30-day window. For each of her colleagues' realized forecast errors, we classify it as *optimistic* if the forecasted value is above the actual EPS, and *pessimistic* if the forecasted value is below the actual EPS. If the forecast error is zero, its classification is neither optimistic nor pessimistic although it is still counted as a forecast error in the 30-day window. Thereafter, we define *Peer Pessimism* as the proportion of pessimistic forecast errors in the 30-day window. *Peer Optimism* is defined symmetrically to *Peer Pessimism*. Note that the sum of *Peer Optimism* and *Peer Pessimism* needs not be unity since some forecasts have no errors. Given an analyst i in brokerage G covering firm f and making a forecast revision on date d, we let $r(j,d)$ equate to unity if colleague j has a realized forecast error 30 days prior to d, and zero otherwise.

$$Peer\ Optimism_{i,f,d} = \sum_{j\neq i, j\in G} \frac{r(j,\ d|\ Forecast\ Value_j > Actual\ Earnings)}{r(j,d)} \quad (10a)$$

$$Peer\ Pessimism_{i,f,d} = \sum_{j\neq i, j\in G} \frac{r(j,\ d|\ Forecast\ Value_j < Actual\ Earnings)}{r(j,d)} \quad (10b)$$

We also include separate interaction terms of *Analyst Centrality* with *Peer Optimism* and *Peer Pessimism*. Apart from stock performance during the run-up to the forecast revision date, we also include *Net Competitors' Firm-week Revisions* as a control variable. In the week of the forecast revision date, we count the number of positive and negative forecast revisions made by other competing analysts covering the same firm. Subsequently, we subtract the proportion of negative forecast revisions from that of positive forecast revisions to compute *Net Competitors' Firm-week Revisions.* By construction, *Net Competitors' Firm-week Revisions* is bounded between zero and unity, with more positive values translating to more positive forecast revisions of firm earnings in aggregate. We argue that this control not only addresses the concern that analysts may mimic the revision strategies of their competitors but it may also capture changes in the information environment of the firm.

Table 5 presents our results. In Column (1), we find that *Peer Pessimism* is associated with more positive *Analyst Revision* while *Peer Optimism* predicts more negative *Analyst Revision*. Further, we find that the effect of *Peer Pessimism* is amplified by *Analyst Centrality*, suggesting that high centrality analysts revise their forecasts more positively after learning *ex-post* that their colleagues have provided pessimistic forecasts. However, the mediating effect of *Analyst Centrality* on *Peer Optimism* is not statistically distinguishable from zero. This implies that the effect of peer-learning on an analyst's revision activity is prevalent only in instances when the forecasts of her colleagues are revealed to be pessimistic. In Column (2), we add control variables to the model and find similar results. Notably, *Peer Optimism* ceases to be statistically significant upon the addition of control variables while *Peer Pessimism* and its interaction with *Analyst Centrality* continue to predict positive forecast revisions.

Our results in Columns (1) and (2) are consistent with the notion that high centrality analysts update their forecasts to unravel the information errors of their colleagues via a peer-learning mechanism. However, an alternative explanation is that *Analyst Centrality* is correlated with the ability to process information. This alternative is in conflict with our preferred peer-learning hypothesis because the public availability of ex-post forecasting performance does not necessitate information exchange between an analyst and her brokerage peers. To rule out the information-processing hypothesis, we introduce 2 other variables – *Global Pessimism* and *Global Optimism* – in our tests. *Global Pessimism* is the proportion of all *pessimistic* forecast errors (notably, including those of non-colleagues) in the same GICS sector within the 30-day window before the analyst revision date. *Global Optimism* is defined symmetrically. Similar to their *Peer*-level analogs, these *Global*-level measures of forecasting performance are also publicly available. If the information-processing hypothesis is true, we ought to find that the revision activity of high centrality analysts is also responsive to *Global Pessimism* and *Global Optimism*. However, the contained information in these *Global*-level variables is not admissible in a peer-learning mechanism because it is implausible that analysts engage in systematic information exchange with analysts from other brokerages.

We replace *Peer Pessimism* (*Optimism*) with *Global Pessimism* (*Optimism*) in Column (3). We find that the interactions between the *Global*-level variables and *Analyst Centrality* fail to predict the revision activity of analysts. Aligned with our predictions, the effect of peer-learning is not observed when we employ a source of *ex-post* forecasting performance that is publicly available but which analysts cannot learn from. In the final test, we run a horse race among the *Peer*-level variables, *Global*-level variables, and their respective interactions with *Analyst Centrality* in Column (4). Consistent with our findings in the previous columns, the interaction between *Peer Pessimism* and *Analyst Centrality* predicts positive analyst revisions but interactions terms between *Global*-level variables and *Analyst Centrality* are not statistically significant.

In summary, our results are consistent with the notion that analysts learn from their colleagues' pessimistic forecast errors and the extent of learning is increasing in *Analyst Centrality*. Additionally, we find that analysts learn from *pessimistic* forecast errors of their colleagues, but not from *optimistic* ones. This pattern of asymmetric learning is consistent with the management-relations cultivation hypothesis put forth in prior studies.[15] According to this hypothesis, analysts may be incentivized to issue optimistic forecasts to cultivate management relations in a multi-task environment. The implication to our study

[15] See for example – Francis and Philbrick, 1993; Das et al., 1998; Lim, 2001; Matsumoto, 2002; Richardson et al., 2004; Ke and Yu, 2006

is that *ex-post* optimism may be less informative in our proposed peer-learning mechanism since analysts may understand, either explicitly or implicitly, the motives behind their colleagues' optimism. This may explain why only *Peer Pessimism*, and not *Peer Optimism*, drives our results on peer-learning. This finding also helps us to substantiate a peer-learning hypothesis over other alternative explanations by suggesting that there is strategic incorporation of new information learned through colleagues.

5. Analyst Centrality and Innate Ability

While our preceding evidence supports a peer-learning explanation, we have not disentangled the influence of innate ability from *Analyst Centrality*. An alternative explanation is that higher *Analyst Centrality* scores are indicative of higher analyst ability. If we liken the analyst labor market to a tournament in which performance determines an analyst's continuation in the profession, we expect high-ability analysts to have longer tenures in brokerages. Owing to their longer tenures, high-ability analysts may consequently develop more within-brokerage connections. This can happen in at least two ways. First, a brokerage may strategically build a network of analysts around a high-ability analyst to lever on her forecasting ability. Second, a brokerage may confer more coverage responsibilities (e.g. coverage across sectors) to a high-ability analyst over time. Under this alternative explanation, superior forecasting outcomes are spuriously driven by innate ability, and not by peer-learning.

While analyst ability is unobservable, prior research suggests that membership in the Institutional Investor All-American Research Team (all-star) is reflective of analyst ability.[16] We separate the effect of innate ability from peer-learning by adding an *All-Star* indicator alongside *Analyst Centrality* in horse-race regressions. Due to data limitations, we end our sample period in 2008. Under the peer-learning hypothesis, we expect *Analyst Centrality* to retain its associations with superior forecasting outcomes.

[Insert Table 6]

[16] All-star analysts produces more accurate forecasts (Stickel, 1992), elicit stronger market reactions around their forecasts (Gleason and Lee, 2003), and attract more investment banking deal flows (Clarke et al., 2007)

Table 6 presents our results. In Column (1), we find that *All-Star*, in the absence of *Analyst Centrality*, is positively associated with *Clement-Tse Accuracy*. This is consistent with prior research that all-star analysts produce more accurate forecasts than unranked analysts. We include *Analyst Centrality* alongside *All-Star* in Column (2) and continue to find a positive relation between *All-Star* and forecast accuracy. Crucially, *Analyst Centrality* is also a positive predictor of forecast accuracy. We repeat the above analysis for *Clement-Tse Herding Rate.* In Columns (3) and (4), we find that the effects of *Analyst Centrality* and *All-Star* on herding behavior are not only statistically significant but also distinct from each other. Similarly, Columns (5) and (6) show that our prior results on market reactions are robust to the inclusion of *All-Star* and its interaction with *Consensus Deviation*. In aggregate, our results suggest that the effect of *Analyst Centrality* on forecasting performance is incremental to that of innate analyst ability.

6. Causal Effects of Analyst Centrality on Forecast Accuracy

We exploit brokerage closures (Kelly and Ljungqvist, 2012) from years 2000 to 2007 as an exogenous employment shock to *Analyst Centrality* scores. Derrien and Kecskes (2013) find that most brokerage closures stem from breakdowns in business strategy, suggesting that these individual analysts are unlikely to have systematically different forecasting ability. We only use brokerage closures because they are cleaner employment shocks for our purpose. For each closure event, we track all analysts who subsequently find employment in another brokerage and cover the same firm pre- and post- the closure event. Therefore, our unit of observation in this quasi-natural experiment is an analyst-firm.

The treatment in this test is *Analyst Centrality Up*, an indicator that equal to one if the analyst's average post-closure *Analyst Centrality* is higher than her average pre-closure value, and to zero otherwise. Since brokerage closures 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 following model.

$$Clement - Tse\ Accuracy_{g,t} = \gamma_g + \tau_t + \sum_{j=-m, j\neq 0}^{+n} \beta_j \cdot D_{g,t}(t = k + j) + \theta_{g,t} \cdot \delta + \varepsilon_{g,t} \quad (11)$$

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 closure 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 closures in our sample occur at various points in time. We let $j \neq 0$ because we skip the year of the brokerage closure. 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 10-year window[17] centered on the brokerage closure event. We choose *Clement-Tse Accuracy* 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 five temporal leads (*Pre-Treatment*) and five temporal lags (*Post-Treatment*) of the treatment.[18] The m^{th} temporal *Pre-Treatment* is an indicator that equates to unity only in the m^{th} year before the brokerage closure 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 after he finds new employment and only if the analyst-firm is treated, and zero otherwise. We add year dummies and analyst-firm dummies in all specifications.

[Insert Table 7]

We present results from the generalized difference-in-difference model in Table 7. In Columns (1) to (3), we find that the coefficients of the *Pre-Treatment* indicators are largely

[17] For each analyst-firm, we look back up to 5 years before brokerage closure and skip the year in which the brokerage closure occurs. We only start the 5-year post-closure window when the analyst is employed at a new brokerage because analysts take different lengths of time to find new employment. Most analysts in our sample find employment within a year. This convention essentially lumps the brokerage closure events and the treatment into one single time point.

[18] In this study, our choice of $n = m$ is arbitrary but the generalized model allows n to differ from m. Our results hold when we choose an alternative window width of $n = 3$ and $m = 3$.

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 employment shocks, is unlikely to be dependent on their pre-shock performance. On the other hand, we find that the coefficients of the *Post-Treatment* indicators are all statistically significant and positively loaded on *Clement-Tse Accuracy*. This suggests that the positive effects of higher within-brokerage network centrality on forecast accuracy only occur after treatment has been administered. As robustness checks, we include analyst time trends in Column (2) and further add analyst-firm time trends in Column (3) to help control for confounding heterogeneity. In summary, our findings suggest that among analysts who found subsequent employment after brokerage closures, analysts who experience an increment in *Analyst Centrality* show higher forecast accuracy than those who did not. Moreover, evidence shows that the differences in forecast accuracy manifest largely in the post-treatment period, and not in the pre-treatment period.

7. Conclusion

Overall, our evidence suggests that an analyst's colleagues are valuable sources of information which help analysts make more accurate forecasts. Using 2-digit GICS sector overlaps in their coverage portfolios, we build an information network within a brokerage house to measure potential information flow among analysts. We find that analysts who are better-connected in their brokerage networks produce more accurate forecasts and less herding forecast revisions. We also find that analyst who are more centrally located in these information networks generate greater market reactions to their forecast revisions. This suggests that market participants believe that high centrality analysts provide more informative revisions, which is consistent with better information access.

We provide more direct evidence that high centrality analysts are more likely to learn from their colleagues. If high centrality analysts are tapping into their colleagues' knowledge and expertise, we expect that they would be more likely to include such information in their forecast revisions. Consistent with this argument, we find that high centrality analysts are 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 analysis suggests 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 better understand the causal relation between within-brokerage network centrality and forecast accuracy, we exploit exogenous brokerage closures from 2000 to 2007. These closures force analysts to join new brokerage houses and exogenously change their network centrality scores. We find that analysts who experienced increases in centrality scores at their new employers improve their forecast accuracy.

REFERENCES

Bhojraj, S., C. Lee, and D. K. Oler, 2003. What's My Line? A Comparison of Industry Classification Schemes for Capital Market Research. Journal of Accounting Research 41, 745-774.

Bradshaw, M. T., 2011. Analysts' Forecasts: What Do We Know After Decades of Work? Available at SSRN 1880339.

Brown, L. D., 2001. A Temporal Analysis of Earnings Surprises: Profits versus Losses. Journal of Accounting Research 39, 221-241.

Brown, L. D., A. C. Call, M. B. Clement, and N. Y. Sharp, 2015. Inside the "Black Box" of Sell-Side Financial Analysts. Journal of Accounting Research 53, 1-47.

Chen, S., and D. A. Matsumoto, 2006. Favorable Versus Unfavorable Recommendations: The Impact on Analyst Access to Management-Provided Information. Journal of Accounting Research 44, 657-689.

Clement, M. B., 1999. Analyst Forecast Accuracy: Do Ability, Resources, and Portfolio Complexity Matter? Journal of Accounting and Economics 27, 285-303.

Clement, M. B., and S. Y. Tse, 2005. Financial Analyst Characteristics and Herding Behavior in Forecasting. The Journal of Finance 60, 307-341.

Clement, M. B., L. Koonce, and T. Lopez., 2007. The Roles of Task-Specific Forecasting Experience and Innate Ability in Understanding Analyst Forecasting Performance. Journal of Accounting and Economics 44, 378–398.

Cohen, L., A. Frazzini, and C. Malloy, 2010. Sell-Side School Ties. The Journal of Finance 65, 1409-1437.

Das, S., C. B. Levine, and K. Sivaramakrishnan, 1998. Earnings Predictability and Bias in Analysts' Earnings Forecasts. The Accounting Review 2, 277-294.

David, H., 2003. Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing. Journal of Labor Economics 21, 1-42.

Derrien, F., and A. Kecskés, 2013. The Real Effects of Financial Shocks: Evidence from Exogenous Changes in Analyst Coverage. The Journal of Finance 68, 1407-1440.

Francis, J., and D. Philbrick, 1993. Analysts' Decisions as Products of a Multi-Task Environment. Journal of Accounting Research 31, 216-230.

Hilary, G., and R. Shen, 2013. The Role of Analysts in Intra-Industry Information Transfer. The Accounting Review 88, 1265-1287.

Jacob, J., T. Z. Lys, and M. A. Neale, 1999. Expertise in Forecasting Performance of Security Analysts. Journal of Accounting and Economics 28, 51-82.

Ke, B., and Y. Yu, 2006. The Effect of Issuing Biased Earnings Forecasts on Analysts' Access to Management and Survival. Journal of Accounting Research 44, 965-999.

Kelly, B., and A. Ljungqvist, 2012. Testing Asymmetric-Information Asset Pricing Models. Review of Financial Studies 25, 1366-1413.

Larcker, D. F., E. C. So, and C. C. Wang, 2013. Boardroom Centrality and Firm Performance. Journal of Accounting and Economics 55, 225-250.

Lawyer, G., 2015. Understanding the Influence of All Nodes in a Network. Scientific Reports 5.

Lees, F. A., 1981. Public Disclosure of Corporate Earnings Forecasts. Conference Board.

Lim, T., 2001. Rationality and Analysts' Forecast Bias. The Journal of Finance 56, 369-385.

Matsumoto, D. A., 2002. Management's Incentives to Avoid Negative Earnings Surprises. The Accounting Review 77, 483-514.

Mikhail, M. B., B. R. Walther, and R. H. Willis, 1997. Do Security Analysts Improve Their Performance with Experience? Journal of Accounting Research 35, 131-157.

Mikhail, M. B., B. R. Walther, and R. H. Willis, 1999. Does Forecast Accuracy Matter to Security Analysts? The Accounting Review 74, 185–200.

Newman, M. E., 2003. The Structure and Function of Complex Networks. SIAM Review 45, 167-256.

Richardson, S., S. H. Teoh, and P. D. Wysocki, 2004. The Walk-down to Beatable Analyst Forecasts: The Role of Equity Issuance and Insider Trading Incentives. Contemporary Accounting Research 21, 885-924.

Soltes, E., 2014. Private Interaction between Firm Management and Sell-Side Analysts. Journal of Accounting Research 52, 245-272.

Stein, J. C., 2008. Conversations among Competitors. American Economic Review 98, 2150– 34 2162.

Yin, H., and H. Zhang, 2014. Tournaments of Financial Analysts. Review of Accounting Studies 19, 573-605.

APPENDIX I

Figure 1.

Example of a poorly-connected analyst (Analyst A) and a well-connected analyst (Analyst B) in a within-brokerage network. Analyst A is only connected to 2 other colleagues. Analyst B has 17 connections.

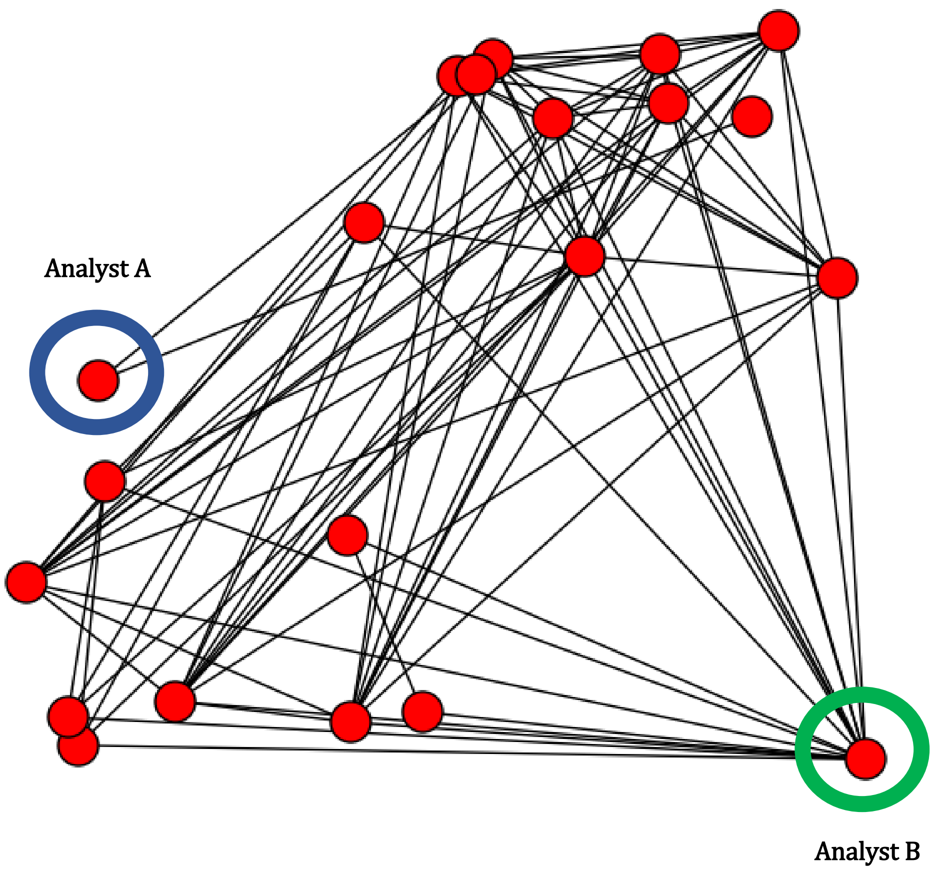

APPENDIX II

II.A. Analyst Centrality and Forecast Accuracy

Under the framework of network theory, social power in a network is relational in nature. A given node has social power only because it can dominate other nodes. Since social power is a function of relational patterns in the network, the amounts of social power across differently-structured networks may be highly heterogeneous. For instance, in a low-density network where there are few connections among nodes, the potential for social power to emerge there is low. On the other hand, in a high-density network where nodes are intricately linked to one another, there is a comparatively higher potential for the exertion of social power by nodes. Besides variation in the network-level (macro) amounts of power, the distribution of social power across nodes (micro) in a given network may also be unequal. The micro-level analysis of social power is the focus of this paper as we investigate the propensities of information transfer afforded to individual analysts by their embedded positions in their brokerages.

In a given network, various embedded positions confer divergent tradeoffs between benefits and constraints to nodes. However, there is no single measure to describe all benefits and constraints consummately because social power is a multi-faceted construct (Newman, 2003). Lawyer (2015) further suggested that the optimality of every network theory measure is dependent on the structure of the most important (or central) nodes, and may be sub-optimal in analyzing the remainder of the network. For example, a particular node in a network may be favored by the degree centrality measure but is less favored by the eigenvector centrality measure. Since it is ex-ante challenging to identify the most important nodes in networks, we employ 4 widely-used measures of network centrality to increase the likelihood that centrality is optimally measured in networks. The 4 measures of network centrality used in this study are degree, betweenness, closeness, and eigenvector centrality. In this section, we provide formal mathematical definitions, and discussions of each centrality measure in the context of this paper.

II.B. Degree Centrality

Degree centrality is related to the number of colleagues that an analyst is immediately connected to in a brokerage network.

For example, in Figure 1, 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.

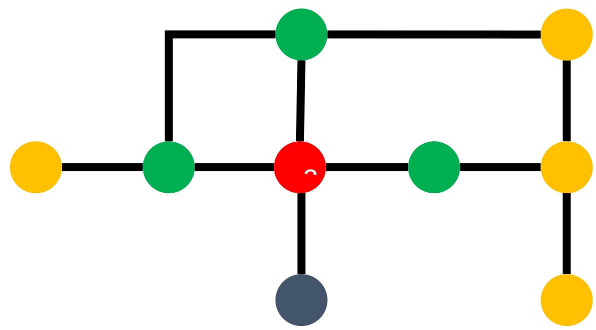

Figure 2

Degree centrality favors analysts who have relatively more opportunities. 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 3 neighbors – represented as green nodes. However, ***Z*** does not have an alternative connection if A is unable to provide resources. Under degree centrality, ***A*** possesses more social power than ***Z*** because the former is less dependent on any single colleague. 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.

$$Centrality_{Degree}(i,\ G) = \frac{\sum_{j \neq i, j \in G}^{N_G - 1} f(i,\ j)}{N_G - 1}$$

$$N_G = Number\ of\ analysts\ in\ network\ G$$

$$f(i,\ j) = \begin{cases} +1, & i\ is\ connected\ to\ j \\ 0, & i\ is\ not\ connected\ to\ j \end{cases}$$

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

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

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

II.C. 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.

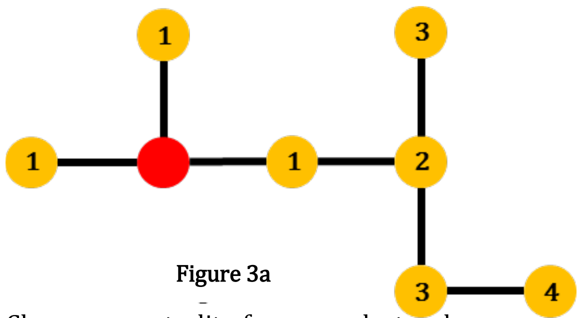

Figure 3a

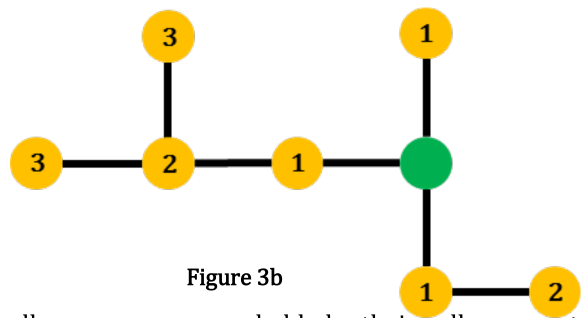

Figure 3b

Closeness centrality favors analysts who can access their colleagues, or are reachable by 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. This advantage is salient if there are isolated or disconnected components (cluster of nodes) in the network. When such components are present, closeness centrality, unlike degree centrality, can differentiate global (network-wide) centrality from local centrality. 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 in the orange nodes. We show below 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.

$$Centrality_{Closeness}(i,\ G) = \frac{N_G - 1}{\sum_{j \neq i, j \in G}^{N_G - 1} d(i,j)} \qquad \begin{array}{l} N_G = Number\ of\ analysts\ in\ network\ G \\ d(i,j) = Shortest\ path\ distance\ between\ i\ and\ j \end{array}$$

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

$$Centrality_{Closeness}\ of\ analyst\ \boldsymbol{A} = \frac{8-1}{1+1+1+2+3+3+4} = 0.467$$

$$Centrality_{Closeness}\ of\ analyst\ \boldsymbol{Z} = \frac{8-1}{1+1+1+2+2+3+3} = 0.538$$

II.D. Betweenness Centrality

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

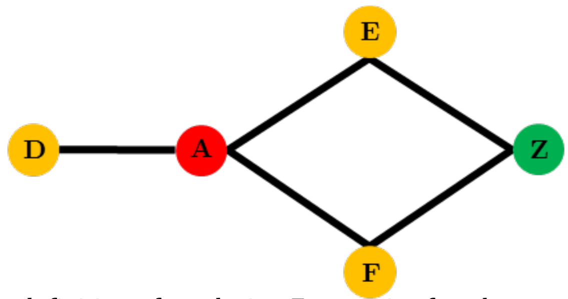

Figure 4

We first elaborate on the definition of geodesics. For a pair of analysts, a geodesic between them is a path of the shortest possible length. An analyst pair may have more than 1 geodesic. Consider analysts ***E*** and ***F***. Analyst ***E*** may reach analyst ***F*** via 2 paths – ***EZF*** and ***EAF***. Since both paths have lengths of 2, and the shortest possible path length between the analyst pair is 2, 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***.

Betweenness centrality favors analysts who have brokering capacity. In Figure 4, analyst ***A*** is in an advantageous brokering position relative to analysts ***D*** and ***E***. Should ***D*** and ***E*** choose to interact with each other, they must do so via ***A***. In contrast, if ***A*** chooses to interact with either ***D*** or ***E***, he may do so without the need to pass through any colleagues. The formal mathematical definition of betweenness centrality for a given analyst ***i*** is as follows.

$$Centrality_{Betweenness}(i,\ G) = \frac{2}{(N_G-1)(N_G-2)} \sum\nolimits_{x \neq i, y \neq i, x,y \in G} \frac{s(x,y|i)}{s(x,y)}$$

$$N_G = Number\ of\ analysts\ in\ network\ G$$

$$s(x,y) = Number\ of\ geodesics\ between\ any\ pair\ x\ and\ y$$

$$s(x,y|i) = Number\ of\ geodesics\ between\ any\ pair\ x\ and\ y\ passing\ through\ i$$

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

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

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

II.E. 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 application 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 $\boldsymbol{P}$, $\boldsymbol{Q}$, $\boldsymbol{R}$, and $\boldsymbol{S}$. 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 $\boldsymbol{P}$ and $\boldsymbol{Q}$ are connected. On the other hand, element $\mathbf{m_{1,3}}$ equates to zero because there is no connection between $\boldsymbol{P}$ and $\boldsymbol{R}$.

Figure 5

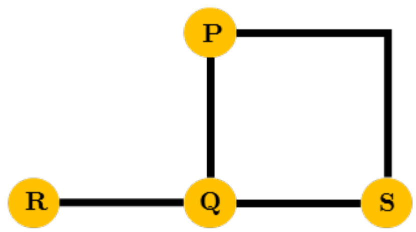

$$\mathbf{M_G} = \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix}$$

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.

$$\mathbf{K_G} = \begin{bmatrix} 2 \\ 3 \\ 1 \\ 2 \end{bmatrix} \quad \textit{where } \mathbf{k_{1,1}}, \mathbf{k_{2,1}}, \mathbf{k_{3,1}}, \textit{and } \mathbf{k_{4,1}} \textit{ are the unnormalized degree centralities of } \boldsymbol{P}, \boldsymbol{Q}, \boldsymbol{R}, \textit{and } \boldsymbol{S} \textit{ respectively}$$

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 $\boldsymbol{P}$, of the resulting matrix carries a value of 5, the cumulative centrality of her connections. This value is derived from $\boldsymbol{P}$s connections – $\boldsymbol{Q}$ and $\boldsymbol{S}$ – 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.

$$\boldsymbol{M_G} \cdot \boldsymbol{M_G} \cdot \boldsymbol{K_G} = \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 5 \\ 5 \\ 3 \\ 5 \end{bmatrix} = \begin{bmatrix} 10 \\ 13 \\ 5 \\ 10 \end{bmatrix} \qquad \mathbf{M_G} \cdot \mathbf{M_G} \cdot \mathbf{M_G} \cdot \mathbf{K_G} = \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 10 \\ 13 \\ 5 \\ 10 \end{bmatrix} = \begin{bmatrix} 23 \\ 25 \\ 13 \\ 23 \end{bmatrix}$$

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.

$$Centrality_{Eigenvector}(i,\ G) = element\ v_i\ of\ matrix\ \mathbf{V_G}$$

$$\mathbf{V_G}\ is\ an\ eigenvector\ of\ the\ network's\ adjacency\ matrix\ \mathbf{M_G}$$

$$or, \qquad \mathbf{M_G} \cdot \mathbf{V_G} = \boldsymbol{\lambda} \cdot \mathbf{V_G}$$

where $\boldsymbol{\lambda}$ is a scalar

Suppose we perform a matrix multiplication between $\mathbf{M_G}$ and its eigenvector $\mathbf{V_G}$.

$$\mathbf{M_G} \cdot \mathbf{V_G} = \boldsymbol{\lambda} \cdot \mathbf{V_G}$$

And multiply the resulting vector with $\mathbf{M_G}$.

$$\mathbf{M_G} \cdot \boldsymbol{\lambda} \cdot \mathbf{V_G} = \boldsymbol{\lambda} \cdot \mathbf{M_G} \cdot \mathbf{V_G} = \boldsymbol{\lambda} \cdot \boldsymbol{\lambda} \cdot \mathbf{V_G} = constant \cdot \mathbf{V_G}$$

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.

Appendix III. 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 that the analyst is immediately connected to in the brokerage network (see Appendix II 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 II 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 II for details on centrality measures).
Eigenvector Centrality	Eigenvector Centrality is related to the notion that the centrality of the analyst is high if her connected colleagues are also well-connected in the brokerage network (see Appendix II 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.
General Experience	Number of months elapsed between the analyst's first appearance in the I/B/E/S dataset and her firm-year forecast.
Firm Experience	Number of months elapsed between the analyst's earliest forecast of the firm in the I/B/E/S dataset 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	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.

Appendix III. (Continued)

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 Firm
Consensus Deviation	Absolute difference between the analyst's revision value and the prevailing firm-year consensus value, scaled by the absolute value of the latter variable.
Forecast Revision	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	Number of firm-year forecasts issued by all analysts in the week of the analyst's forecast revision.
Disagreement	In the week of the analyst's revision, we collect firm-year forecasts issued by all analysts and assign a value of +1 (-1) to an issued forecast if its value is higher (lower) than the prevailing consensus. Disagreement is the standard deviation of these assigned values in the week. If there is only 1 forecast issued in the week, Disagreement is set to zero.
Peer Pessimism	For a given forecast revision of an analyst, we collect all instances of her brokerage colleagues' realized forecast errors within the past 30 days. Alternatively, the firms covered by the analyst's brokerage colleagues must have announced their actual earnings in the same 30-day window. For each forecast error of the analyst's colleagues, we classify it as optimistic if the forecasted value is above the actual earnings, and categorize them as pessimistic if the forecasted value is below the actual earnings. If the colleague's forecast error is zero, the instance is neither optimistic nor pessimistic but is still counted in the window. Following this, Peer Pessimism is the proportion of pessimistic forecast errors issued by brokerage colleagues in the 30-day window.
Peer Optimism	Peer Optimism is defined symmetrically to Peer Pessimism.
Global Pessimism	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, and categorize them as pessimistic if the forecasted value if below the actual earnings. If the forecast error is zero, it is neither optimistic nor pessimistic but is still counted in the window. Following this, Global Pessimism is the proportion of pessimistic forecast errors issued by all analysts in the 30-day window.
Global Optimism	Global Optimism is defined symmetrically to Global Pessimism.

Appendix III. (Continued)

Competitors' Revisions	In the week of the analyst's revision, we collect firm-year forecasts issued by all other analysts and assign a value of +1 (-1) to an issued forecast if its value is higher (lower) than the analyst's previous firm-year forecast. Competitors' Revisions is the sum of these assigned values.
All-Star	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.
Forecast Error	Absolute difference between the analyst's final firm-year forecast and the actual firm-year EPS value.
Normalized Forecast Error	Forecast Error scaled by the average Forecast Error across analysts in the firm-year.
Clement-Tse Accuracy	In each firm-year, the maximum (Max FE) and minimum (Min FE) values of Forecast Error are computed. Clement-Tse Accuracy of an analyst in a firm-year is the ratio of the difference between Max FE and Forecast Error to the difference between Max FE and Min FE. Hence, Clement-Tse Accuracy is bounded between 0 and unity, with higher values reflecting a higher degree of forecast accuracy.
Herding Rate	Following Clement and Tse (2005), an analyst's forecast revision in the firm-year is classified as non-herding only 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. Otherwise, the revision is classified as herding. Herding Rate is the ratio of herding revisions to the total number of revisions made by an analyst in the firm-year.
Clement-Tse Herding Rate	In each firm-year, the maximum (Max HR) and minimum (Min HR) values of Herding Rates are computed. Clement-Tse Herding Rate of an analyst in a firm-year is the ratio of the difference between Herding Rate and Min HR to the difference between Max HR and Min HR. Hence, Clement-Tse Herding Rate is bounded between 0 and unity, with higher values reflecting a greater degree of herding behavior.
Abs. CAR	Market-adjusted cumulative abnormal returns, centered on the forecast revision date.
Analyst Revision	Difference between the forecast revision value and the previous forecast value, scaled by the absolute value of the latter variable. Unlike Forecast Revision (see above for definition), Analyst Revision is a signed variable. Positive (negative) values of Analyst Revision reflects an increment (a decline) in the analyst's forecasted value from her previous forecast.

Table 1.

Panel A. Summary Statistics

Variable	N	Mean	S.D	p10	p25	p50	p75	p90
Analyst Centrality	12075	0.178	1.007	-1.095	-0.589	0.127	0.866	1.456
Degree Centrality	12075	0.356	0.240	0.108	0.174	0.291	0.471	0.714
Closeness Centrality	12075	0.549	0.175	0.367	0.456	0.535	0.626	0.774
Betweenness Centrality	12075	0.025	0.063	0	0	0	0.019	0.076
Eigenvector Centrality	12075	0.142	0.116	0.007	0.041	0.124	0.216	0.304
Horizon	12075	126.57	79.045	49	90	99	128	273
Brokerage Size	12075	68.739	52.657	14	25	56	103	129
Revision Frequency	12075	3.393	3.087	1	2	3	4	6
Log (General Experience)	12075	4.544	0.493	3.850	4.159	4.575	4.927	5.198
Log (Firm Experience)	12075	4.106	0.467	3.555	3.689	4.078	4.443	4.787
Firm Breadth	12075	14.773	6.716	7	10	14	18	23
Industry Breadth	12075	1.620	0.849	1	1	1	2	3
Lowball	12075	0.458	0.687	0	0	0	1	1
Loss	12075	0.088	0.283	0	0	0	0	0
Forecast Dispersion	12075	0.146	0.228	0.016	0.031	0.069	0.157	0.346
Analyst Coverage	12075	19.337	9.466	8	12	18	25	33
Leverage	12075	0.561	0.229	0.253	0.396	0.558	0.717	0.883
Book to Market Ratio	12075	0.480	0.370	0.126	0.240	0.397	0.629	0.925
Log (Total Assets)	12075	8.366	1.673	6.226	7.187	8.288	9.512	10.574
Forecast Error	12075	0.139	0.295	0.005	0.018	0.040	0.120	0.320
Normalized Forecast Error	12057	101.34	90.958	10.185	41.667	83.582	127.83	207.77
Herding Rate	12075	28.400	29.793	0	0	25.000	50.000	66.667

Panel B. Univariate Correlations

	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	(i)
Analyst Centrality [A]	1								
Degree Centrality [B]	0.735	1							
Closeness Centrality [C]	0.759	0.893	1						
Betweenness Centrality [D]	0.475	0.332	0.286	1					
Eigenvector Centrality [E]	0.591	0.756	0.619	0.358	1				
Brokerage Size [F]	0.052	-0.417	-0.207	-0.131	-0.489	1			
General Experience [G]	0.020	0.004	0.009	0.030	0.026	-0.035	1		
Firm Experience [H]	-0.014	-0.025	-0.022	0.015	0.006	-0.011	0.615	1	
Firm Breadth [I]	0.106	0.058	0.100	0.074	-0.011	0.015	0.248	0.127	1
Industry Breadth [J]	0.616	0.473	0.424	0.551	0.399	-0.084	0.063	0.023	0.224

Table 2.
The Effect of Analyst Centrality on Forecast Accuracy

This table presents results from an ordinary least squares model. The dependent variable in Column (3) of Panel A is *Forecast Error*. *Forecast Error* is the absolute difference between the analyst's final firm-year forecast and the actual firm-year EPS value. *Forecast Error* is measured in percentage points. The dependent variable in Column (4) of Panel A is *Normalized Forecast Error*, defined as *Forecast Error* scaled by the average forecast error in the firm-year. The dependent variable in Columns (1) and (2) of Panel A is *Clement-Tse Accuracy* (see Clement and Tse, 2005). In each firm-year, the maximum (*Max FE*) and minimum (*Min FE*) values of *Forecast Error* are computed. *Clement-Tse Accuracy* of an analyst in a firm-year is the ratio of the difference between *Max FE* and *Forecast Error* to the difference between *Max FE* and *Min FE.* Hence, *Clement-Tse Accuracy* is bounded between 0 and unity, with higher values reflecting a higher degree of forecast accuracy. The key independent variable in Panel A is *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 2 for details on centrality measures). In Panel B, we regress each of the above 4 network centrality measure individually on *Forecast Error*. All independent variables in Column (1) of Panel A are normalized according to Clement and Tse (2005). In each firm-year, the maximum (*Max Var*) and minimum (*Min Var*) values of each variable are computed. Subsequently, each independent variable is normalized as the ratio of the difference between the variable value and *Min Var* to the difference between *Max Var* and *Min Var*. Hence, the normalized value of each independent variable is bounded between 0 and unity, with higher values reflecting a greater magnitude of the variable. This normalization is not performed for all other Columns in Panel A and Panel B. To facilitate presentation, estimated coefficients are multiplied by a factor of 100 in Column (2) of Panel A. Where firm-year fixed effects are employed, control variables specific to the firm-year are excluded from the regressions. Detailed definitions of other variables are in the Appendix. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 2. (Continued)
Panel A. Analyst Centrality and Forecast Accuracy

	(1) Clement-Tse Accuracy	(2) Clement-Tse Accuracy	(3) Forecast Error	(4) Norm. Forecast Error
Analyst Centrality	0.026***	0.616***	-0.003***	-1.983***
	(0.004)	(0.002)	(0.001)	(0.460)
Revision Frequency	0.081***	0.047	-0.003*	-0.095
	(0.004)	(0.000)	(0.002)	(0.067)
Horizon	-0.218***	-0.143***	0.001***	0.456***
	(0.003)	(0.000)	(0.000)	(0.006)
Log (General Experience)	0.025***	-0.187	-0.003	1.846**
	(0.004)	(0.003)	(0.002)	(0.908)
Log (Firm Experience)	0.015***	-0.096	-0.003*	0.026
	(0.003)	(0.003)	(0.002)	(0.790)
Firm Breadth	0.005	0.017	-0.000	-0.111**
	(0.003)	(0.000)	(0.000)	(0.055)
Industry Breadth	-0.038***	-1.120***	0.004***	2.123***
	(0.003)	(0.002)	(0.001)	(0.527)
Lowball	-0.021***	-1.842***	0.007***	5.456***
	(0.003)	(0.002)	(0.002)	(0.564)
Brokerage Size	0.061***		0.000***	
	(0.003)		(0.000)	
Loss		-1.771***		-3.709***
		(0.004)		(1.015)
Forecast Dispersion		11.831***		-1.960
		(0.005)		(1.727)
Analyst Coverage		0.441***		0.147***
		(0.000)		(0.046)
Prev. C-T Accuracy	-0.024***	6.727***		
	(0.003)	(0.004)		
Prev. Forecast Error			0.160***	
			(0.014)	
Prev. Norm. Forecast Error				0.092***
				(0.007)
Firm financial variables	No	Yes	No	Yes
Observations	77,130	108,228	120,756	107,545
R-squared	0.111	0.197	0.734	0.206
Brokerage-Year fixed effects	No	Yes	No	Yes
Firm-Year fixed effects	No	No	Yes	No
Brokerage-Year cluster	No	Yes	No	Yes
Firm-Year cluster	No	No	Yes	No
Analyst-Firm cluster	No	Yes	Yes	Yes

Table 2. (Continued)
Panel B. Individual Measures of Network Centrality and Forecast Accuracy

	(1)	(2)	(3)	(4)	(5)
	Norm. Forecast Error	Norm. Forecast Error	Norm. Forecast Error	Norm. Forecast Error	Norm. Forecast Error
Degree Centrality	-5.014	-2.514			
	(3.673)	(1.673)			
Closeness Centrality	1.625		-2.429		
	(4.045)		(1.976)		
Betweenness Centrality	-12.007**			-11.174**	
	(5.420)			(5.328)	
Eigenvector Centrality	5.655				-0.528
	(4.282)				(3.520)
Revision Frequency	-0.927**	-0.929**	-0.930**	-0.931**	-0.931**
	(0.413)	(0.414)	(0.414)	(0.414)	(0.415)
Horizon	0.576***	0.576***	0.576***	0.576***	0.576***
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Log (General Experience)	-1.253	-1.256	-1.236	-1.191	-1.205
	(0.871)	(0.871)	(0.871)	(0.871)	(0.871)
Log (Firm Experience)	-1.802**	-1.783**	-1.792**	-1.795**	-1.796**
	(0.903)	(0.903)	(0.903)	(0.903)	(0.903)
Firm Breadth	-0.105*	-0.100*	-0.096*	-0.104*	-0.099*
	(0.055)	(0.054)	(0.054)	(0.054)	(0.055)
Industry Breadth	2.752***	2.346***	2.213***	2.466***	2.037***
	(0.570)	(0.520)	(0.492)	(0.514)	(0.496)
Lowball	7.615***	7.615***	7.617***	7.613***	7.616***
	(0.728)	(0.728)	(0.728)	(0.728)	(0.728)
Brokerage Size	0.018**	0.017**	0.020***	0.019***	0.020***
	(0.008)	(0.007)	(0.007)	(0.007)	(0.007)
Prev. Norm. Forecast Error	0.089***	0.089***	0.089***	0.089***	0.089***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Observations	120,232	120,232	120,232	120,232	120,232
R-squared	0.263	0.263	0.263	0.263	0.263
Firm-Year fixed effects	Yes	Yes	Yes	Yes	Yes
Firm-Year cluster	Yes	Yes	Yes	Yes	Yes
Analyst-Firm cluster	Yes	Yes	Yes	Yes	Yes

Table 3.
The Effect of Network Centrality on Herding Behavior

This table presents results from an ordinary least squares model. The dependent variable in Columns (3) and (4) of Panel A is *Herding* Rate. Following Clement and Tse (2005), an analyst's forecast revision in the firm-year is classified as *non-herding* only 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. Otherwise, the revision is classified as *herding*. *Herding Rate* is the ratio of herding revisions to the total number of revisions made by an analyst in the firm-year. The dependent variable in Columns (1) and (2) of Panel A is *Clement-Tse Herding Rate*. In each firm-year, the maximum (*Max HR*) and minimum (*Min HR*) values of *Herding Rates* are computed. *Clement-Tse Herding Rate* of an analyst in a firm-year is the ratio of the difference between *Herding Rate* and *Min HR* to the difference between *Max HR* and *Min HR*. Hence, *Clement-Tse Herding Rate* is bounded between 0 and unity, with higher values reflecting a greater degree of herding behavior. The key independent variable in Panel A is *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 2 for details on centrality measures). In Panel B, we regress each of the above 4 network centrality measure individually on *Herding Rate*. All independent variables in Column (1) of Panel A are normalized according to Clement and Tse (2005). In each firm-year, the maximum (*Max Var*) and minimum (*Min Var*) values of each variable are computed. Subsequently, each independent variable is normalized as the ratio of the difference between the variable value and *Min Var* to the difference between *Max Var* and *Min Var*. Hence, the normalized value of each independent variable is bounded between 0 and unity, with higher values reflecting a greater magnitude of the variable. This normalization is not performed for all other Columns in Panel A and Panel B. To facilitate presentation, estimated coefficients are multiplied by a factor of 100 in Column (2) of Panel A. Where firm-year fixed effects are employed, control variables specific to the firm-year are excluded from the regressions. Detailed definitions of other variables are in the Appendix. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 3. (Continued)
Panel A. Analyst Centrality and Herding Behavior

	(1) Clement-Tse Herding Rate	(2) Clement-Tse Herding Rate	(3) Herding Rate	(4) Herding Rate
Analyst Centrality	-0.027***	-0.619***	-0.540***	-0.468***
	(0.005)	(0.002)	(0.135)	(0.160)
Revision Frequency	0.210***	1.006*	0.551*	0.641**
	(0.005)	(0.006)	(0.331)	(0.289)
Horizon	-0.020***	-0.071***	-0.044***	-0.037***
	(0.004)	(0.000)	(0.005)	(0.005)
Log (General Experience)	-0.004	-0.311	-0.472*	-0.734**
	(0.005)	(0.004)	(0.280)	(0.317)
Log (Firm Experience)	0.000	-1.391***	-0.618**	-0.513**
	(0.004)	(0.003)	(0.286)	(0.260)
Firm Breadth	0.009**	-0.014	0.019	-0.042**
	(0.004)	(0.000)	(0.017)	(0.021)
Industry Breadth	0.032***	0.657***	0.880***	0.466**
	(0.004)	(0.003)	(0.183)	(0.190)
Lowball	0.027***	1.074***	0.698***	1.105***
	(0.003)	(0.002)	(0.173)	(0.145)
Brokerage Size	-0.019***		-0.008***	
	(0.004)		(0.002)	
Loss		0.801		1.376***
		(0.005)		(0.384)
Forecast Dispersion		-0.135		-3.152***
		(0.009)		(0.707)
Analyst Coverage		-0.271***		-0.099***
		(0.000)		(0.016)
Prev. C-T Herding Rate	0.057***	4.575***		
	(0.004)	(0.004)		
Prev. Herding Rate			0.052***	0.056***
			(0.004)	(0.004)
Firm financial variables	No	Yes	No	Yes
Observations	75,651	102,123	120,756	108,228
R-squared	0.044	0.075	0.269	0.060
Brokerage-Year fixed effects	No	Yes	No	Yes
Firm-Year fixed effects	No	No	Yes	No
Brokerage-Year cluster	No	Yes	No	Yes
Firm-Year cluster	No	No	Yes	No
Analyst-Firm cluster	No	Yes	Yes	Yes

Table 3. (Continued)
Panel B. Individual Measures of Centrality and Herding Behavior

	(1) Herding Rate	(2) Herding Rate	(3) Herding Rate	(4) Herding Rate	(5) Herding Rate
Degree Centrality	-2.186*	-0.992**			
	(1.183)	(0.502)			
Closeness Centrality	1.196		-0.821		
	(1.318)		(0.605)		
Betweenness Centrality	-4.383**			-4.232**	
	(1.788)			(1.750)	
Eigenvector Centrality	1.473				-0.713
	(1.402)				(1.099)
Revision Frequency	0.550*	0.549*	0.549*	0.548*	0.548*
	(0.331)	(0.330)	(0.330)	(0.330)	(0.330)
Horizon	-0.044***	-0.044***	-0.044***	-0.044***	-0.044***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Log (General Experience)	-0.458	-0.459	-0.449	-0.434	-0.441
	(0.280)	(0.280)	(0.280)	(0.280)	(0.280)
Log (Firm Experience)	-0.625**	-0.622**	-0.625**	-0.626**	-0.624**
	(0.286)	(0.286)	(0.286)	(0.286)	(0.286)
Firm Breadth	0.014	0.017	0.019	0.016	0.017
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
Industry Breadth	0.796***	0.631***	0.567***	0.671***	0.534***
	(0.184)	(0.168)	(0.162)	(0.170)	(0.163)
Lowball	0.698***	0.700***	0.701***	0.699***	0.701***
	(0.173)	(0.173)	(0.173)	(0.173)	(0.173)
Brokerage Size	-0.012***	-0.011***	-0.010***	-0.010***	-0.010***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Prev. Herding Rate	0.052***	0.052***	0.052***	0.052***	0.053***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Observations	120,756	120,756	120,756	120,756	120,756
R-squared	0.269	0.269	0.269	0.269	0.269
Firm-Year fixed effects	Yes	Yes	Yes	Yes	Yes
Firm-Year cluster	Yes	Yes	Yes	Yes	Yes
Analyst-Firm cluster	Yes	Yes	Yes	Yes	Yes

Table 4.
The Effect of Analyst Centrality on Market Reactions to Forecasts

This table presents results from an ordinary least squares model. The dependent variable in this table is the absolute 3-day market-adjusted cumulative abnormal returns, centered on the forecast revision date. The key independent variables in Columns (1) and (2) are *Analyst Centrality*, *Consensus Deviation,* and their interaction *Analyst Centrality* X *Consensus Deviation*. The key independent variables in Columns (3) and (4) are *Analyst Centrality*, *Forecast Revision*, and their interaction *Analyst Centrality* X *Forecast Revision*. *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 2 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. *Forecast Revision* is the absolute difference between an analyst's revision value and her prior forecast value, normalized by the absolute value of the latter variable. All return variables are unitized in percentage points. Detailed definitions of other variables are in the Appendix. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 4. (Continued)

	(1)	(2)	(3)	(4)
	Abs. CAR [-1, +1]	Abs. CAR [-1, +1]	Abs. CAR [-1, +1]	Abs. CAR [-1, +1]
Analyst Centrality	0.003	0.003	-0.010	-0.011
	(0.010)	(0.010)	(0.010)	(0.010)
Consensus Deviation [A]	0.737***	0.737***		
	(0.046)	(0.046)		
Analyst Centrality X [A]	0.251***	0.250***		
	(0.051)	(0.051)		
Forecast Revision [B]			0.851***	0.852***
			(0.047)	(0.047)
Analyst Centrality X [B]			0.348***	0.348***
			(0.057)	(0.057)
Log (General Experience)	0.077***	0.077***	0.076***	0.075***
	(0.016)	(0.016)	(0.017)	(0.017)
Log (Firm Experience)	-0.011	-0.011	-0.003	-0.003
	(0.018)	(0.018)	(0.019)	(0.019)
Firm Breadth	-0.001	-0.000	-0.001	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Industry Breadth	-0.014	-0.014	-0.011	-0.011
	(0.011)	(0.011)	(0.011)	(0.011)
Number of Forecasts	0.340***	0.340***	0.332***	0.332***
	(0.007)	(0.007)	(0.007)	(0.007)
Disagreement	-0.227***	-0.227***	-0.231***	-0.232***
	(0.030)	(0.030)	(0.031)	(0.031)
Loss	0.438***	0.439***	0.457***	0.458***
	(0.079)	(0.079)	(0.080)	(0.080)
Abs. CAR [-5, -2]	0.034***		0.032***	
	(0.005)		(0.005)	
Abs. CAR [-10, -2]		0.021***		0.020***
		(0.003)		(0.003)
Firm financial controls	Yes	Yes	Yes	Yes
Observations	530,623	530,623	484,934	484,934
R-squared	0.290	0.290	0.293	0.293
Firm fixed effects	Yes	Yes	Yes	Yes
Industry-Week fixed effects	Yes	Yes	Yes	Yes
Firm-Week cluster	Yes	Yes	Yes	Yes
Analyst-Firm cluster	Yes	Yes	Yes	Yes

Table 5.
Analyst Centrality and Learning from Peers' Ex-Post Forecast Errors

This table presents results from an ordinary least squares model. 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, positive (negative) values of *Analyst Revision* reflects an increment (a decline) in the analyst's forecasted value from her previous forecast. The key independent variables in Columns (1) to (4) are *Peer Pessimism*, *Peer Optimism*, *Global Pessimism*, *Global Optimism* and their respective interactions with *Analyst Centrality*. For a given forecast revision of an analyst, we collect all instances of her brokerage colleagues' *realized* forecast errors within the past 30 days. Alternatively, the firms covered by the analyst's brokerage 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, and categorize them as *pessimistic* if the forecasted value is below the actual earnings. If the colleague's forecast error is 0, the instance is neither *optimistic* nor *pessimistic* but is still counted in the window. Following this, *Peer Pessimism* is the proportion of *pessimistic* forecast errors in the 30-day window. *Peer Optimism* is defined symmetrically. Since the analyst's colleagues may have forecast errors of 0, the sum of *Peer Pessimism* and *Peer Optimism* needs not be unity. To construct *Global Pessimism*, we capture all realized forecast errors in the same GICS sector in the same 30-day window. Following this, *Global Pessimism* is the proportion of *pessimistic* forecast errors made by all analysts (including non-colleagues) in the 30-day window. *Global Optimism* is defined symmetrically. *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 2 for details on centrality measures). *Analyst Revision*, *Peer Pessimism*, *Peer Optimism*, and stock return variables are expressed in percentage points. Detailed definitions of other variables are in the Appendix. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 5. (Continued)

	(1) Analyst Revision	(2) Analyst Revision	(3) Analyst Revision	(4) Analyst Revision
Analyst Centrality	-0.224	-0.263	0.019	-0.162
	(0.178)	(0.180)	(0.367)	(0.387)
Peer Pessimism [A]	0.974***	0.544***		0.463***
	(0.182)	(0.173)		(0.173)
Analyst Centrality X [A]	0.373**	0.347**		0.309*
	(0.186)	(0.175)		(0.174)
Peer Optimism [B]	-0.595***	-0.304		-0.228
	(0.203)	(0.194)		(0.194)
Analyst Centrality X [B]	0.268	0.224		0.255
	(0.204)	(0.193)		(0.193)
Global Pessimism [C]			1.350***	1.265***
			(0.408)	(0.408)
Analyst Centrality X [C]			0.059	-0.001
			(0.384)	(0.383)
Global Optimism [D]			-0.361	-0.372
			(0.419)	(0.419)
Analyst Centrality X [D]			-0.233	-0.257
			(0.393)	(0.393)
Competitors' Revisions		0.077***	0.077***	0.076***
		(0.001)	(0.001)	(0.001)
General Experience		0.935	0.899	0.877
		(0.688)	(0.688)	(0.688)
Firm Experience		-1.048**	-0.976*	-0.981*
		(0.527)	(0.527)	(0.527)
Firm Breadth		0.010	0.009	0.009
		(0.018)	(0.018)	(0.018)
Industry Breadth		0.163	0.192	0.176
		(0.172)	(0.172)	(0.172)
Brokerage Size		-0.005**	-0.005*	-0.005**
		(0.003)	(0.003)	(0.003)
X`Abs. CAR [-10, -2]		0.150***	0.150***	0.150***
		(0.009)	(0.009)	(0.009)
Observations	409,397	409,397	409,397	409,397
R-squared	0.170	0.251	0.251	0.251
Analyst-firm fixed effects	Yes	Yes	Yes	Yes
Analyst-firm cluster	Yes	Yes	Yes	Yes

Table 6.
Analyst Centrality and All-Star Status

This table presents results from OLS regressions. Due to data limitations on the identities of all-star analysts, our sample period ends in 2008. The dependent variables are *Clement-Tse Accuracy*, *Clement-Tse Herding Rate*, and *Abs. CAR*. In each firm-year, the maximum (*Max FE*) and minimum (*Min FE*) values of *Forecast Error* are computed. *Clement-Tse Accuracy* of an analyst in a firm-year is the ratio of the difference between *Max FE* and *Forecast Error* to the difference between *Max FE* and *Min FE.* Hence, *Clement-Tse Accuracy* is bounded between 0 and unity, with higher values reflecting a higher degree of forecast accuracy. *Forecast Error* is the absolute difference between the analyst's final firm-year forecast and the actual firm-year EPS value. *Clement-Tse Herding Rate* is constructed in a similar fashion using *Herding Rate*. Following Clement and Tse (2005), an analyst's forecast revision in the firm-year is classified as *non-herding* only 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. Otherwise, the revision is classified as *herding*. *Herding Rate* is the ratio of herding revisions to the total number of revisions made by an analyst 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 in Columns (1) to (4) are *All-Star* and *Analyst Centrality*. In Columns (5) and (6), we add *Consensus Deviation* and its respective interactions with *All-Star* and *Analyst Centrality* as key independent variables. *All-Star* 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 2 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. To facilitate presentation, we do not present the estimated coefficients of control variables in the table. Apart from the addition of *All-Star,* Columns (1) and (2) follow the specification in *Table 2A Column 1*; Columns (3) and (4) follow the specification in *Table 3A Column 1*; Columns (5) and (6) follow the specification in *Table 4 Column 1*. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 6. (Continued)

	(1) Clement-Tse Accuracy	(2) Clement-Tse Accuracy	(3) Clement-Tse Herding Rate	(4) Clement-Tse Herding Rate	(5) Abs. CAR [-1, +1]	(6) Abs. CAR [-1, +1]
Specification	Table 2A Column 1		Table 3A Column 1		Table 4 Column 1	
All-Star [A]	0.011***	0.011***	-0.016***	-0.016***	0.071***	0.068***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.024)	(0.024)
Analyst Centrality [B]		0.017***		-0.017**		0.003
		(0.005)		(0.007)		(0.015)
Consensus Deviation [C]					0.900***	0.855***
					(0.076)	(0.075)
[A] X [C]					0.645***	0.616***
					(0.167)	(0.167)
[B] X [C]						0.315***
						(0.086)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	33,596	33,596	32,435	32,435	250,922	250,922
R-squared	0.146	0.146	0.061	0.061	0.287	0.287
Firm fixed effects	No	No	No	No	Yes	Yes
Industry-Week fixed effects	No	No	No	No	Yes	Yes
Firm-Week cluster	No	No	No	No	Yes	Yes
Analyst-Firm cluster	No	No	No	No	Yes	Yes

Table 7.
The Causal Effect of Increases in Analyst Centrality on Forecast Accuracy: Subsequent Analyst Employment after Brokerage Closures

This table presents results from a difference-in-difference model with multiple groups, and multiple shocks across time. We define shocks to analyst employment as brokerage closures (Kelly and Ljungqvist, 2012) from years 2000 to 2007. For each closure event, we retain all analysts who subsequently find employment in another brokerage. Furthermore, these analysts are required to cover the same firm before and after the closure event. Therefore, our unit of observation in the model is an analyst-firm. The treatment in the model is *Analyst Centrality Up*, an indicator variable which equates to unity if the average post-closure *Analyst Centrality* is higher than the average pre-closure *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 [-5 years, +5 years] window centered on the closure event. The dependent variable in this table is *Clement-Tse Accuracy* (see Clement and Tse, 2005)*.* In each firm-year, the maximum (*Max FE*) and minimum (*Min FE*) values of *Forecast Error* are computed. *Clement-Tse Accuracy* of an analyst in a firm-year is the ratio of the difference between *Max FE* and *Forecast Error* to the difference between *Max FE* and *Min FE.* Hence, *Clement-Tse Accuracy* is bounded between 0 and unity, with higher values reflecting a higher degree of forecast accuracy. The key independent variables are the temporal leads (*Pre-Treatment*) and lags (*Post-Treatment*) of the treatment. We also include analyst and analyst-firm time trends in some specifications to help control for confounding trends. Analyst time trends are *Log* (*General Experience*), *Brokerage Size*, *Firm Breadth*, and *Industry Breadth*. Analyst-firm time trends are *Horizon*, *Revision Frequency*, and *Log* (*Firm Experience*). Detailed definitions of the variables are in the Appendix. To facilitate presentation, coefficient estimates of analyst and analyst-firm time trends are not presented. Robust standard errors are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5% and 10% levels respectively.

Table 7. (Continued)

	(1) Clement-Tse Accuracy	(2) Clement-Tse Accuracy	(3) Clement-Tse Accuracy
Leads and Lags of Increases in Analyst Centrality:			
Pre-Treatment 5	12.743	6.492	10.937
	(0.139)	(0.150)	(0.087)
Pre-Treatment 4	11.699	8.629	11.364
	(0.093)	(0.091)	(0.105)
Pre-Treatment 3	15.089	11.839	17.278
	(0.098)	(0.091)	(0.112)
Pre-Treatment 2	21.187**	18.178**	19.536**
	(0.067)	(0.056)	(0.078)
Pre-Treatment 1	1.938	0.761	6.595
	(0.043)	(0.036)	(0.052)
Post-Treatment 1	13.388**	12.475**	16.565***
	(0.049)	(0.043)	(0.044)
Post-Treatment 2	16.504**	16.655**	22.791**
	(0.058)	(0.057)	(0.086)
Post-Treatment 3	15.733*	15.995*	21.135**
	(0.079)	(0.082)	(0.082)
Post-Treatment 4	15.374**	15.396**	14.508*
	(0.046)	(0.057)	(0.071)
Post-Treatment 5	16.915**	17.184**	18.111**
	(0.063)	(0.063)	(0.062)
Other Covariates:			
Analyst time trends	No	Yes	Yes
Analyst-Firm time trends	No	No	Yes
Observations	1,148	1,148	1,148
R-squared	0.251	0.254	0.339
Analyst-Firm dummies	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes
Analyst cluster	Yes	Yes	Yes
Industry cluster	Yes	Yes	Yes