Social Networks under Stress: Specialized Team Roles and Their Communication Structure

DANIEL M. ROMERO, University of Michigan

BRIAN UZZI, Northwestern University

JON KLEINBERG, Cornell University

Social network research has begun to take advantage of fine-grained communications regarding coordination, decision-making, and knowledge sharing. These studies, however, have not generally analyzed how external events are associated with a social network's structure and communicative properties. Here, we study how external events are associated with a network's change in structure and communications. Analyzing a complete dataset of millions of instant messages among the decision-makers with different roles in a large hedge fund and their network of outside contacts, we investigate the link between price shocks, network structure, and change in the affect and cognition of decision-makers embedded in the network. We also analyze the communication dynamics among specialized teams in the organization. When price shocks occur the communication network tends not to display structural changes associated with adaptiveness such as the activation of weak ties to obtain novel information. Rather, the network "turtles up." It displays a propensity for higher clustering, strong tie interaction, and an intensification of insider vs. outsider and within-role vs. between-role communication. Further, we find changes in network structure predict shifts in cognitive and

This research was sponsored by the Northwestern University Institute on Complex Systems (NICO), the U. S. Army Research Laboratory and the U. S. Army Research Office under grant number W911NF-09-2-0053, Defense Advanced Research Projects Agency grant BAA-11-64, a Simons Investigator Award, a Google Research Grant, a Facebook Faculty Research Grant, an ARO MURI grant "QUANTA: Quantitative Network-based Models of Adaptive Team Behavior", and NSF grants IIS-0910664 and IIS-1617820. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. government. All our summary statistics and programs are available on request. Northwestern University IRB Approved the Study (#STU00200578). All data were previously collected, accessed from the firm's archive, anonymized, and involved no manipulation or interaction with subjects. Subjects knew the data were collected and available for research purposes. The firm provided the data under written agreement for research purposes contingent on firm's identifying characteristics remaining confidential and anonymous. This article is an extension of a article published in the Proceedings of the 25th ACM International World Wide Web Conference (Romero et al. 2016). This extended version of the article includes a new analysis and set of results on the communication dynamics among employees of the hedge fund with different titles and roles: portfolio managers, traders, and analysts. Our results show that the daily distribution of active IM users by role changes significantly with stock price changes, exhibiting more active analysis and less active traders. This is consistent with the firm spending more effort planning and analyzing their decision making than executing their decisions when price shocks occur. Additionally, we observe new evidence that is consistent with the firm "turtling up"—when price shocks occur, there is an increased volume of communication among employees of the same role. These new findings provide a more detailed view of the social dynamics in the organization during times of stress by taking into account the roles of the employees of the firm and their inter- and intra-group communication.

Authors' addresses: D. M Romero, 3340 North Quad, 105 S. State Street, University of Michigan, Ann Arbor, MI 48109; email: drom@umich.edu; B. Uzzi, 2211 Campus Drive, Room 5329, Northwestern University, Evanston, IL 60208; email: uzzi@kellogg.northwestern.edu; J. Kleinberg, Computing and Information Science, Gates Hall, Cornell University, Ithaca, NY 14853; email: kleinber@cs.cornell.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

1559-1131/2019/02-ART6 \$15.00

<https://doi.org/10.1145/3295460>

affective processes, execution of new transactions, and local optimality of transactions better than prices, revealing the important predictive relationship between network structure and collective behavior within a social network.

CCS Concepts: • **Information systems** → **Web log analysis**; **Social networks**; *Data mining*; • **Networks** → **Online social networks**; • **Human-centered computing** → **Collaborative and social computing theory, concepts and paradigms**; **Collaborative and social computing devices**;

Additional Key Words and Phrases: Social networks, organizations, temporal dynamics, collective behavior

ACM Reference format:

Daniel M. Romero, Brian Uzzi, and Jon Kleinberg. 2019. Social Networks under Stress: Specialized Team Roles and Their Communication Structure. *ACM Trans. Web* 13, 1, Article 6 (February 2019), 24 pages.

<https://doi.org/10.1145/3295460>

1 INTRODUCTION

The emergence of detailed, time-resolved data on large social networks makes it possible to study their structure and dynamics in new ways. Work has identified stable network structures, communities, influentials in networks, information flows, networked teams (Easley and Kleinberg 2010; Jackson 2010; Lazer et al. 2009; Newman 2010, 2003; Platt and Romero 2018), and the temporal evolution of large social networks (Ahn et al. 2007; Chakrabarti and Faloutsos 2012; Ke and Ahn 2014; Kossinets et al. 2008; Kossinets and Watts 2006; Leskovec et al. 2007; Newman et al. 2011; Peel and Clauset 2015; Saavedra et al. 2011; Viswanath et al. 2009; Weng et al. 2013).

Despite the richness of the available datasets and the computational techniques for analyzing them, work has generally not focused on a crucial feature of dynamic behavior—how changes in the social network’s structure are linked to external shocks. Little research thus far has examined how social networks operate in a reactive capacity, as they respond to such stimuli from their broader environments. By external shocks we mean events that are extreme relative to average events, or unexpected (Gilbert 2005). Such questions are critical to understanding a social network’s capacity to respond to uncertainty when the membership of the network remains essentially fixed. The real-life conditions reflecting these dynamics are varied, including the response of intelligence and law enforcement personnel to emergency situations such as terrorist attacks (Butts et al. 2007), health professionals responding to outbreaks (Dutta and Rao 2015), organizations facing sudden competitive threats or “normal accidents” (Perrow 2011), or the sudden government censorship of popular websites (Zhang et al. 2017). The effect has been to leave a set of basic questions largely unanswered. How does a network respond to external shocks? What can shock-related responses reveal about adaptive collective behavior?

We address these questions in the context of an organization’s social network (Adamic and Adar 2005; Brass et al. 2004; Diesner et al. 2005; Dodds et al. 2003; Kilduff and Brass 2010; Klimt and Yang 2004; Uzzi 1996). We consider the complete instant-messaging corpus, including content, among the decision-makers of a hedge fund and their outside contacts. The analysis of financial investment behavior is of interest because of its criticality to the economy and society (Fenton-O’Creevy et al. 2011; Gell-Mann and Lloyd 1996; Lo and Repin 2002; Malkiel and Fama 1970; Preis et al. 2013; Saavedra et al. 2013).

We focus on a few central questions. First, we study the link between environmental changes—encoded in the movement of the prices of specific stocks at particular points in time—and changes in the network structure, specifically the subgraph consisting of instant messages about these particular stocks during the times in question. For financial services firms, price shocks can represent threatening and stressful external stimuli. Experiments measuring the stress level of traders using conductance technology show that price changes lead to involuntary and significant physiological

changes indicative of stress. Shifts in prices cause traders to have increased heart rates and electrodermal responses relative to control conditions (Lo and Repin 2002). Qualitative observations of traders echo experimental results. In the industry, the VIX, a widely used measure of market price volatility, is referred to as the “Fear Index” (Whaley 2000).

Following sociological theory, one expectation for the association between price shocks and network change is that price shocks are associated with a propensity to activate connections that improve access to novel and diverse information and manage risk (Burt 2009; Granovetter 1973). In this case, it would be expected that actors in the network would rely relatively more on weak ties, network outsiders, and relationships with low closure. Conversely, shocks may be associated with reflective network structure and behavior (Staw et al. 1981). Following this view, expectations are that actors disproportionately favor contacts repeatedly seen in the past, in-group rather than out-group relations and highly interconnected contacts. Our analysis supports the latter expectation. In the face of shocks, we find that networks exhibit high levels of clustering, out-group communication, and strong ties.

Where our first question explores the link between external factors and network structure, our second main question asks whether these changes to the network structure can yield additional insight into the behavior of the organization—beyond what is provided by the external changes themselves. Specifically, if we seek to predict how the organization will respond to a price change, does knowing properties of the network structure improve the performance of the prediction even if we already have access to the time series of stock prices?

We show that knowledge of the network provides strong improvements in prediction of collective behavior for both individual-level actions and firm-level actions. We evaluate individual-level actions by analyzing the way traders express themselves in the collection of instant messages, using a standard set of linguistic measures. These measures show that when a stock’s price changes significantly (in either the positive or negative direction), the messages associated with that stock display increased emotion and cognitive complexity. While these changes can already be inferred to a limited extent from the price changes alone, we are able to predict their direction and magnitude much more effectively when we incorporate network-level features into the prediction. We find analogous effects in the context of firm-level actions, where we analyze the trading decision employees make for the firm. We introduce a simple measure of local optimality in the trading price and show that network features provide significantly improved performance in predicting whether the firm will perform a locally optimal action. We find a similar pattern when we attempt to predict whether trades that have not been observed for a number of days are suddenly traded.

Taken together, our findings suggest a key role for electronic communication networks in the analysis of external shocks and events affecting an organization: The network displays a consistent set of changes in response to these external events, emphasizing strong ties and clustered structures; and knowledge of the network structure provides significant leverage in predicting the organization’s collective behavior.

2 TURTLED-UP AND OPEN NETWORKS

Current frameworks suggest different conjectures for relating social network structure to external risks. Networks could open up structurally, with actors tapping acquaintances who are most likely to provide novel information and diverse perspectives on problem-solving (Burt 2009). Experiments show that persons facing threatening job changes, for example, disproportionately turn to their weak ties and reduced their triadic closure (Menon and Smith 2014; Smith et al. 2012), broadening their options and access to novel job information (Granovetter 1973). Conversely, external shocks could be associated with networks of highly clustered, frequent, and familiar

Fig. 1. Structure of *open* (a) and *turtled up* (b) networks. Turtled up networks exhibit a propensity toward strong ties over weak ties, high clustering, and insider vs. outsider links. Networks tend to turtle up on days with large or unexpected stock price changes.

relationships (Coleman 1988; Granovetter 1985) that promote trustworthiness but narrow social cognition (Coleman 1988; Ellis 2006; Granovetter 1973; Menon and Smith 2014).

These frameworks predict two vastly different network structures. In the context of the hedge fund, we define two types of structures that represent each of the two conjectures: an *open* network and a *turtled-up* network. We characterize turtled-up network structures as containing many triangles of relationships (friends of friends linked to each other), a high proportion of strong ties (high-frequency relationships), and a small fraction of edges that connect employees of the hedge fund to outsiders. Open network structures are the opposite—they contain few triangles of relationships, have relatively high numbers of weak ties, and high volume of communication between hedge fund employees and outsiders. Figure 1 illustrates key differences between closed or turtled up networks and open networks.

While we have defined turtled-up and open networks using features that are directly relevant to the setting we are studying (e.g., low volume of outside communication), we note that the concept can be applied to other settings where a network of communication can be constructed. In the case of the hedge fund, we could conveniently identify outsiders and measure their communication with the inside of the firm as well as communication between actors that play different roles within the firm. In other settings, it may be necessary to identify communities using community detection techniques (Newman 2006) to measure the volume of communication across different groups, but a similar characterizations of open and turtled-up structures can be made.

3 DATA

Our data include the complete instant-messaging communication history among the personnel at a hedge fund from January 2010 to December 2011. According to the firm, instant-message is the main way in which employees discuss day-to-day decision making around trading both inside the firm and with outside contacts. The data consist of approximately 22 million instant messages (IMs) sent by 8,646 people, of whom 184 are employees of the hedge fund, and the rest are outside contacts. There are three types of titles among the employees of the hedge fund: portfolio manager (pm), trader, and analyst. Portfolio managers are in charge of overseeing investments, developing trading strategies, and placing large orders of trades; traders are in charge of executing orders made by pm's and have some flexibility on the timing of the execution; and analysts provide analytic support to both portfolio managers and traders. Our dataset contains 63 pm's, 26

traders, and 95 analysts. We use the full content of each message and unique person and time of transmission identifiers to analyze the IMs. We know each trade completed at the firm, the stock symbol involved, and date and time of execution. All transactions in our data were executed by the fund's employees, not algorithms. On average, 559 transactions were performed daily. We merged the above data with public data on daily stock prices. All of our data analysis was consistent with guidelines from the relevant Institutional Review Board.

4 MEASURES

4.1 Network

IMs define a network of information exchange where nodes are communicating actors and IMs between actors define edges. Our primary interest is in the subgraphs of the larger network designated by mentions of particular stocks at particular points in time. For a stock symbol s and a day d , we define an undirected graph $G_{s,d}$ as follows: For each IM between company insiders that mentions stock symbol s on day d , we include the two participants of the IM as nodes in $G_{s,d}$, and we join them by an edge. We eliminate parallel edges, so that two nodes in $G_{s,d}$ will have at most one edge between them. Multiple messages between pairs were examined when relevant. $G_{s,d}$ thus consists of all employees who mentioned stock s on day d , with messages containing s on day d forming the edges between these people.

We define an edge to be *internal* if both nodes are company insiders and *border* if one node is a company outsider. All IMs include one company employee, because communication between company outsiders cannot be recorded. Our definition of $G_{s,d}$ uses only internal edges; the analogous construction using both internal and border edges results in a larger graph $G_{s,d}^+$ that contains $G_{s,d}$.

The notation and terminology used in the analysis for graph size and connectivity and quantifications of *turtled up* or *open* networks is as follows. Let $N_{s,d}$ and $E_{s,d}$ be the number of nodes and edges respectively in $G_{s,d}$. We normalize relevant measures in relation to comparable quantities in the data; in particular, for a function $f(G_{s,d})$ that we compute on $G_{s,d}$, we will define $\nu(f(G_{s,d}))$ to be the ratio of $f(G_{s,d})$ to the average value of $f(G_{s',d'})$ for all pairs (s', d') such that $N_{s',d'} = N_{s,d}$ and $d' < d$. This lets us discuss whether $f(G_{s,d})$ is large or small relative to other graphs from previous days with the same number of nodes. We can define the analogous normalization based on the number of edges; we write $\epsilon(f(G_{s,d}))$ to denote the ratio of $f(G_{s,d})$ to the average value of $f(G_{s',d'})$ for all pairs (s', d') such that $E_{s',d'} = E_{s,d}$ and $d' < d$. Furthermore, since some stocks tend to be more popular in traders' conversations than others, we normalize the number of nodes. We define $\bar{N}_{s,d}$ as the ratio of $N_{s,d}$ to the average values of $N_{s,d'}$ for all $d' < d$. Since the earliest part of the dataset does not have sufficient history to have a robust set of baseline graphs, we do not make measurements for the first 30 days of the study period.

Beyond the number of nodes and edges, a basic quantity is connectivity—whether $G_{s,d}$ is connected or whether most nodes belong to few connected components. We write $L_{s,d}$ for the fraction of nodes in the largest connected component of $G_{s,d}$, and $K_{s,d}$ for the minimum number of components required to account for at least 90% of nodes in $G_{s,d}$.

A measure for capturing triangles of relationship is the *clustering coefficient*; formally the fraction of pairs of a node's neighbors that are connected by an edge. We define $C_{s,d}$ as the average clustering coefficient over all nodes in $G_{s,d}$. To capture strength of ties we use frequency of communication, which is a relevant measure for the structure of a social network (De Choudhury et al. 2010). For each node x in $G_{s,d}$, we consider the set of all nodes y with whom x has participated in any messages on days $d' < d$, and we sort these nodes y in descending order by the number of such messages they have participated in x . We define $U_{x,d,\alpha}$ as the highest α fraction of this sorted list:

Table 1. Description of Network, Role Communication, and Price Change Features Used in the Study

	Feature	Notation	Mean	SD
Network Features	Nodes	$\tilde{N}_{s,d}$	1.76	1.02
	Edges	$v(E_{s,d})$	1.41	0.87
	Frac. Large Con. Comp.	$v(L_{s,d})$	1.40	0.74
	Num. Comp. for 90% of Nodes	$v(K_{s,d})$	1.36	0.63
	Clustering Coeff.	$v(C_{s,d})$	1.30	7.74
	Strength of Ties	$S_{s,d,0.1}$	0.51	0.34
	Ratio of Border Edges	$O_{s,d}$	0.86	0.26
Role Communication Features	Ratio Analysts	$P_{s,d,Analys}$	0.22	0.27
	Ratio Traders	$P_{s,d,Trader}$	0.34	0.22
	Ratio Portfolio Managers	$P_{s,d,PM}$	0.44	0.21
	Comm. Traders/Traders	$DC_{s,d}^{Trader, Trader}$	-0.11	0.13
	Comm. Traders/Analysts	$DC_{s,d}^{Trader, Analyst}$	-0.01	0.05
	Comm. Traders/PMs	$DC_{s,d}^{Trader, PM}$	0.00	0.03
	Comm. Analysts/Analysts	$DC_{s,d}^{Analyst, Analyst}$	-0.05	0.11
	Comm. Analysts/PMs	$DC_{s,d}^{Analyst, PM}$	0.00	0.20
	Comm. PMs/PMs	$DC_{s,d}^{PM, PM}$	-0.16	0.13
	Comm. Different Roles	$DC_{s,d}^{R_1 \neq R_2}$	0.25	0.29
Price Change Features	Price Change (%)	$\Delta_{s,d}$	0.02	3.66
	Absolute Price Change (%)	$ \Delta_{s,d} $	2.08	3.01

the α fraction of x 's communication partners from days prior to d , measured by communication volume these are x 's strongest edges. We quantify whether $G_{s,d}$ favors strong or weak ties using the measure $S_{s,d,\alpha}$, the fraction of edges (x, y) for which $y \in U_{x,d,\alpha}$.¹ To the extent that closed and open structures are related to an actor's reach for information across boundaries (Burt 2009), we define the *openness* $O_{s,d}$ as the fraction of edges in $G_{s,d}^+$ that are border edges.

In summary, we compute the following daily features, which we refer to as the *network features* for the article: number of nodes, number of edges, fraction of nodes in the largest connected component, minimum number of connected components required to account for 90% of the nodes, average clustering Coefficient, fraction of strong ties, and openness. As described above, some of these measured are normalized using baselines from graphs from previous days. Table 1 shows the mean and standard deviation of these features.

4.2 Communication between Actors of Different Roles within the Firm

Our network measures describe the general structure of who talks to whom and the rates of internal and external communication. However, they do not describe how the volume of communication is distributed internally among employees with different roles. We expect that the role of the employees who are active on a given day and the role of those they choose to communicate with carries a signal about the dynamics of decision making within the firm. To capture this, we define a set of features that describe communication between roles. First, to capture who is active on a given day, we define $P_{s,d,R}$ as the fraction of employees of role R who mention stock s on day d out of all employees who mention stock s on day d .

¹We use $\alpha = 0.1$ for the analysis presented in this article, but other reasonably small values of α produce qualitatively similar results.

Next, we consider the flow of communication between employees of different roles in the hedge fund. The intuition is that communication between two groups can signal a specific type of activity in the trading dynamics of the firm. For example, communication between analysts and traders could signal that analysts have found potential trading strategies or insight on how to execute trading orders and are conveying this information to traders; communication among analysts can signal that the firm is actively planning and developing strategies; and communication among pm's could signal that the firm is discussing high level decision-making outside of any particular portfolio.

Given that we already measure the fraction of active employees and that this fraction can impact the rate of inter role communication (e.g., more active analysts and more active traders would naturally lead to more communication between analysis and traders), we must control for the composition of active employees when we measure the rate of communication between the different subgroups. Thus, given a stock s and a day d , we compute the expected fraction of IMs between employee roles R_1 and R_2 , $EC_{s,d}^{R_1,R_2}$, assuming IMs are generated from a random sender to a random receiver, while fixing the number of employees from each group. That is, $EC_{s,d}^{R_1,R_2} = P_{s,d,R_1} P_{s,d,R_2}$ if $R_1 \neq R_2$ and $EC_{s,d}^{R_1,R_2} = 2P_{s,d,R_1} P_{s,d,R_2}$ if $R_1 = R_2$. We then compute the difference between the actual and the expected rate of communication between the groups—we let $AC_{s,d}^{R_1,R_2}$ be the fraction of IMS between employees with role R_1 and R_2 and let $DC_{s,d}^{R_1,R_2} = AC_{s,d}^{R_1,R_2} - EC_{s,d}^{R_1,R_2}$.

Finally, we define a similar measure as above, but we now measure the fraction of IMs between employees of *different roles*, while controlling for the expected fraction given the number of active employees of each role. We let $DC_{s,d}^{R_1 \neq R_2}$ be this measure.

In summary, our *role communication* features consist of the fraction active employees of each role, the normalized fraction of communication between roles, and the normalized fraction of communication between employees of different roles. Table 1 shows the mean and standard deviation of these features.

4.3 Shocks

To define the extremeness and unexpectedness of price shocks, we defined for each stock s and day d , $a_{s,d}$ and $b_{s,d}$ to be the opening and closing prices respectively of stock s on day d . We define $\Delta_{s,d} = \frac{b_{s,d} - a_{s,d}}{a_{s,d}}$ as the proportional change in the price of s on day d .

Some price changes are disruptive—the price change magnitudes are greater than recent changes, and therefore they can be characterized as unexpected. They intuitively correspond to shocks (Gilbert 2005). We operationalize the notion of a “shock” as follows: a stock-day pair (s, d) is an x -shock if $|\Delta_{s,d}| > x$ and $|\Delta_{s,d'}| \leq x$ for $d' = d - 3, d - 2, d - 1$. That is, (s, d) is an x -shock if stock s 's price change on day d was higher than x , and its price change was lower than x on the previous three days. We investigate how continuous values of $\Delta_{s,d}$ and discrete x -shocks relate to the properties of $G_{d,s}$.

5 FINDINGS

5.1 Price Changes and Network Dynamics

We observe a substantial relationship between changes in stock price and changes in network structure. Figure 2 shows how each network feature—clustering, tie strength, and the relative intensity of insider to outsider contacts—changes with $\Delta_{s,d}$. The horizontal axis of each figure shows the percentage price change Δ . When $\Delta > 0$ ($\Delta < 0$): the vertical axis indicates the mean measure for networks $G_{s,d}$ such that $\Delta_{s,d} \geq \Delta$ ($\Delta_{s,d} \leq \Delta$). When $\Delta = 0$, the vertical axis indicates the mean

Fig. 2. Change in stock prices vs. network structure for networks with two or more nodes.

measure for all networks. Throughout the article, figures also include 95% confidence intervals. Aggregating changes in network properties with respect to price changes over all stocks and days, we observe that as price changes increase in intensity, the network exhibits a higher clustering coefficient.² Price changes are also related to increasing levels of tie strength.³ Border edges also drop as a percentage of edges in the face of price changes. Note that the trends show in Figure 2 are clear as we move away from a change in price of 0% to about 5%. However, the error bars become very large for changes price that are larger than 5%. Thus, the change in network structure as prices change from 5% to 10% is far less clear. These findings reveal that decision-makers in the network tend to *turtle up* their communication in the face of stress, rather than *open up*.

Changes in connectivity are also significant but substantively slight. The fraction of nodes in the largest component increases slightly and the number of components needed to account for 90% of the nodes decreases slightly with price changes. We also find a trend with respect to network size—as the price changes increase, the number of nodes ($\bar{N}_{s,d}$) and edges increases ($v(E_{s,d})$). These effects are reported for graphs $G_{s,d}$ with at least two nodes; all results are consistent if we restrict to graphs containing a larger number of nodes. We note that changes in network structure are symmetric with respect to the direction of the price change, which is curious in nature given that one might expect that price increases and decreases would be associated with different structural changes—an issue we address below when we examine the actual changes in actors' cognition and affect.

We further examined the results in Figure 2 by disaggregating the analysis on a stock-by-stock and industry-by-industry basis subject to control variables. We run OLS regressions of the form $f(G_{s,d}) = \beta_0|\Delta_{s,d}| + \beta_1f(G_{s,d-1}) + \beta_2f(G_{s,d-2}) + \beta_3VIX + \alpha_1D_s^1 + \dots + \alpha_SD_s^S + r_{s,d}$, where S is the number of stocks in our dataset. For each stock s' the indicator function $D_s^{s'}$ is defined as follows:

$$D_s^{s'} := \begin{cases} 1 & : s = s' \\ 0 & : s \neq s' \end{cases}.$$

Including this indicator function as an independent variable measures the association between the change in stock prices, $|\Delta_{s,d}|$ and the network properties $f(G_{s,d})$ on a stock-by-stock basis (i.e., fixed effects model (Bhargava et al. 1982)). We included fixed effect variables for day of the week to control communication patterns explained by the day of the week (e.g., earning announcements and quarterly reports typically are made on specific days for specific stocks); a variable measuring

²We observe consistent results when we measure clustering while controlling for nodes ($v(C_{s,d})$) and edges ($\epsilon(C_{s,d})$).

³Repeating the test with various values of α showed that the results are similar regardless of the choice of α .

Table 2. Results Form OLS Regressions of the Form
 $f(G_{s,d}) = \beta_0|\Delta_{s,d}| + \beta_1f(G_{s,d-1}) + \beta_2f(G_{s,d-2}) + \beta_3VIX + r_{s,d}$

Independent Variables	Model 1 $f = \text{Nodes}$	Model 2 $f = \text{Clustering}$	Model 3 $f = \text{Perc. border edges}$	Model 4 $f = \text{Strength of ties}$
Stock price change	0.0540***	0.1590***	-0.0009***	0.0010***
$f\text{-lag}(-1)$	-0.0009***	0.0089***	0.1415***	0.0448***
$f\text{-lag}(-2)$	0.0644***	0.0026	0.0959***	0.0270***
VIX	-0.0034***	-0.0213***	-0.0006***	-0.0001
5 Day of week fixed effects	Y	Y	Y	Y
Stock fixed effects	Y	Y	Y	Y

Regressions include fixed effects at the stock and day of the week level. Each column represents a regression with f as the dependent variable and the rows show the value and significance of the independent variable. Asterisks indicate coefficient significance (***) p value < 0.0001).

Table 3. Results Form OLS Regressions of the Form
 $f(G_{s,d}) = \beta_0|\Delta_{s,d}| + \beta_1f(G_{s,d-1}) + \beta_2f(G_{s,d-2}) + \beta_3VIX + r_{s,d}$

Independent Variables	Model 1 $f = \text{Nodes}$	Model 2 $f = \text{Clustering}$	Model 3 $f = \text{Perc. border edges}$	Model 4 $f = \text{Strength of ties}$
Stock price change	0.0557***	0.1925***	-0.0004***	0.0009***
$f\text{-lag}(-1)$	0.2579***	0.0209***	0.2008***	0.0729***
$f\text{-lag}(-2)$	0.1277***	0.0146***	0.1555***	0.0542***
VIX	-0.0038***	-0.0230***	-0.0006***	-0.0001
5 Day of week fixed effects	Y	Y	Y	Y
Industry fixed effects	Y	Y	Y	Y

Regressions include fixed effects at the industry and day of the week level. Each column represents a regression with f as the dependent variable and the rows show the value and significance of the independent variable. Asterisks indicate coefficient significance (***) p value < 0.0001).

the daily market wide volatility using the VIX index; and the lagged values $f(G_{s,d-1})$ and $f(G_{s,d-2})$ to control for possible autocorrelation in the time series of network properties.

We run regressions with the described stock level fixed effects and an additional version where we use a fixed effects model at the industry level instead of the stock level. Tables 2 and 3 show the value and significance of each independent variable using fixed effects at the stock and industry level, respectively. The results confirm the analyses presented in Figure 2. Changes in price are associated with increases in network size, clustering, and strength of ties, and decreases in the percentage of outside contacts. A post-test analysis for autocorrelation of the residuals was also run. For each stock s , we ran a separate regression of the form $f(G_{s,d}) = \beta_0|\Delta_{s,d}| + \beta_1f(G_{s,d-1}) + \beta_2f(G_{s,d-2}) + r_{s,d}$. Based on the Durbin–Watson test (Durbin and Watson 1951), there was no statistical evidence of positive or negative serial correlation in 99.2% and 99.9% of the stocks, respectively.

The relationship between unexpected price changes and network structure further substantiate the inference that social networks in the face of uncertainty, as measured here, exhibit a propensity to turtle up. We compare the value of each feature of $G_{s,d}$ when (s,d) is an x -shock and when it is not. Further, we measure the number of days a graph feature takes to return to its mean value following a shock. Figure 3 shows the values of network features in graphs $G_{s,d}$ on the day of an x -shock ($x = 5\%$, as defined above), and then on subsequent days until the feature approximately returns to its average over all networks. The insets in Figure 3 compare network values on shock

Fig. 3. Network structure on days with a x -shock and on the days following a an x -shock. Insets compare network properties on shock and non-shock days. $x = 5\%$ was used to define x -shocks.

and non-shock days. We find that all features are significantly different when there are shocks, and the direction of the differences mirror the previous Figure 2. The same patterns arise when different values of x are used. Our analysis of network recovery indicates that most network properties return to their average value one or two days after the shock, suggesting that normalization in relation to these shocks is relatively fast acting.

These results suggest that social networks in the face of stress broadly exhibit features associated with turtling up rather than opening up. Social networks become more intensely interconnected among third parties, rely more on information from strong rather than weak ties, and disproportionately attend to organizational insiders.

Fig. 4. Percentage of words reflecting affective and cognitive processed by sender's role.

Fig. 5. Price changes vs. percentage of active employees by role.

5.2 Price Changes and Communication between Actors of Different Roles within the Firm

We now explore the differences in communication patterns across employees of different roles. We begin by comparing the type of language used by each employee type. To infer affect and cognition of the actors in our network, we use the Linguistic Inquiry and Word Count (LIWC) dictionary, which identifies words in the content of communications (i.e., instant messages) that reflect affective and cognitive states.⁴ Affect includes positive emotion, negative emotion, anxiety, anger, and sadness. Cognition includes insight, causation, discrepancy, tentative, certainty, inhibition, inclusive, and exclusive. the LIWC dictionary is well validated and regularly used (Coviello et al. 2014; Kramer et al. 2014; Romero et al. 2015).

Figure 4(a) and (b) shows the percentage of words related to cognitive and affective processes used by employees of each role. We observe that there are significant differences—portfolio managers exhibit the highest rate of cognitive language while analysts exhibit the highest rate of affective language. Traders use cognitive and affective language at the lowest rates.

Given the observed differences in language between employees of different roles, which suggest that they serve different functions in the social dynamics of the firm, we now investigate the participation rate of employees of each role in IM communication when stocks change in price.

Figure 5 shows $P_{s,d,R}$ as function of the change in price of stock s on day d . The x -axis shows the relative price change Δ . In Figure 5(a), when $\Delta > 0$ ($\Delta < 0$) the y -axis indicates the fraction of users who mentioned s on day d who are portfolio managers such that $\Delta_{s,d} \geq \Delta$ ($\Delta_{s,d} \leq \Delta$). Figure 5(b) and (c) show the corresponding plots for traders and analysts, respectively.

⁴See <http://www.liwc.net> for more information on LIWC 2007.

Fig. 6. Price changes vs. fraction of active employee pairs by role controlling for expected fraction.

We observe a very clear trend. When the price of a stock s changes, regardless of direction, there is an increase in the fraction of analysts and a decrease in the fraction of traders who mention s . The fraction of pm's who mention s does not appear to depend significantly on the extent of the price change. If we think of the role of traders as actors in the firm who execute trading orders, then analysts help traders and pm's with planning and strategy, and pm's oversee and manage the full decision-making pipeline, then these findings are consistent with the following. When price shocks occur, there is a larger emphasis on *planning*, represented by the increased participation of analysts, and a reduced emphasis on *acting*, represented by the decreased participation of traders. When price changes are small, we observe the opposite effect. Portfolio managers are expected to be involved in both acting and planning, and we do not see a significant change in their participation with price shocks. This finding is also consistent with the increase in cognitive language when price changes occur, which we will discuss in Section 5.3. An increase in cognitive language and participation of analysts in IM communication suggests that the firm is emphasizing planning when faced with large changes in stock prices.

We now turn to the volume of inter-role communication. Figure 6 shows how $DC_{s,d}^{R_1, R_2}$ changes with price changes in s for all pairs of employee roles. The x-axis shows the relative price change Δ . When $\Delta > 0$ ($\Delta < 0$) the y-axis shows $DC_{s,d}^{R_1, R_2}$ such that $\Delta_{s,d} \geq \Delta$ ($\Delta_{s,d} \leq \Delta$). We observe some interesting patterns. There is clearly a sharp increase in communication between traders when prices change (regardless of direction) and an increase in pm communication when prices go up (Figure 6(b) and (a)). There is also a slight, but observable, increase in analyst communication when prices go up (Figure 6(c)). Furthermore, communication between pm's and traders and between analysts and traders decreases when prices change in both direction (Figure 6(d) and (f)). Finally, there is no significant change in communication between analysts and pm's with price changes (Figure 6(e)).

These patterns suggest an interesting trend—prices appear to generate a surge in communication between employees with the same role. Figure 7 shows $DC_{s,d}^{R_1 \neq R_2}$ as a function of the change in price of s on day d . The resulting trend confirms that price changes result in an overall increased rate of communication between employees of with the same role.

Fig. 7. Price changes vs. percentage of active employee pairs with different role controlling for expected fraction.

Taking this finding together with the results in previous sections, we can interpret it as another way in which the firm turtles up. Not only do we observe an increased rate of communication with internal members of the firm, strong ties, and people with high overlap in neighbors, but we also observe an increased rate of communication within groups—that is, between people of the same role—and a corresponding reduction in communication between people with different roles.

5.3 Cognitive and Affective Content

In the previous section, we identified changes on network structure that result from changes in the stock market. Given this observation, it is important to know whether network structure has an impact on other activities of the firm. To explore this question, we set up prediction tasks to compare the predictive power of changes in the stock market and the structure of the communication network on various activities of the firm. Our goal is to assess whether network structure impacts these activities above and beyond what we can already predict purely from changes in the stock market.

We begin by considering the language used in the messages exchanged by the employees of the firm. Language can often signal a person’s psychological and emotional states (Tausczik and Pennebaker 2010). In the setting of a hedge fund, cognition and affect are particularly relevant as they are known to shape decision making (Agarwal et al. 2012; Kahneman 2011; Tetlock 2007). If changes in network structure are associated with actual changes in behavior, then changes in network structure should predict changes in actors’ cognition and affect.

Before incorporating the role of network features, we investigate the relationship between stock prices and IM content by measuring how the usage of LIWC categories varies with changes in the stock prices. In particular, research holds that price changes are stressful for the employees of the firm from a cognitive and emotional perspective (Lo and Repin 2002). For each instant message among insiders, we computed the percentage of words in each LIWC category. Formally, given a stock-day pair (s, d) and the IMs that mention s on day d , we test whether the percentage of LIWC words in the IMs varies with $\Delta_{s,d}$. The x -axis in Figures 8 and 9 show the relative price change Δ . When $\Delta > 0$ ($\Delta < 0$) the y -axis indicates the mean percentage of words in each LIWC category in IMs mentioning s on day d such that $\Delta_{s,d} \geq \Delta$ ($\Delta_{s,d} \leq \Delta$). When $\Delta = 0$, the y -axis indicates the mean value for all IMs.

We find that price changes are associated with expected changes in cognition and affect. Figure 8(a) shows that as price changes intensify, either upward or downward, decision-makers’ communications express higher levels of cognitive processes—presumably because they face greater risk and complexity in their judgments. Figure 8(b)–(i) shows how the different

Fig. 8. Price changes vs. percentage change in words reflecting various cognitive processes.

subcategories of cognitive processes change with price. The general trend is that the change in cognitive processes is symmetric in the direction of the price change. This symmetric pattern is consistent with the structural changes in Figure 2.

By contrast, affect expressed during ups and downs in price are asymmetric. Figure 9(a) shows that expressed affect surrounding a stock during its price changes differ depending on whether prices rise or fall. Figure 9(b)–(f) shows the change in the affective subcategories with changes in prices. Words related to positive emotions are used more when stock prices rise while words related to negative emotions, anger, anxiety, and sadness are used more often when stock prices drop. One explanation of this is that while funds can make as much money when stock prices fall

Fig. 9. Price changes vs. percentage change in words reflecting various affective processes.

as when they rise by selling short, more negative affect is expected when prices drop, because falling prices generally sound alarms among retail investors who put their capital in the hands of hedge fund decision-makers and take money out of the market when stock prices fall.

If message content varies with price changes, then does network structure further predict the psychology of traders?

To address this question, we formulate the following prediction task. For each stock s and each LIWC category C , let C_s be the fraction of all IMs containing stock symbol s that include a word from category C . For each day d , let $C_{s,d}$ be the fraction of all IMs containing stock symbol s on day d that include a word from category C . We say that the pair (s, d) *conforms* to category C if $C_{s,d} > C_s$ —in other words, if words from category C are used at a higher rate on day d than is typical for stock s .

We use binary classifiers to predict whether each pair (s, d) conforms to each of the cognitive and emotional LIWC categories using the properties of the network and the stock price changes as predictors. We run the prediction test using three feature sets: the network features (as described in Section 3.1), price change features (the percent change in price of stock s on day d , $\Delta_{s,d}$, and its absolute value, $|\Delta_{s,d}|$), and the two sets of features together. Each set of features includes lagged values for 3 days before day d . In this analysis, we leave out the role communication features as predictors as we do not expect these variables to be related to linguistic patterns. To test the accuracy of the classifiers, we split time into 100-day bins, using each bin as a test set and all the previous bins as training data. We balance the testing and training data by including all positive examples and selecting a random sample of negative examples of the same size as the set of control cases. Table 4 shows the number of cases in each class before being balanced. We use seven binary

Table 4. Number of Positive and Negative Examples of Pairs (s, d) That Conform to Each LIWC Category

Category	Num. Positive	Num. Negative
Affective processes	27,229	65,704
Anger	3,047	89,886
Anxiety	1,570	91,363
Causation	14,794	78,139
Certainty	13,113	79,820
Cognitive processes	50,088	42,845
Discrepancy	19,281	73,652
Exclusive	27,383	65,550
Inclusive	36,033	56,900
Inhibition	7,556	85,377
Insight	19,663	73,270
Negative emotion	11,478	81,455
Positive emotions	23,551	69,382
Sadness	4,531	88,402
Tentative	27,784	65,149

Fig. 10. Prediction accuracy of logistic regression classifier for LIWC categories using network features, price change features, and all features combined. Category key: Affective processes (A), Anger (B), Anxiety (C), Causation (D), Certainty (E), Cognitive processes (F), Discrepancy (G), Exclusive (H), Inclusive (I), Inhibition (J), Insight (K), Negative emotion (L), Positive emotions (M), Sadness (N), Tentative (O).

classifiers: Random Forest, Linear Discriminant Analysis, Quadratic Discriminant Analysis, Naive Bayes, Decision Trees, Support Vector Machines, and Logistic Regression. Choice of classifier does not change the results; logistic regression results are presented.

Figure 10 shows the accuracy of the logistic regression classifier using each of the feature sets for the affective and cognitive categories, and for the positive-emotion, negative-emotion, and insight subcategories. Network properties alone provide significantly better predictions of the psychology of the decision-makers than the price changes alone. Furthermore, the figure indicates that

combining the two types of features does not yield a significant improvement over the network features alone. This pattern suggests that network structure is more predictive of a network's collective affect and cognition than the price shock itself.

5.4 Decision-Making Behavior

We now turn to the question of whether the structure of the communication networks provides insight about the trading behavior of the firm. Since trading requires coordination and communication among the employees, it is likely that some of the high level trading decisions of the firm are latent in the structure of the firm's communication. We look at two important aspects of the decision making process when the firm decides to make a trade—the quality of its trading decision in terms of timing, and the decision to begin trading a stock that has not been traded for a long period of time. Stock prices are crucial in both trading time and in deciding to trade a new stock, hence when we test the predictive value of the networks, we always compare it against the stocks' price changes. We find that the network structure is indeed predictive of these trading behaviors, above and beyond what is predicted by price changes alone.

5.4.1 Predicting Performance. We begin by measuring whether the timing of each trade was *locally optimal*. To measure local optimality of each transaction, we ask whether the firm would have benefited from waiting until the next day to make the transaction. For a stock s traded on day d we let $p_{s,d}$ denote the price of the stock for that transaction. We let $p_{s,d}^{max}$ and $p_{s,d}^{min}$ be the maximum and minimum price of stock s on day d . If a stock s was bought on day d at price $p_{s,d}$ and the maximum price of s the following day ($p_{s,d+1}^{max}$) was less than $p_{s,d}$, then the company would have benefited from waiting until the next day to buy stock s . Following this reasoning, we label each *buy* transaction $(s, d, p_{s,d})$ as *locally suboptimal* if $p_{s,d} > p_{s,d+1}^{max}$ and *locally optimal* otherwise. Similarly, we label each *sell* transaction $(s, d, p_{s,d})$ as locally suboptimal if $p_{s,d} > p_{s,d+1}^{min}$ and locally optimal otherwise. We observe that about 20% and 23% of all buy and sell transaction are locally suboptimal.

We expect that the decision process of the firm does not always involve optimizing exactly the day on which to make a transaction. Instead, the firm may focus most of its efforts on other objectives—trying to decide *which* stocks to trade, *how much* to trade, or on long term profit as opposed to small gains. To validate our measure of local optimality, we test whether the company's trades reflect an effort to trade on a locally optimal day by comparing the firm's real performance with their performance if the trades occurred on a random day.

We create a random set of transactions by taking each transaction in the dataset and generating an alternative transaction of the same stock and same number of shares, but on a randomly selected day. The price of the alternative transaction is selected uniformly at random between the minimum and maximum price on the selected day.

The number of locally suboptimal transactions in the random set is 2% higher than in the actual set of transactions. While 2% may not initially appear to be very large, it is easier to assess the size of the difference when we compute the loss in profit that these locally suboptimal transaction generate. For each locally suboptimal transaction $(s, d, p_{s,d})$, we let the number of shares traded be $V_{s,d,p}$. The loss generated by this transaction is $V_{s,d,p} * |p_{s,d} - p_{s,d+1}|$. The difference in total loss generated by the set of actual transactions and the set of random transaction is about \$40 million. This shows that the company performs much better than they would if they traded the same stock and the same number of shares but on a randomly selected day. This validates our measure; we now test if it is related to features of the networks.

We use our set of classifiers to predict whether transactions are locally optimal using the properties of the network, role communication features, and the price change features as predictors.

Fig. 11. Prediction accuracy of logistic regression classifier when predicting if a trade is locally optimal vs. the number of consecutive days the stock was traded. The curves show the results using network features, price change features, and all features combined.

For each transaction $t_{s,d}$ of stock s on day d , we use the features of the graph $G_{s,d}$ (as described in Section 3.1), the change in price $\Delta p_{s,d}$, and the absolute change in price $|\Delta p_{s,d}|$ to predict whether $t_{s,d}$ is locally optimal. Each feature also includes lagged values for 3 days before day d .

As we did before, we split time into 100 bins, and use each bin as a test set and all the previous bins as training data. We also use a balanced set of positive and negative examples. Since some stocks are traded very often and others are rarely traded, we further split the prediction task into subtasks by the number of consecutive days on which transactions occurred. Letting T_k be the set of transactions that have occurred on at least the previous k days, we run the classifiers on each set of transactions T_k for $k = 0 \dots 6$.

Figure 11 shows the accuracy of the Logistic Regression classifier for each set of transactions T_k . The accuracy of the classifiers increases with the number of days of consecutive transactions, suggesting the local optimality of routine transactions is easier to predict than that of unexpected transactions. We also observe that the network features are significantly more predictive than the price changes and role communication features. Furthermore, combining all the features only marginally improves the accuracy of the classifier with network features alone. This shows that the local optimality of the firm's decisions are better aligned with the properties of their communication network than the changes in stock prices.

Feature Analysis. Having found that network structure can be predictive of locally optimal performance in the hedge fund, we now explore which network features are predictive of high performance. In particular, we analyze the coefficients of the network features in the Logistic regressions we run on above. Since we found the network features outperform other features regardless of the number of days of consecutive transactions, we aggregate the coefficients of all the regressions. We focus our analysis on features that measure whether the network turtles up—clustering coefficient, fraction of border edges, and strength of ties.

Figure 12 shows the average value of the coefficients of these three features. The figure includes the coefficients of the feature based on the network from the day when the performance was measured, as well as the network from one and two days prior. This allows us to explore the relationship between performance and network structure on the days following the transactions. We find a very clear trend. On the day of the transaction, a turtles up network—high clustering, low fraction of border edges, and high tie strength, is associated with lower performance. However, on the two days prior to the transaction, the trend reverses—clustering and strength of ties have a positive (but small) coefficient and fraction of border edges has a positive coefficient but much

Fig. 12. Mean coefficient value of network features in Logistic regression predicting performance.

lower than on the day of the transaction. This suggests that on days prior to the transaction a more turtled up network is predictive of high performance. Put together, these trends suggest that high performance occurs when employees first discuss and coordinate their trading decisions with their most frequent contacts, people in their clusters, and insiders on the days leading up to their trades, and open up their communication to less frequent contacts, people from other clusters, and outsiders as they get closer to trading. In other words, high performance occurs when networks are initially turtled-up, but slowly open up as actions become more imminent.

5.4.2 Predicting Sudden Trading. We observed that it's easier to predict the local optimality of stocks that are traded consecutively. We now turn to a second, more basic question to study the firm-level actions with respect to trading: Are we able to predict whether or not a stock is traded on a given day? Our first observation is that many stocks are traded at very high frequencies, and hence the best predictor of a stock being traded on given day is whether it was traded during the previous few days. Given this, we pose the task of predicting *new transactions*—trades of stocks that have not been traded for a given number of days prior to the transaction being considered.

Engaging on a new transaction requires coordination and discussion among the employees to decide on details such as how much to trade and in what type of position to take on the stock. While volume and position also need to be decided for sequential transactions, they are often determined by the trend of transactions from previous days. We expect that when traders are planning on executing a new trade, they will need to exchange information about a particular stock for several days before actually executing a trade. The nature of such discussion may impact the structure of the communication network, and thus, such structure may be predictive of a new trade occurring.

We let NT_k^d be the set stocks that have not been traded for k weeks prior to day d . We say that a stock s is k -unobserved on day d if it is in the set NT_k^d . That is, s is k -unobserved on day d if it has not been traded during the past k weeks preceding d . Note that a k -unobserved stock on day d is also k' -unobserved on day d for all $k' < k$. We say that all stocks are 0-unobserved on all days.

We use our binary classifiers to predict whether stocks that are k -unobserved on day d will be traded on day d . The setup of the prediction task is the same as in the last two sections—we split time into 100 bins and use each bin as a test set and all the previous bins as training data, and we use a balanced set of positive and negative examples. We use network features, role communication features, and price change features (as described in Section 3.1), but in this case we also include features that indicate whether the stock was traded on the 7 days prior to the k -week period when the stock was not traded. For example, when we predict if a 1-unobserved stock s is traded on day d , we know s was not traded on days $d - 1, \dots, d - 7$, but we include features that indicate if s was traded on days $d - 8, d - 9, \dots, d - 14$.

Figure 13 shows the accuracy of the classifiers for different values of k . When $k = 0$ we are predicting whether stocks are traded on day d , regardless of whether they have been traded on the

Fig. 13. Prediction accuracy of logistic regression classifier when predicting if k -unobserved stocks are traded vs. k . The curves show the results using network features, price change features, and all features combined.

Fig. 14. Mean coefficient value of network features in Logistic regression predicting sudden trading.

previous days. We observe that with no minimum window on the time since the last trade, the 7-day trading history of the stocks provides over 80% accuracy, and adding information about the network, role communication, or stock prices does not significantly increase the accuracy. However, as k increases and we begin predicting trading of stocks that have not been traded for some time, the accuracy of the stocks' trading history alone drops significantly. When we add the price change and role communication features to the trading history features, we only observe a small increase in accuracy. However, adding the network features to the trading history features yields a large increase in accuracy. For example, for $k = 4, \dots, 9$ the accuracy of the price change, role communication and trading history features is always less than 60%. However, the accuracy of the network and trading history features is over 69%. Finally, combining all the features does not significantly improve the performance of the network features alone. This pattern is consistent with what we found when predicting affective and cognitive content and trading performance.

Feature Analysis. We now explore the relationship between network features and sudden trading. Figure 14 shows the average coefficient of clustering, strength of ties, and fraction of border edges in the Logistic regressions that predict sudden trading. The figure includes the coefficients based on the network from the day when sudden trading is being predicted, as well as the network from one and two days prior. We observe that a low fraction of border edges and low tie strength on the day of trading are predictive of sudden trading. The trend on days prior to trading is different. Both fraction of border edges and tie strength have coefficients of much smaller magnitude, suggesting that they are not as predictive on days preceding sudden trades. Clustering coefficient is never a significant predictor of sudden trading. Overall, we do not find that sudden trading is predicted by how turtles up the networks become on the day of the trades or on the days prior to it. Instead, a

combination of high within-company and weak tie communication is predictive of sudden trades. A possible explanation is that, when preparing to explore new stocks that have not been traded in a while, employee prefer to discuss their strategy more with their colleagues than with outsiders, but within the company, they explore the expertise of their weak connections.

6 CONCLUSION

Network science has examined the reaction of networks to internal stresses, particularly nodal loss, but has given considerably less attention to the relationship between external shocks in a network of stable members (Saavedra et al. 2008). Using all the instant messages among stock traders in an investing organization and their outside contacts to define structural, cognitive, and affective properties of their social network, we found that shocks—in the form of extreme price changes—were not associated with conventional adaptive network responses to uncertainty such as activating weak ties and external sources to obtain novel and diverse information. Rather, the network turtled up. Relationships within the network favored strong ties, high clustering, company insiders, and colleagues with the same role within the company. We also find evidence that emphasis within the company during times of stress is on planning and analyzing decision making rather than acting and trading, which could reveal the mechanism responsible to the observed changes in network structure.

Implications of this work relate to networks facing disruptive environments and “normal accidents” (Perrow 2011). One often cited benefit of networks is that they are more agile than hierarchies and more coordinated than markets in solving collective action problems (Powell 2003; Uzzi 1996). While one basis for this benefit has been to show that institutions organized as networks do better than hierarchies in turbulent environments, there has been little work on the actual network dynamics that arise in organizations facing environmental disruptions. Indeed, the 2013 DARPA “robolympics” challenge was instituted to investigate whether machines can work hand in hand with humans to address problems of organizing in the face of “normal accidents” better than humans can on their own. Our case indicated that networks facing shocks turtle up rather than open up. Nevertheless, we find that networks are relatively elastic—turtled up states return to normal states relatively fast. Whether the combination of these changes leads to better or worse performance relative to a set criterion beyond the metrics we investigated and, provided by theory, is a logical research extension and a broad goal for improving knowledge about the scientific functionality and practical management of social networks.

Moreover, we find that the network structure is diagnostic of important patterns of behavior—including the emotional and cognitive content of individual communications, local optimality of transactions, and the sudden execution of new transactions. It is noteworthy that the structure of the networks is more effective at predicting these behavioral patterns than the price changes in the market. This suggests that the network-level changes we observe are not simple offshoots of the underlying price changes, but instead that they carry additional rich information that can be used to analyze the organization’s behavior. Understanding the network’s reaction to shocks can thus be an important factor in understanding the organization more broadly.

The role of the network in this analysis raises a number of further open questions. In particular, while there is a clear relationship between changes in the market and changes in the network and inter- and intra-group communication dynamics, it is interesting to consider what the lower-level mechanisms that produce these effects might be. As we come to better understand the links between shocks, networks, communication dynamics, and behavior, we can thus arrive at a clearer picture of social and organizational dynamics in the context of their surrounding environments.

A limitation of this work is that our findings are limited to a single hedge fund. It would be interesting to investigate whether the patterns we observe in this firm generalize to other type

of organizations and other communities that face exogenous shocks. In particular, testing the universality of the turtling-up effect we observe on others types of social networks is a promising future direction as it would have consequences for the group's ability to process information and react to the shock. Settings such as natural disasters, financial crisis, and political instability are examples of settings where the communication structure of a group could be affected by the event. Furthermore, the types of activities and behaviors we analyze in this work are analogous to activities that can take place in many other settings. For instance, affect and cognitive processes, performance, and engaging in new or underexplored activities can be studied in other contexts beyond hedge funds or even organizations. Exploring whether communication structure is predictive of such behaviors is also an interesting direction of future research.

REFERENCES

- Lada Adamic and Eytan Adar. 2005. How to search a social network. *Soc. Netw.* 27, 3 (2005), 187–203.
- Sumit Agarwal, Ran Duchin, and Denis Sosyura. 2012. In the mood for a loan: The causal effect of sentiment on credit origination. (2012). SSRN: <https://ssrn.com/abstract=2141030> or <http://dx.doi.org/10.2139/ssrn.2141030>.
- Yong-Yeol Ahn, Seungyeop Han, Haewoon Kwak, Sue Moon, and Hawoong Jeong. 2007. Analysis of topological characteristics of huge online social networking services. In *Proceedings of the 16th International Conference on World Wide Web*. ACM, 835–844.
- Alok Bhargava, Luisa Franzini, and Wiji Narendranathan. 1982. Serial correlation and the fixed effects model. *Rev. Econ. Stud.* 49, 4 (1982), 533–549.
- Daniel J. Brass, Joseph Galaskiewicz, Henrich R. Greve, and Wenpin Tsai. 2004. Taking stock of networks and organizations: A multilevel perspective. *Acad. Manage. J.* 47, 6 (2004), 795–817.
- Ronald S. Burt. 2009. *Structural Holes: The Social Structure of Competition*. Harvard University Press.
- Carter T. Butts, Miruna Petrescu-Prahova, and B. Remy Cross. 2007. Responder communication networks in the World Trade Center disaster: Implications for modeling of communication within emergency settings. *Math. Sociol.* 31, 2 (2007), 121–147.
- Deepayan Chakrabarti and Christos Faloutsos. 2012. Graph mining: Laws, tools, and case studies. *Synth. Lect. Data Min. Knowl. Discov.* 7, 1 (2012), 1–207.
- James S. Coleman. 1988. Social capital in the creation of human capital. *Am. J. Sociol.* 94 (1988), S95–S120.
- Lorenzo Coviello, Yunkyu Sohn, Adam D. I. Kramer, Cameron Marlow, Massimo Franceschetti, Nicholas A. Christakis, and James H. Fowler. 2014. Detecting emotional contagion in massive social networks. *PLoS ONE* 9, 3 (2014), e90315.
- Munmun De Choudhury, Winter A. Mason, Jake M. Hofman, and Duncan J. Watts. 2010. Inferring relevant social networks from interpersonal communication. In *Proceedings of the 19th International Conference on World Wide Web*. ACM, 301–310.
- Jana Diesner, Terrill L. Frantz, and Kathleen M. Carley. 2005. Communication networks from the Enron email corpus? It's always about the people. Enron is no different? *Comput. Math. Organiz. Theory* 11, 3 (2005), 201–228.
- Peter Sheridan Dodds, Duncan J. Watts, and Charles F. Sabel. 2003. Information exchange and the robustness of organizational networks. *Proc. Natl. Acad. Sci. U.S.A.* 100, 21 (2003), 12516–12521.
- James Durbin and Geoffrey S. Watson. 1951. Testing for serial correlation in least squares regression. II. *Biometrika* 38, 1/2 (1951), 159–177.
- Sunasir Dutta and Hayagreeva Rao. 2015. Infectious diseases, contamination rumors and ethnic violence: Regimental mutinies in the Bengal Native Army in 1857 India. *Organiz. Behav. Hum. Decision Process.* 129 (2015), 36–47.
- David Easley and Jon Kleinberg. 2010. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press.
- Aleksander P. J. Ellis. 2006. System breakdown: The role of mental models and transactive memory in the relationship between acute stress and team performance. *Acad. Manage. J.* 49, 3 (2006), 576–589.
- Mark Fenton-O'Creevy, Emma Soane, Nigel Nicholson, and Paul Willman. 2011. Thinking, feeling and deciding: The influence of emotions on the decision making and performance of traders. *J. Organiz. Behav.* 32, 8 (2011), 1044–1061.
- Murray Gell-Mann and Seth Lloyd. 1996. Information measures, effective complexity, and total information. *Complexity* 2, 1 (1996), 44–52.
- Clark G. Gilbert. 2005. Unbundling the structure of inertia: Resource versus routine rigidity. *Acad. Manage. J.* 48, 5 (2005), 741–763.
- Mark Granovetter. 1985. Economic action and social structure: The problem of embeddedness. *Am. J. Sociol.* 91, 3 (1985), 481–510.
- Mark S. Granovetter. 1973. The strength of weak ties. *Am. J. Sociol.* 78, 6 (1973), 1360–1380.

- Matthew O. Jackson. 2010. *Social and Economic Networks*. Princeton University Press.
- Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Macmillan.
- Qing Ke and Yong-Yeol Ahn. 2014. Tie strength distribution in scientific collaboration networks. *Phys. Rev. E* 90, 3 (2014), 032804.
- Martin Kilduff and Daniel J. Brass. 2010. Organizational social network research: Core ideas and key debates. *Acad. Manage. Ann.* 4, 1 (2010), 317–357.
- Bryan Klimt and Yiming Yang. 2004. The enron corpus: A new dataset for email classification research. In *Machine Learning: ECML 2004 (ECML '04)*. Lecture Notes in Computer Science, Vol 3201. Springer, Berlin, Heidelberg.
- Gueorgi Kossinets, Jon Kleinberg, and Duncan Watts. 2008. The structure of information pathways in a social communication network. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 435–443.
- Gueorgi Kossinets and Duncan J. Watts. 2006. Empirical analysis of an evolving social network. *Science* 311, 5757 (2006), 88–90.
- Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock. 2014. Experimental evidence of massive-scale emotional contagion through social networks. *Proc. Natl Acad. Sci. U.S.A.* 111, 24 (2014), 8788–8790.
- David Lazer, Alex Sandy Pentland, Lada Adamic, Sinan Aral, Albert Laszlo Barabasi, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, et al. 2009. Life in the network: The coming age of computational social science. *Science* 323, 5915 (2009), 721.
- Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. 2007. Graph evolution: Densification and shrinking diameters. *ACM Trans. Knowl. Discov. Data* 1, 1 (2007), 2.
- Andrew W. Lo and Dmitry V. Repin. 2002. The psychophysiology of real-time financial risk processing. *J. Cogn. Neurosci.* 14, 3 (2002), 323–339.
- Burton G. Malkiel and Eugene F. Fama. 1970. Efficient capital markets: A review of theory and empirical work. *J. Finance* 25, 2 (1970), 383–417.
- Tanya Menon and Edward Bishop Smith. 2014. Identities in flux: Cognitive network activation in times of change. *Soc. Sci. Res.* 45 (2014), 117–130.
- Mark Newman. 2010. *Networks: An Introduction*. Oxford University Press.
- Mark Newman, Albert-Laszlo Barabasi, and Duncan J. Watts. 2011. *The Structure and Dynamics of Networks*. Princeton University Press.
- Mark E. J. Newman. 2003. The structure and function of complex networks. *SIAM Rev.* 45, 2 (2003), 167–256.
- Mark E. J. Newman. 2006. Modularity and community structure in networks. *Proc. Natl. Acad. Sci. U.S.A.* 103, 23 (2006), 8577–8582.
- Leto Peel and Aaron Clauset. 2015. Detecting change points in the large-scale structure of evolving networks. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI'15)*. AAAI Press, 2914–2920.
- Charles Perrow. 2011. *Normal Accidents: Living with High Risk Technologies*. Princeton University Press.
- E. Platt and D. M. Romero. 2018. Network structure, efficiency, and performance in WikiProjects. In *Proceedings of International AAAI Conference on Web and Social Media (ICWSM'18)*.
- Walter Powell. 1990. Neither market nor hierarchy: Network forms of organization. *Res. Organ. Behav.* 12 (1990), 295–336.
- Tobias Preis, Helen Susannah Moat, and H. Eugene Stanley. 2013. Quantifying trading behavior in financial markets using Google Trends. *Sci. Rep.* 3, Article 1684 (2013), 01684.
- Daniel M. Romero, Roderick I. Swaab, Brian Uzzi, and Adam D. Galinsky. 2015. Mimicry is presidential: Linguistic style matching in presidential debates and improved polling numbers. *Pers. Soc. Psychol. Bull.* 41, 10 (2015), 1311–1319.
- Daniel M. Romero, Brian Uzzi, and Jon Kleinberg. 2016. Social networks under stress. In *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 9–20.
- Serguei Saavedra, Kathleen Hagerty, and Brian Uzzi. 2011. Synchronicity, instant messaging, and performance among financial traders. *Proc. Natl. Acad. Sci. U.S.A.* 108, 13 (2011), 5296–5301.
- Serguei Saavedra, R. Dean Malmgren, Nicholas Switanek, and Brian Uzzi. 2013. Foraging under conditions of short-term exploitative competition: The case of stock traders. *Proc. R. Soc. B* 280, 1755 (2013) 2012–2901.
- Serguei Saavedra, Felix Reed-Tsochas, and Brian Uzzi. 2008. Asymmetric disassembly and robustness in declining networks. *Proc. Natl. Acad. Sci. U.S.A.* 105, 43 (2008), 16466–16471.
- Edward Bishop Smith, Tanya Menon, and Leigh Thompson. 2012. Status differences in the cognitive activation of social networks. *Organiz. Sci.* 23, 1 (2012), 67–82.
- Barry M. Staw, Lance E. Sandelands, and Jane E. Dutton. 1981. Threat rigidity effects in organizational behavior: A multi-level analysis. *Admin. Sci. Quart.* 26, 4 (1981), 501–524.
- Yla R. Tausczik and James W. Pennebaker. 2010. The psychological meaning of words: LIWC and computerized text analysis methods. *J. Lang. Soc. Psychol.* 29, 1 (2010), 24–54.

- Paul C. Tetlock. 2007. Giving content to investor sentiment: The role of media in the stock market. *J. Finance* 62, 3 (2007), 1139–1168.
- Brian Uzzi. 1996. The sources and consequences of embeddedness for the economic performance of organizations: The network effect. *Am. Sociol. Rev.* 61, 4 (1996), 674–698.
- Bimal Viswanath, Alan Mislove, Meeyoung Cha, and Krishna P. Gummadi. 2009. On the evolution of user interaction in facebook. In *Proceedings of the 2nd ACM Workshop on Online Social Networks*. ACM, 37–42.
- Lilian Weng, Jacob Ratkiewicz, Nicola Perra, Bruno Gonçalves, Carlos Castillo, Francesco Bonchi, Rossano Schifanella, Filippo Menczer, and Alessandro Flammini. 2013. The role of information diffusion in the evolution of social networks. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 356–364.
- Robert E. Whaley. 2000. The investor fear gauge. *J. Portf. Manage.* 26, 3 (2000), 12–17.
- Ark Fangzhou Zhang, Livneh Danielle, Ceren Budak, Lionel P. Robert Jr., and Daniel M. Romero. 2017. Shocking the crowd: The effect of censorship shocks on Chinese Wikipedia. In *Proceedings of the 11th International AAAI Conference on Web and Social Media*.

Received August 2017; revised August 2018; accepted November 2018