Journal of Financial Economics 113 (2014) 235–251

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Trading networks and liquidity provision ☆

Ethan Cohen-Cole [a,*], Andrei Kirilenko [b], Eleonora Patacchini [c,d,e]

[a] Econ One Research, 2040 Bancroft Way, Berkeley, CA 94704, United States
[b] Massachusetts Institute of Technology, Sloan School of Management, Cambridge, MA, United States
[c] Syracuse University, United States; University of Rome, La Sapienza, Italy
[d] Einaudi Institute for Economics and Finance, Italy
[e] Centre for Economic Policy Research, UK

ARTICLE INFO

Article history:
Received 12 January 2012
Received in revised form
25 July 2013
Accepted 27 October 2013
Available online 30 April 2014

JEL classification:
G10
C21

Keywords:
Financial interconnections
Contagion
Spatial autoregressive models
Network centrality
Trading limits

ABSTRACT

We study the profitability of traders in two fully electronic and highly liquid markets: the Dow and Standard & Poor's 500 e-mini futures markets. Using unique information that identify counterparties to a transaction, we show and seek to explain the fact that the network pattern of trades captures the relations between behavior in the market and returns. Our approach includes a simple representation of how much a shock is amplified by the network and how widely it is transmitted. This representation provides a possible shorthand for understanding the consequences of a fat-finger trade, a withdrawing of liquidity, or other market shock.

1. Introduction

In this paper, we analyze a unique data set of transactions from two financial futures contracts traded on the Chicago Mercantile Exchange (CME). The dataset contains information about transactions from the month of August for the September 2008 e-mini Standard & Poor's (S&P) 500 and Dow contracts. The data set has time-stamped transaction-level quantities, prices and counterparty identifiers for all transactions during August 2008. This includes more than seven million trades across more than 30 thousand accounts for the S&P 500 and more than one million trades across more than seven thousand accounts for the Dow.

The unique feature of the data is the availability of precise counterparty information. We are able to identify who traded, when, and with whom. We exploit this feature of the data to discuss the relation between the counterparty connections and a variety of market features of interest to financial economists. We characterize the topology of a trading network to help understand how traders' positions in the network influence their profitability and how shocks are transmitted across the market.

In spite of a growing literature on financial interconnections and a widespread belief in the importance of financial linkages, no consensus has been reached on how network

☆ We are very grateful to Ana Babus, Lauren Cohen, Ernst Eberlein, Rod Garratt, Thomas Gehrig, Dilip Madan, Todd Prono, Uday Rajan, Julio Rotemberg, and Jose Scheinkman, as well as conference participants at the Centre for Financial Analysis and Policy (CFAP) conference on Financial Interconnections and the FRAIS conference on Information, Liquidity, and Trust in Incomplete Financial Markets for constructive comments. Nicholas Sere and Kyoung-sun Bae provided research assistance. We are grateful for financial support from the Europlace Institute of Finance. All errors are our own.

* Corresponding author. Tel.: +1 301 541 7227.
E-mail addresses: ecohencole@gmail.com (E. Cohen-Cole), akirilenko@mit.edu (A. Kirilenko), epatacch@maxwell.syr.edu (E. Patacchini).

http://dx.doi.org/10.1016/j.jfineco.2014.04.007
0304-405X/

structure is related to liquidity or risk. A growing understanding exists of extreme cases such as repo runs (Brunnermeier and Pederson, 2009; Brunnermeier, 2009) or sequential default (Allen and Gale, 2000) or linkages in outcome across types of firms (Billio, Getmansky, Lo, and Pellizon, 2012), but these successes remain relatively rare in the literature.

We estimate the importance of market topology on trader-level returns using an approach that captures the correlation in returns between counterparties, the actual network topology of the entire market, and the importance of each transaction. Central to this approach is the introduction of the Bonacich centrality measure (Bonacich, 1987, 2007) to the financial economics literature. We believe that this network centrality measure is particularly salient in financial markets as it provides a way to understand the relative importance of direct and indirect links and thus helps explain the propagation of shocks in the system. As shown in Liu and Lee (2010), a close link exists between a spatial autoregressive model with network data and Bonacich centrality. This type of regression model captures recursively the network effects at any degree of separation (see also Lee, Liu, and Lin, 2010). In our application, a network regression model can explain more than 70% of the cross section of trader-level returns.

Why do networks emerge in this context? And why do they explain returns and shock amplification? We show that the (observed) network of trades is a characterization of the (unobserved) strategic interactions at work in the market. Traders with similar strategies trade amongst themselves as well as with others. As they do so, and form links with one another, correlation in trading strategies leads to a connection between strategies and network position. That is, certain types of traders are more frequently central in the network and other types are more frequently peripheral. A trader's network position thus predicts profitability and the network topology drives the transmission of shocks.

In Section 2, we present data and institutional features of the markets that we study. Section 3 contains the empirics of trader-level returns and highlights the role of network position for a better understanding of markets and trader profitability. Section 4 is devoted to describing our estimation results, and Section 5 discusses the causal nature of our empirical work. Section 6 extends the work to implement a policy experiment on the impact of trading limits. We discuss our contribution to the existing literature in Section 7 and conclude in Section 8.

2. Data and institutional features

Our data of interest are the actual trades completed on the CME for two contracts, the S&P 500 and Dow futures. The trades we observe are the result of orders placed by traders that have been matched by a trading algorithm implemented by the CME. Using the audit trail from the two markets, we uniquely identify two trading accounts for each transaction: one for the trader who booked a buy and the opposite for the trader who booked a sale. For these two markets, First In, First Out (FIFO) is used. FIFO uses price and time as the only criteria for filling an order: all orders at the same price level are filled according to time priority.

Each financial transaction has two parties, a direction (buy or sell), a transaction identification number, a time stamp, a quantity, and a price. We have transaction-level data for all regular transactions that took place in August 2008 for the September 2008 e-mini S&P 500 futures and the Dow futures contracts. The transactions take place during August 2008, when the markets for stocks underlying the indices are open. Both markets are highly liquid, are fully electronic, and have cash-settled contracts traded on the CME GLOBEX trading platform.

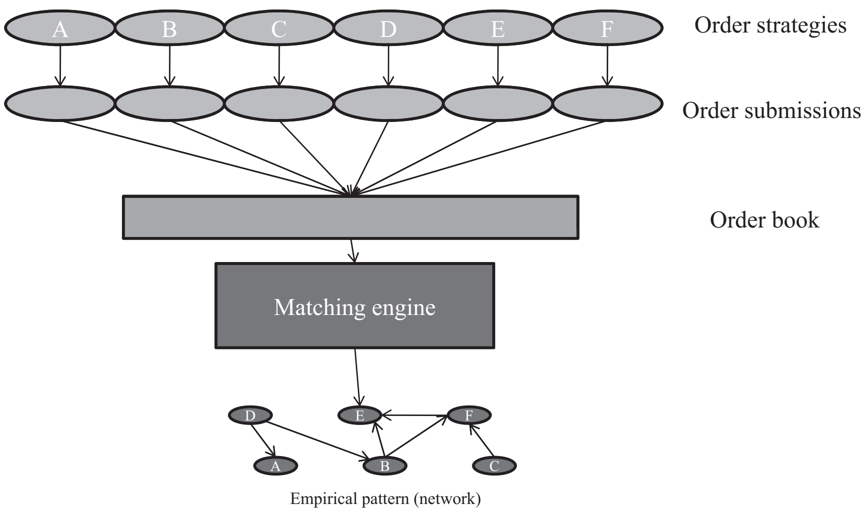

Fig. 1. Each node in the section labeled “order strategies” represents a single trader's plans for trading. The ovals beneath each trader, next to the label “order submissions,” represent actual placed orders. Below this, we denote with a box the complete order book. This is the aggregation at each time of all the orders submitted by traders. This order book is passed through the box beneath it, which we have labeled a “matching engine.” This computer matches orders based on price and time priority. Finally, beneath the matching engine, we provide a sample representation of the network patterns that could emerge from a set of six completed transactions.

Because these two markets are characterized by the use of price and time priority alone in determining trading partners, the only phenomenon that generates networks is the pattern of trading strategies that links traders with each other. Particular patterns of trading lead to different probabilities of being at the center or periphery of the network, as well as to distinct chances of trading with different types of counterparties. While, for each period, we do not observe the limit order book itself, we know that transactions occurred because market orders or limit orders were matched with existing orders in the limit order book. We can then trace the pattern of order execution – a trading network. Fig. 1 illustrates this pattern.

We empirically define a trading network as a set of traders engaged in conducting financial transactions within a period of time. The presence of a link is simply a reflection of the ex post realization of a cleared trade.

The choice of the period of time within which a network is defined is important, as it contains valuable information on the resulting network structure. With more time, more transactions are formed and more participants can form accurate beliefs about the valuation of a given asset.

Our approach is to define the network as a given number of transactions among traders that are either directly or indirectly linked. Then, throughout the remainder of the paper, we use a range of network densities to ensure that our results are robust to this choice. More specifically, we designate a network as a sequence of consecutive transactions. What we call sparse networks are defined as containing 250 transactions, moderately dense networks contain five hundred transactions, and dense networks contain one thousand transactions. Parsing trading activity in this way allows for avoiding variations in returns that could occur solely due to the ebbs and flows of trading.

While one could imagine alternate approaches, our evidence supports the above choice, i.e., defining networks as a given number of transactions.[1] Our results on the existence of network effects are strongly robust when we vary the number of transactions. As well, the fact that we find our chosen network definition has enormous empirical salience suggests that we have chosen a reasonable concept for the network. In addition, there is no reason to believe that an incorrect choice of network timing would lead to the spurious finding of a strong relation between networks and returns. The opposite is true: a randomly defined network shows no evidence of network effects by construction.

The networks that we define are distinct from one another over time. This occurs both because agents can be inactive in each time period and because their transactions are matched by the trading algorithm in each time period.

2.1. Returns and descriptive statistics

Each trader in the market that we study earns a return. For example, buying a contract for a price of \$1.00 and selling it for \$1.10 yields a profit of \$0.10 and a return of 10%. Because some positions are left open at the end of a given network time period, we report realized returns when positions clear during a network time period. When they do not clear, we report the mark-to-market returns for the trader in question.

Our S&P 500 futures data set consists of over 7,224,824 transactions that took place among more than 31,585 trading accounts. The DOW futures dataset consists of 1,163,274 transactions between approximately 7,335 trading accounts. We show in Table 1 some simple statistics of the data for each of the two markets that we analyze.

For each definition of networks, we compute returns for each trader, volumes for each trader, and the variance of returns across traders over the course of a trading day. Returns are shown as absolute levels of holding at the end of the time period, based on an initial investment of \$1.00. Thus, a return of one indicates that the trader broke even during the time period. Average returns vary from a loss of 4 basis points to a gain of 11 basis points. Individual-level results vary more widely. We report the returns unweighted by volume. The weighted average return across traders is, by construction in futures markets, equal to one. The average return across trading accounts is below one, suggesting that traders with high volume, on average, earn higher returns. To be more specific, we measure volume as the total number of contracts traded over the time period. In our data, these high volume traders are those that transact repeatedly, with regularly low order sizes. This would suggest that market makers and traders with regular interactions with the market are those that profit the most.

3. Empirics of trader-level returns

Before proceeding with the formal analysis, we provide a heuristic description of the market to illustrate the relation between trading behavior and the network topology. Table 2 reports information on transactions by type of trader. Three categories of traders are classified by frequency of trading. A high-frequency trading (HFT) group is composed of very active traders.[2] A group of irregular traders is referred to as hedge traders. Finally, a residual group with intermediate trading frequency is denoted as mid-size. These types also are mapped by type of activity. We have traders that only sell or only buy, which include most of those in the hedge group, as well as traders that both buy and sell. The HFT type are universally part of the buy and sell group. The table shows that the high-frequency group has very active, high-volume, high-profit traders; the hedge group collects irregular traders with larger individual transaction sizes that are not particularly profitable; and the mid-size group is composed of profitable traders that are regular participants in the markets. Based on trading frequency, these appear more likely to be broker-dealers than the HFT group. It confirms that more active traders tend to be more profitable in aggregate and transact in smaller sizes for each transaction.

[1] For example, an alternative would be to define the network based on some period of time or number of transactions beginning at a market shock, such as a significant price change.

[2] While we do not have a way to precisely identify traders as high frequency, we assume that any trader with more than five thousand transactions in a day is using some type of high-frequency trading.

Table 1
Summary statistics.

Sparse networks are defined as containing 250 transactions each, moderately dense networks as containing five hundred transactions each, and dense networks as containing one thousand transactions each. The table reports statistics from the Standard & Poor's (S&P) 500 e-mini futures market and from the Dow futures market. The columns report the mean, standard deviation, minimum, and maximum of each variable. Returns are defined as the gross return on an investment. Thus, a value of one indicates no change in value. Values greater than one are net gains and those less than one are net losses. For each density of network in each market, we report the average daily return as well as the total daily volume at the trader level. Thus, we report the mean return across individual-level traders, where for each trader we have calculated their own average return over the course of the trading day. These trader-level returns are unweighted by volume. Because the futures markets are zero-sum, volume-weighted returns are zero by construction. Volumes statistics are average daily volumes at the level of the trader. Standard deviations are measured as the variance over the returns at the trader level, again unweighted. Minimums and maximums are the smallest and largest for a trader on any day.

	Mean	Standard deviation	Minimum	Maximum
S&P 500 e-mini futures				
Sparse networks				
Average returns	0.98	0.01	0.97	1.05
Volume	5.94	4.98	1.00	1215
Moderately dense networks				
Average returns	0.96	0.02	0.96	1.09
Volume	5.73	7.90	1.00	1,518
Dense networks				
Average returns	0.92	0.02	0.96	1.106
Volume	5.32	12.68	1.00	2,060
Total number of trading accounts	31,585			
DOW futures				
Sparse networks				
Average returns	0.99	0.03	0.99	1.02
Volume	6.39	1.42	1.00	150
Moderately dense networks				
Average returns	0.98	0.05	0.98	1.03
Volume	6.33	2.60	1.00	190
Dense networks				
Average returns	0.95	0.07	0.98	1.04
Volume	5.91	4.86	1.00	341
Total number of trading accounts	7,335			

Table 2
Summary statistics by type of trader.

This table shows an example of trading patterns over a period of ten minutes. Each sample period has about 11 thousand transactions. Reported on are three types of traders: those with fewer than five trades per day, those with 50 to five thousand per day, and those with more than five thousand per day. The sample is divided into three types as well. Sell only is defined as traders that have no buy transactions during the ten minute period. Buy only is similarly defined. Buy and sell is defined as traders that conduct both a buy and a sell during the time period. Correlation for both panels is calculated as a simple correlation between the average returns for traders in the group and average returns for market markers. For example, the correlations in Panel B show that the market maker returns are positively correlated within its group and negatively correlated with sell only and buy only.

	Average transaction size	Average number of transactions per period	Average profit per transaction	Average correlation between group and most common counterparty
	(1)	(2)	(3)	(4)
Sample statistics by frequency of trade S&P 500 e-mini market				
Hedge	4.60	1.01	0.9448	−0.20
Mid-size	1.83	1,967.58	0.9999	0.29
High frequency	1.40	13,230.54	1.0001	0.40
Sample statistics by market role S&P 500 e-mini market				
Sell only	5.19	24.95	0.9999	−0.20
Buy and sell	1.69	59.80	1.0005	0.30
Buy only	3.79	11.37	0.9999	−0.20

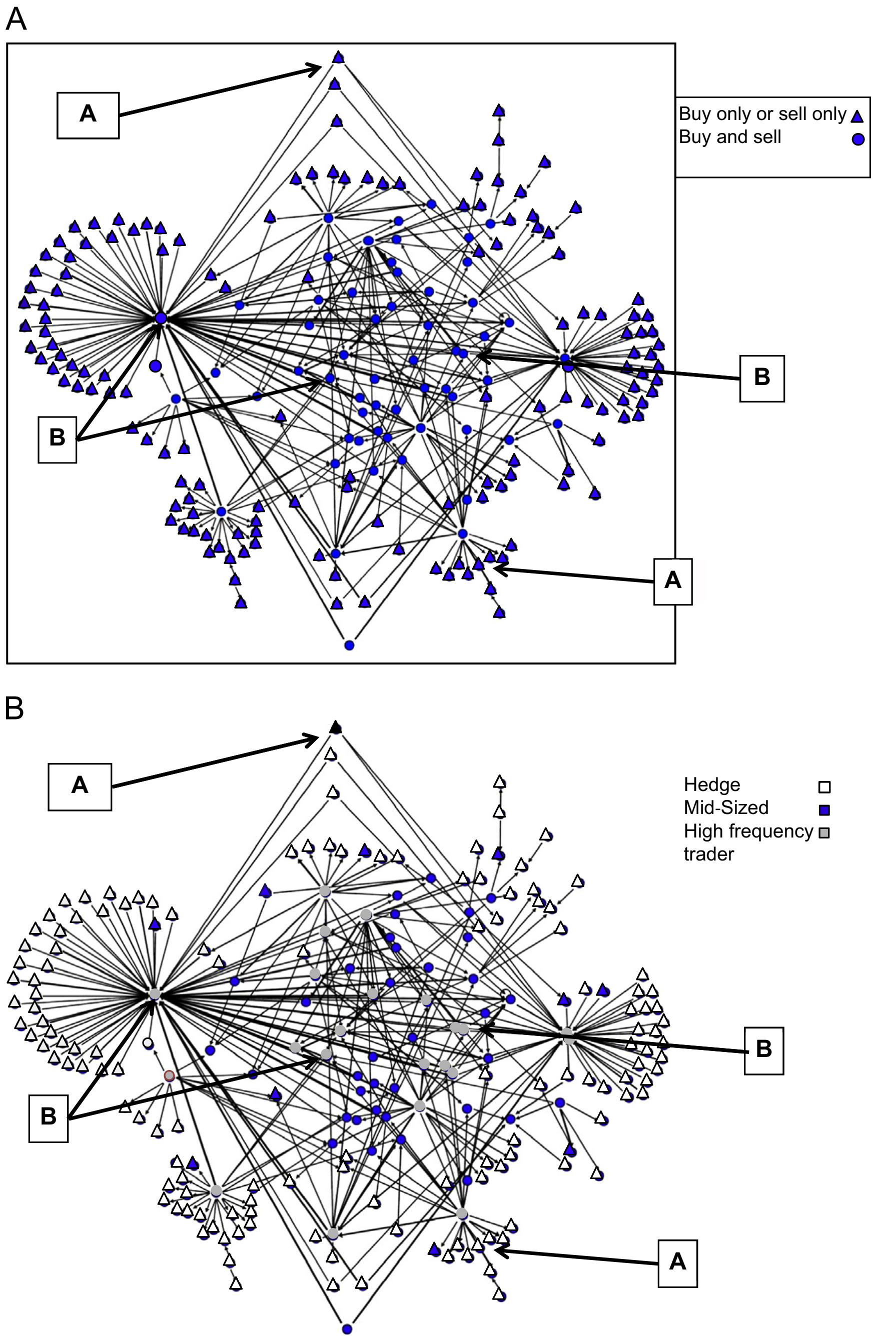

Fig. 2. Panel A and Panel B. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

These different trader types (or behaviors) are reflected in the network topology of the transactions. Panel A of Fig. 2 shows a representative network. Each node represents a trader and each arrow represents a trade, with the arrow pointing toward the buyer of a contract. We denote different trader types A and B. Traders A only sell or buy. These traders are examples of those with fundamental liquidity needs. We mark these traders with triangles. Each of these participate in the futures market by placing one-sided orders to either buy or sell contracts.[3] A separate set of

[3] Of particular emphasis, note that they do not trade with each other directly. If buy-only and sell-only traders did trade with each other directly, smaller networks of traders would be observed. The diagram shown is a fully connected, single networks. We do observe some small, isolated network, in our data, but they are very rare, a phenomenon that

traders, denoted B, implements rapid offers to buy and sell. These traders are indicated with circles in the diagram. These market makers typically trade with the objective to provide the liquidity needed by the traders A that have fundamentals demands. Because the buy-only and sell-only traders might not appear on the market at the same time, the liquidity providers can earn returns from them by being willing to transact when needed. The combination of the liquidity traders' actions can generate a diamond-shaped network pattern, illustrated in this figure. On one side, the buy-only traders buy when needed and, on the other, sell-only traders sell as needed. By being willing to buy and sell, the agents in the center can generate profits. The actions of B are known by market participants to be a profitable strategy. For this reason, more than a single agent conducts business in this fashion. As a result, the demands of A are not always intermediated by a single market maker. Often, there will be many before the contract reaches its end holder. That is, there are a large number of B traders. Effectively, the B traders hope to intermediate between the two traders with fundamental demand. In the process, they often end up trading with other agents of type B. In Panel B we color traders differently accordingly to the frequency of trades. White shows hedge (infrequent) traders, blue depicts mid-size traders and red denotes very active (HFT) traders. As can be seen from the picture, the group of very active traders (in red) includes a significant overlap with type B traders (circles in Panel A). It also appears that the group of irregular traders (white) is largely type A (triangles in Panel A).

Now, the relevant questions here are: How do these different types interact in the market? And what are the implications for returns?

The last column of Table 2 reports, for each type of trader, the average correlation in returns with the most common counterparty. In a single transaction, buyers and sellers have negatively correlated returns. When the market moves, one profits and the other loses. As a result, these buy-only and sell-only traders have negatively correlated outcomes. As it appears from Panel B, the transactions between the market makers and these buy-only and sell-only traders also are negatively correlated. Critically, within the buy-and-sell group, the returns are positively correlated. An explanation consistent with these features of the market is that the positive correlation emerges from the similarity in strategy between the market makers themselves. The intuition is simple. Traders enter the market each day with a set of trading strategies. These strategies can be either formal or informal, automated or manual. The market contains some of each. Among the formal strategies, for example, are high-frequency traders. These computerized high-frequency traders compose approximately one-third of volume (Kirilenko, Kyle, Samadi, and Tuzun, 2011) on any given day. The strategy of any given trader depends on the anticipated strategies of other traders as well as the observed actions during the day. As successful strategies become known, followers emerge and copy these strategies. As long as traders either use strategies that are broadly similar to each others' or condition their strategies on like information, their behaviors could be correlated in equilibrium and thus, too, in the observed data. These correlated bidding patterns lead to similarity in returns. Table 2 and Fig. 2 show that a substantial fraction of trades are intermediation ones.

To better understand the role of network structure in shaping returns and in propagating shocks, we introduce some network analysis tools.

3.1. A network regression model

Consider a model to explain the return, $r_{i,\kappa}$, of a trader, i, in network, k. We define returns as the log change in price over the time period defined as a network.

Assume that N traders are divided into $k=1,\ldots,K$ networks, each with n_k members, $i=1,\ldots,n_k$, $\sum_{k=1}^{K} n_k = N$.

Consider the influence on i of only a single other agent j. A basic specification would read

$$r_{i,\kappa} = \alpha_0 + \sum_{m=1}^{M} \beta^m x_{i,\kappa}^m + \gamma r_{j,k} + \upsilon_{i,\kappa}, \tag{1}$$

where x^m denotes a set of explanatory variables and r_j denotes the returns of the trading partner. So, an estimated coefficient γ greater than zero indicates that returns for trader j are positively correlated with returns for trader i. Extended to a simple network of three agents (i,j,s),

the equation becomes

$$r_{i,\kappa} = \alpha_0 + \sum_{m=1}^{M} \beta^m x_{i,\kappa}^m + \gamma_1 r_{j,k,d} + \gamma_2 r_{s,k,2d} + \upsilon_{i,\kappa}, \tag{2}$$

where the subscripts d and $2d$ indicate agents j and s at one node and two nodes distant from i, respectively. The coefficient γ_1 captures correlation in returns between directly connected traders, and γ_2 captures the correlation between agents further away in the network structure. These multiple steps are important. They are similar in spirit to multiple lags in a time series regression. The set x^m now also includes additional regressors for the characteristics of every other agent. Thus, as the number of agents increases and the network expands, we can continue to add regressors to the right-hand side of this specification for each agent and each degree of separation from agent i. Eventually, we add $n-1$ regressors for each degree of separation, leading to a complex specification that takes into account each type of influence of every agent on every other.

To include every other agent and every degree of separation, and to simplify notation, we can introduce a matrix that keeps track of the links between agents. This is an N-square adjacency matrix $G=\{g_{ij}\}$ whose generic element g_{ij} would be one if i is connected to j (i.e., interacts with j) and zero otherwise. Here $g_{ij}=1$ if trader i and j have concluded a transaction during a period of time and $g_{ij}=0$ otherwise. This matrix represents the interaction scheme of the traders in the market. The G matrix

(*footnote continued*)
points to the fact that market makers are prevalent in the data and intermediate most transactions.

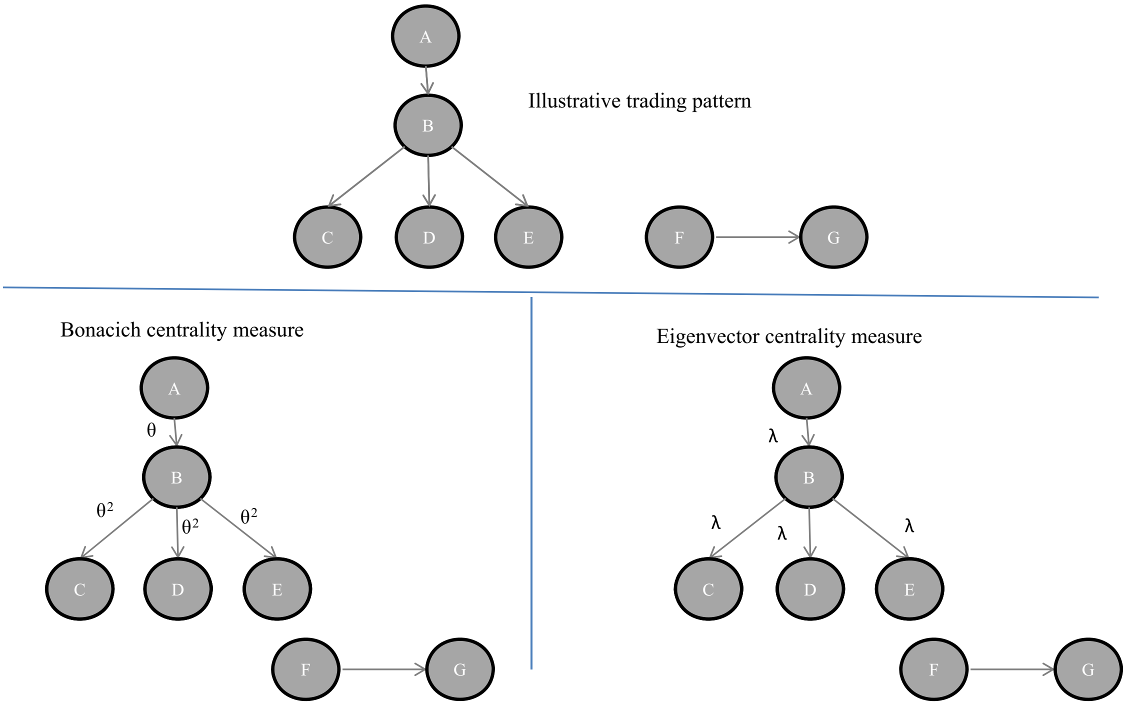

Fig. 3. This figure shows a stylized network of five traders (A, B, C, D, and E) as well as a network of two others (F and G). The lower-left figure shows how one would weigh the links to determine the centrality of trader A. The link to B has a weight of $\theta < 1$, the links from B to its parts, $\theta^2 < \theta$. The disconnected F and G traders receive no weight. On the lower right, we show this calculation using eigenvector centrality. To emphasize that the eigenvector method does not consider distance from A, we show that any connection to A receives the same weight. Again, F and G receive no weight.

associated with the simple network in the picture above is

$$G = \begin{array}{c|ccc} & i & j & s \\ \hline i & 0 & 1 & 0 \\ j & 1 & 0 & 1 \\ s & 0 & 1 & 0 \end{array} \tag{3}$$

indicating that i trades with j; s with j; and j with i and s. We use an undirected network for this analysis.[4]

Then, we can collapse the above specification with all traders at every level of interaction into the following specification:

$$r_{i,\kappa} = \alpha_0 + \sum_{m=1}^{M} \beta^m x_{i,\kappa}^m + \theta \frac{1}{g_{i.,k}} \sum_{j=1}^{n_\kappa} g_{ij,\kappa} r_{j,\kappa} + \upsilon_{i,\kappa}$$

$$\text{for } i = 1, \ldots, n_\kappa; \quad \kappa = 1, \ldots, K \tag{4}$$

[4] Our methodology can also be applied to directed networks. The fact that, using a directed network definition, the matrix G is no longer symmetric has no technical consequences (see, e.g., Lee, Liu, and Lin, 2010; Liu and Lee, 2010). However, our analysis remains roughly unchanged whether we use an undirected or directed network definition. We separately estimate the buy and sell networks to evaluate differences and find that the results are nearly identical. The results are available upon request. On further reflection, this similarity is reasonable. The driver of the network effects is the fact that market makers (who constitute most of the trading) tend to make similar returns as their neighbors by pursuing similar strategies. The agnosticism about buying and selling leads them to pursue a nearly identical amount of both. This agnosticism also means that there are no pressures on one side of the transaction from these central actors that would lead to differences in estimation results.

where $r_{i,\kappa}$ is the idiosyncratic return of trader i in the network k;

$$g_{i.,k} = \sum_{j=1}^{n_\kappa} g_{ij,\kappa} \tag{5}$$

is the number of direct links of i;

$$\frac{1}{g_{i.,k}} \sum_{j=1}^{n_\kappa} g_{ij,\kappa} r_{j,\kappa} \tag{6}$$

is the average returns of trading partners; $\upsilon_{i,k}$ is a random error term; and $x_{i,\kappa}^m$ is a set of M control variables at the individual or network level, or both. This model is the so-called spatial lag model or spatial autoregressive model in the spatial econometrics literature (see, e.g., Anselin, 1988) and can be estimated using standard software via maximum likelihood.

As shown in Lee, Liu, and Lin (2010), Eq. (4) captures recursively the network effects at any degree of separation and it is closely linked with a particular network centrality measure: Bonacich centrality (Bonacich, 1987, 2007).

3.2. Bonacich centrality

Bonacich centrality is a count of the number of all direct and indirect paths starting at node i and ending at node j, where paths of length p are weighted by θ^p. More paths from i to j imply a more central trader. A full description of Bonacich measure, including the connection with our Eq. (4), is contained in the Appendix.

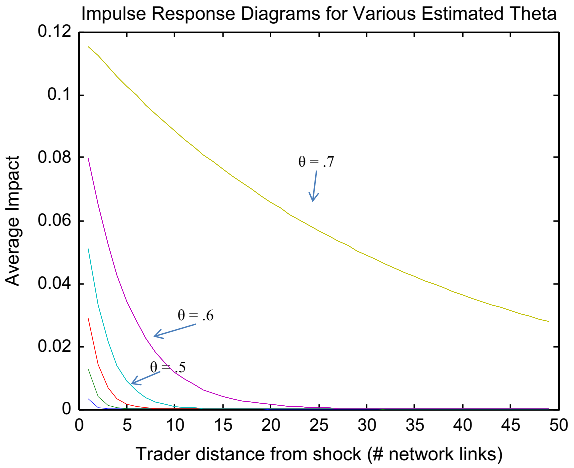

Fig. 4. Impulse response diagram for various estimated theta. This figure shows the impact of a one-unit shock to a network. Each line shows the impact of the shock for a different value of θ. The distance from shock shows the impact as one moves away from the origin of the shock to the remainder of the network, traversing along only trading relations. For example, two distance-degrees from shock would indicate the impact of i on k, with j in between.

To illustrate its relevance to a trading network, we proceed as follows. First, we explain the importance of understanding the role of indirect connections (Fig. 3). Next, we discuss why it can be helpful to describe how shocks are propagated through a financial network (Fig. 4). Finally, we highlight the difference between Bonacich centrality and the more traditional eigenvector centrality (Fig. 5).

Let us begin by considering the following simple network in Fig. 3. In this network, trader A has one direct transaction with another trader. That trader, marked B, has a record of transactions with A in the audit trail data that we use. In addition, B has transactions with traders C, D, and E. While C, D, and E did not trade with A, because of the linkage through B it is easy to see that A could be influenced in her decisions and, as a result, her profitability, by the actions of C, D, and E. As C, D and E trade with others, a complex and long chain of connections emerges. One particular challenge in understanding this type of system is capturing the importance of indirect links. It should be apparent that changes in A's action flow through to B as well as C, D, and E. Bonacich centrality measure allows one both to capture the salience of the direct and indirect links and to characterize the relative importance of counterparties versus the counterparties of counterparties.

To illustrate the relevance of this measure to the analysis of shock propagation, consider the original Bonacich (1987) example. Bonacich uses a network of individuals who communicate with each other. Using the notation of Section 3.2, the parameter θ measures the probability that a communication is transmitted by any individual to any of his contacts. θG is the expected number of these communications that are passed on to direct contacts, $\theta^2 G^2$ are the ones passed on to contacts two links away, and $\theta^p W^p$ is the expected number of messages that reach agents at path-length p. In the context of a trading network of mutual exposures, the magnitude of θ thus reflects the degree to which a shock is transmitted locally or to the structure as a whole. Small values of θ heavily weight the local structure, while large values take into account the position of agents in the structure as a whole.

Our final task is to highlight the differences of Bonacich centrality with respect to the more standard eigenvector centrality measure.[5] In our context, the eigenvector centrality would assume that $\theta = 1$ and, thus, would equally weight the entire network. All actions at distant points of the network impact a trader and with the same importance of actions close by. The bottom panel of Fig. 3 depicts these differences. By allowing θ to be different from one, the degree to which a shock is transmitted locally or to the structure as a whole can be understood. Using our data, we simulate the impact of a shock to a trader for different values of θ. Fig. 4 plots the results. One can see that as θ becomes larger, the shock transmits more widely across the network; i.e., it impacts traders much further away in the network.

Changes in the parameter θ can lead to changes in the agents' role in a network. A network can reveal distinct centrality scores depending on the centrality measure chosen. For illustration, we compare Bonacich and eigenvector centrality measures in the network (see Fig. 5). One can calculate the measures for each agent A, B, and C in the figure. If $\theta > 0.2$, A is more central than B, and eigenvector and Bonacich centrality return the same ranking. However, if $\theta < 0.2$, B is more central than A. Why does this occur? The intuition is that when θ is small, contacts further away are highly discounted and, as a result, the Bonacich measure counts only individuals who are close by. When θ is larger, agents that are far away begin to be counted as they would be in eigenvector centrality. In the eigenvector measure, there is no discounting for distance, so agent A here is as close to agents C as to agents B. Eigenvector centrality treats every connection as having the same weight; that is, two traders that are 20 links apart are similarly important in the measure as ones that are directly connected. The Bonacich measure weighs these links in an exponential fashion instead. Directly connected traders receive weight θ, second degree connected traders, θ^2, etc. This has a number of implications. One, it produces Bonacich centrality scores that are relatively higher than the eigenvector equivalent for traders that have many direct connections. Two, it permits analysis of distinct networks through an additional degree of freedom, the θ parameter, which can be derived from the estimation of Eq. (4). Broadly speaking, Bonacich centrality can be considered as a generalization of eigenvector centrality that permits additional understanding of the role of indirect connections.

3.3. *Weighted networks*

Eq. (4) is based on an unweighted network definition. This implies that the size of the single transaction when establishing a network link is not taken into account. However, trading with large counterparties would be different than trading with smaller ones. One can thus

[5] In recent years, social network studies have proposed different centrality measures to account for the variability in network location across agents. There is no criterion to pick up the right centrality measure. It depends on each particular situation (Borgatti, 2003; Wasserman and Faust, 1994).

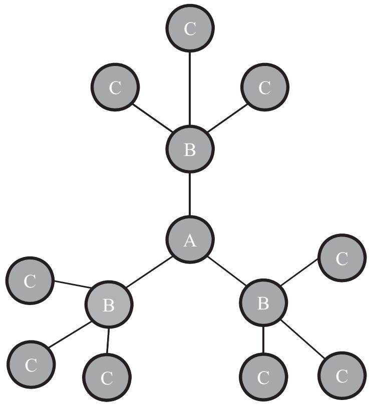

Bonacich and eigenvector centrality scores for various θ

Θ	A		B		C	
	Bonacich	Eigenvector	Bonacich	Eigenvector	Bonacich	Eigenvector
0	1.33	0.50	1.78	0.41	.44	0.17
.1	1.52	0.50	1.67	0.41	.51	0.17
.2	1.65	0.50	1.59	0.41	.55	0.17
.3	1.74	0.50	1.53	0.41	.58	0.17
.4	1.8	0.50	1.48	0.41	.60	0.17

Fig. 5. We compute the Bonacich and eigenvector centralities for the three types of agents in the network above as a function of θ. When using Bonacich centrality, the most central agent changes from B to A as one weighs the distant connections more heavily. Eigenvector centrality is not able to appreciate such a difference as it weighs all connections, near and far, equally. This example follows Bonacich (2007, p. 12) closely.

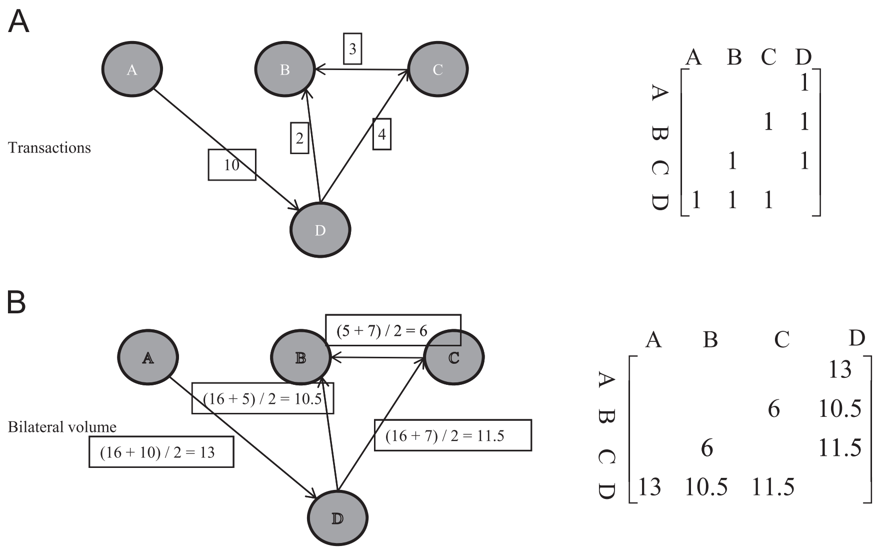

Fig. 6. Panel A shows a set of transaction between four traders. Each arrow is a single transaction, with the arrow pointing toward the buyer of a contract. Each value along an arrow shows the number of contracts traded. The accompanying matrix is an unweighted network representation of the transactions. Each cell contains a “1” where two brokers have transacted. Panel B shows the same set of transactions. Along each arrow is a calculation equal to (total trades of buyer + total trades of seller)/2. These values are then used as weights in the accompanying matrix.

extend the simple network model to use the network equivalent of importance weights in an ordinary least squares regression. We measure the importance of traders by total trading value and replace the binary matrix G with a new matrix capturing both the number of links and the importance of each link. Let the matrix $W = GD$, where G is as defined above and $D = \{d_{ij}\}$ is a matrix that weighs the links within the network. The scalar d_{ij} is a scaling factor, calculated as the total trading volume in the same trading period (the network) of each i and j. Total trading volume

Table 3
Estimation with network effects.
This table estimates the network equation (3). Shown are results from the Standard and Poor's (S&P) 500 futures market and the Dow futures market. The columns distinguish between different levels of network structure complexity. For each type network density and each market, we report the range of maximum likelihood estimation results across 21 trading days. The first row shows the estimates of the parameter θ, the network effect coefficient. t-Statistics are reported below coefficient estimates. We include the adjusted R-squared value from each specification. We denote significance of coefficients at the 10%, 5%, and 1% levels with ***, **, and *, respectively.

	Sparse networks 250 trades per time period		Moderately dense networks Five hundred trades per time period		Dense networks One thousand trades per time period	
	Low	High	Low	High	Low	High
S&P 500 e-mini futures						
Network effect coefficient (θ)	0.02***	004***	0.05***	0.08***	0.10***	017***
t-Statistic	38.01	38.02	37.13	42.66	38.18	44.59
Constant	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.05	0.09	0.09	0.17	0.19	0.37
DOW futures						
Network effect coefficient (θ)	0.02***	0.04***	0.04***	0.05***	0.08***	0.10***
t-Statistic	7.77	49.47	8.19	44.81	9.07	50.68
Constant	Yes	Yes	Yes	Yes	Yes	Yes
R- squared	0.04	0.08	0.07	0.09	0.14	0.18

is defined as the sum of all trades, both buys and sells, made by trader i with all other traders. As a result, $W=\{w_{ij}\}$ is now a weighted network w. Fig. 6 provides an illustration of the calculation of these weights. It shows a set of four transactions among four traders, A, B, C, and D. Each arrow is a single transaction, with the arrow pointing toward the buyer of a contract. In panel A, each value along an arrow shows the number of contracts traded. The accompanying matrix is an unweighted network representation of the transactions; i.e. the G from above. Each cell contains a "1" where two traders have transacted and "0" otherwise. Panel B shows the same set of transactions, but having along each arrow a calculation equal to our measure of importance: (total number of contracts bought or sold by buyer + total number of contracts bought or sold by seller)/2. These values are the ones used for the weights $D=\{d_{ij}\}$ to get a weighted accompanying matrix W.

In the original Bonacich (1987) paper, the centrality measure is presented for unweighted networks. However, the discussions throughout the paper (and the techniques in the Appendix) apply to the weighted network case; i.e., $G=W$ (Newman, 2004).[6]

4. Estimation results

The estimation results of Eq. (4) are collected in Table 3 for different levels of network structure complexity (see Section 2). Shown in the table are the results from the S&P 500 futures market and from the Dow futures market. For each type of network in each market, we separately estimate our model for each trading day and report the range of estimation results and t-statistics across the observed 21 days. The estimated θ coefficients for the S&P are between 0.02 and 0.17 depending on the day and the network type. Similarly, they are between 0.02 and 0.1 for the Dow. Most of these are estimated with a very high degree of precision. At the highest level, this suggests that correlations exist in the network between agents' returns. The magnitudes are economically important, suggesting that increases in trading partner returns could be an important determinant of one's own outcomes. The R-squared coefficients range from 0.05 to 0.37 for the S&P and 0.04 to 0.18 for the Dow. In one case, these regressions explain more than one-third of the variation in trader returns.

4.1. Results for weighted networks

The estimation results for the model with weighted networks are contained in Table 4. We replace G with W in Eq. (4) and run the same regression again. We again follow the format of displaying results by the density of the network.

The qualitative evidence remains unchanged, but the results are stronger. First, the estimated correlation between trader returns is now greater than 0.9 in the S&P and greater than 0.8 in the Dow. That is, the returns a trader earns are very similar to those of her trading partners. Second, across densities of network structure, we find estimates of θ that are large and always statistically significant. Across each specification, the observed t-statistics increase. The estimation is now much more precise than without the weights. Third, these new specifications are able to explain a much larger fraction of the variation in the trader-level returns. The adjusted R-squared values are now uniformly above 70% in both markets. Both the structure of the connections and their importance are important in understanding returns.

The last row of Table 4 reports values for the average multiplier, ϕ. The estimate θ is the average correlation between traders' profits and those of their counterparties. The value θ^2 describes the correlation between traders and their counterparties' counterparties, etc. As a result, we

[6] We are grateful to Jose Scheinkman for calling our attention to this.

Table 4
Estimation with weighted network effects.

This table extends the network model (3) to include weighted networks to reflect the relative importance of traders in the system. We use a weighted matrix W defined as the element-by-element product of the adjacency matrix of realized trades and the sum of trading volume. Reported are results from the Standard and Poor's (S&P) 500 futures market and the Dow futures market. The columns distinguish between different levels of network structure complexity. For each type network density and each market, we report the range of maximum likelihood estimation results across 21 trading days. The first row shows the estimates of the parameter θ, the network effect coefficient. t-Statistics are reported below coefficient estimates. We include the adjusted R-squared value from each specification and the average multiplier. This multiplier is total network impact of a one-unit shock to a trader. Averaging across the impact for all traders in the network produces this estimated effect, which is equal to $\phi = 1/(1-\theta)$. We denote significance of coefficients at the 10%, 5%, and 1% levels with ***, **, and *, respectively.

	Sparse networks 250 trades per time period		Moderately dense networks Five hundred trades per time period		Dense networks One thousand trades per time period	
	Low	High	Low	High	Low	High
S&P 500 e-mini futures						
Network effect coefficient (θ)	0 94***	0.96***	0.96***	0.98***	0.97***	0.98***
t-Statistic	1488.42	2594.92	619.46	619.46	516.89	669.17
Constant	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.74	0.77	0.73	0.77	0.73	0.77
Average multiplier (ϕ)	16.12	26.99	25.61	45.45	37.01	66.46
DOW futures						
Network effect coefficient (θ)	0.82***	0.89***	0.85***	0.92***	0.90***	0.95***
t-Statistic	355.26	475.98	316.51	415.36	233.36	309.37
Constant	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.71	0.80	0.71	0.79	0.71	0.78
Average multiplier (ϕ)	5.52	9.43	6.71	13.15	9.80	21.26

can look at a given trader and evaluate based on the network structure how a counterfactual change in her profits would be reflected in the profits of those near to her. Thus, if a trader gained \$1.00, the multiplier measures how much traders in the network win or lose. Because the coefficient θ measures the average correlation in returns across traders linked by a single node in the network, θ^2 measures the average correlation across two links; θ^3, the average across three, etc. Thus, a simple calculation allows us to measure the impact of a shock to any given trader. Consider a shock of \$1.00. On average, this leads to a change in earnings of directly connected agents of $\theta * 1$, agents two links away of $\theta^2 * 1$, etc. One can see, then, that for each dollar won or lost by a trader, $\phi = 1/(1-\theta)$ is the aggregate amount won or lost by agents connected at any level to the trader. For example, an estimate of θ equal to 0.5 produces a multiplier of two, suggesting that for each dollar lost by a trader hit by an exogenous shock, the individuals connected to the trader lose an aggregate of \$2.00. Because the market is zero-sum, if all agents are marked-to-market at the time of the idiosyncratic loss, the trading partners' \$2.00 of losses is offset by \$2.00 of gains elsewhere in the network. Our measure is thus a calculation of the degree of reallocation of profits.

Notice that ϕ is an average. If one evaluates the counterfactual for a highly central trader, the propagation is much different (and greater) than if the shock hits a relatively isolated trader.

We report the value of ϕ below each specification. The multiplier is between 16 and 66 for the S&P and five and 21 for the Dow. These large numbers imply that these trading networks have very high sensitivity to shocks. Small changes to individuals rapidly spread and magnify. These effects depend on both the structure of the connections and on the strength of the interaction, as captured by θ. As a result, the average multiplier can be helpful to characterizing the transmission of shocks in a trading system. The calculation of an average spillover following a shock defines the degree to which idiosyncratic losses become widespread ones.

4.2. Interpretation of results

Tables 3 and 4 show that traders that transact with each other in this market have highly correlated returns. The correlation emerges in the absence of specific information being shared between agents and of the agents having specific knowledge of the identity of their counterparties. So, the questions arise: Why would individual returns be correlated? How does this relation emerge?

To explain, we return to our evidence in Section 3 (Table 2 and Fig. 2). A consistent explanation is that the (observed) network of realized trades is a tool to describe the (unobserved) strategic interactions at work in the market. Because the matching algorithm used by the CME is blind to identities of the traders, traders with correlated strategies trade amongst themselves as well as with others. As they do so, and form links with one another, correlation in trading strategies leads to a connection between strategies and network position. Traders confirm that sitting between two traders with fundamental liquidity needs can be profitable. The very active, high-profit traders (type B) enter the market with the express purpose of exploiting profitable opportunities. They thus behave similarly at each point. These correlated bidding patterns lead to similarity in

Table 5
Network topology and profitability.
The table shows results from the Standard & Poor's (S&P) 500 futures market and the Dow futures market. The columns distinguish between different levels of network structure complexity. The exercise in this table is to report individual-level variation in centrality and evaluate the difference in returns for traders with different centrality. For each type of network density and each market, we report the range of results across 21 trading days. Individual-level Bonacich centralities are calculated using the formula: $b(w,\theta)=[I-\theta W]^{-1}T$, where the "1" signifies a vector of ones. We report the standard deviation of centrality as well as the change in returns for a trader that changes his centrality by one unit.

	Sparse networks		Moderately dense networks		Dense networks	
	Low	High	Low	High	Low	High
S&P 500 e-mini futures						
Impact of one-unit change in Bonacich centrality	0.06	0.70	0.07	0.78	0.10	0.76
Standard deviation-weighted Bonacich centrality	3.41	4.30	3.42	4.32	3.40	4.33
DOW futures						
Impact of one-unit change in Bonacich centrality	0.39	0.60	0.37	0.57	0.38	0.57
Standard deviation-weighted Bonacich centrality	3.68	4.11	3.68	4.11	3.68	4.11

returns. As noted in Section 3, intermediation transactions are a substantial fraction of trades.

The explanation of returns correlation helps in understanding shock propagation. A change in fundamental demands, for example, by traders A leads to a change in the profit opportunity of traders B. The profitability of the group of traders B changes in a linked fashion as they transact both with the traders of type A and with each other. The raw correlation of 0.3 (last column in Table 2) reflects the degree to which losses would be propagated in the absence of any network structure, while our estimated figures of 0.8 and higher (Table 4) suggest that the linkages in the network cause propagation that is much larger than would occur in the absence of such connections.

The emergence of network patterns from correlated trading strategies provides additional rationale for using the Bonacich centrality measure. As agents' strategies become more correlated, traders become more likely to trade with one another directly or within a small diameter (small number of intermediating traders). The Bonacich centrality measure weighs these closer traders more heavily than those that are very far away in the network. Those that are far away are those that have strategies dissimilar from one another. Measures that weigh all traders similarly regardless of distance, such as eigenvector centrality, would not capture this difference in strategies unless the difference had no impact on the likelihood of trading with any given trader.

4.3. Network centrality and profitability

Taken as a whole, the evidence in Section 4.2 indicates that the number of direct and indirection connections in the network, as weighted by θ, is a relevant factor that plays a role in explaining the cross-sectional variation of returns. To better understand the link between network topology and returns, we can use our estimated θ to calculate Bonacich centrality for each trader in our networks [Eq. (10) in Appendix A}.[7] With this distribution of positions in the network, we can look at the outcome differences across traders of different centrality levels. We do so by looking at the impact of a one-unit change in centrality on returns. Table 5 reports both the impact and the standard deviation of the centrality measure for sparse, moderately dense, and dense networks.[8] We report absolute changes in returns. Because the benchmark return is one, the numbers can also be interpreted as percentage changes. They are changes in returns over one-day time periods. We do not normalize to an annual basis. It appears that high returns are associated with high degrees of centrality irrespective of network complexity.

We note two patterns. One, the standard deviation of the centrality measure is nearly identical across the three network types. It is relatively similar across markets. Two, the impact of a one-unit change (approximately one-third of a standard deviation) is also relatively constant across network densities.

We highlight this finding as it suggests, in part, that our network definition is effective. Even though we construct our networks based on an ad hoc choice of transactions, the impact of the networks that we define remains consistently important throughout the measured time period. While one could improve upon the definition, the strength and consistency over time of these findings suggest that we capture a large portion of the network effect.

4.4. Network structure and distributional effects

Our analysis so far shows to what extent network position (network centrality) of an individual trader is important in explaining the level of individual returns. The more central a trader emerges from the exogenous matching process, the higher his returns.

In the remainder of this subsection, we highlight the implication of differences in network structures in terms of the distribution of outcomes in financial networks. That

[7] This calculation generates a distribution of individual centralities depending on the strength of network interactions and on the heterogeneity of network links [as captured by the estimate of θ and the matrix G, in Eq. (10)].

[8] We do not put centrality on the right-hand side of our regressions. The impacts are derived from a simple transformation of the estimated θ from Eq. (1).

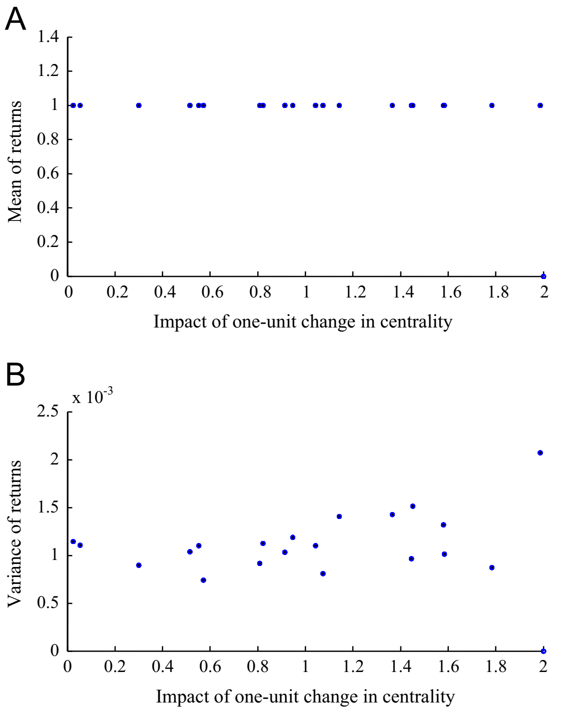

Fig. 7. Panel A shows the relation between a one-unit change in centrality and the mean of returns in the network. To calculate this, we take the average impact of a one-unit change across all traders in a given network and plot it against the mean of returns across traders in the same network. Panel B shows the relation between a one-unit change in centrality and the variance of returns in the network. To calculate this, we take the average impact of a one-unit change across all traders in a given network and plot it against the variance of returns across traders in the same network.

is, does a difference exist in the variance of returns for traders operating in different types of networks?

Recall first a few empirical patterns. One, we find that network structure explains individual-level returns well. Two, we find that the average multiplier, as measured by the ratio of an aggregate impact to the level of an individual shock, is very high in the networks that we analyze. Three, we find that an improvement in terms of centrality for an individual trader is associated with a positive change in returns.

Given these three findings and the fact that futures markets are zero-sum, we can make two claims. First, at the level of a network (250 to one thousand transactions), we should see that a change in the distribution of the centrality measure has no change on the mean return in a network. That is, an arbitrary reallocation of individuals around the network should change the distribution of outcomes, but not the mean.[9] Second, it thus follows that one should find differences in the variance of returns. We find evidence of these two phenomena in our data.

Fig. 7 displays the results. It relates the impact of network centrality to the variance of returns in the network and finds a positive relation. It also shows that the aggregate mean of returns remains roughly unchanged. As centrality becomes more important, the distribution of returns widens. This is a logical implication. If being central leads to greater returns, in a zero-sum market this necessarily means that someone at the periphery must lose out, and the variance of returns widens.

Technically, the relation shows that the distribution of returns of the network with greater sensitivity to centrality stochastically dominates (in a second-order sense only) the distribution of returns for a network with lower sensitivity to centrality.

5. Discussion and robustness checks

The validity of our analysis and its relevance for policy purposes hinges upon the correct identification of the network effect, θ.

The core problem that emerges in estimating linear-in-means models of interactions is the Manski (1993) reflection problem. This arises from the fact that if agents interact in groups, the expected mean outcome is perfectly collinear with the mean background of the group. How can we distinguish between trader i's impact on j and j's impact on i? Effectively, we need to find an instrument: a variable that is correlated with the behavior of i but not of j. Cohen-Cole (2006) notes that complex network structures can be exploited for identification. Bramoullé, Djebbari, and Fortin (2009) highlight the same phenomenon and show that, in network contexts, one observes intransitivities. These are connections that lead from i to j then to s, but not from s to j (see picture). Thus, we can use the partial correlation in behavior between i and j as an instrument for the influence of j on s.

That is, network effects are identified if we can find two agents in the economy that differ in the average connectivity of their direct contacts. A formal proof is in Bramoullé, Djebbari, and Fortin (2009). As a result, the architecture of networks allows us to get an estimate of θ, while eluding the reflection problem. A complex trading network such as the one we are concerned with has a very rich structure of connections, and identification essentially never fails.

Another traditional concern in the assessment of network effects in the social sciences is that network structure can be endogenous for both network self-selection and unobserved common (group) correlated effects. The first problem might originate from the possible sorting of agents. However, given our definition of networks based on high-frequency data and a random matching algorithm, we have no reason to believe that any selection effects exist in this context. Agents are assigned to trading partners as we described above, based on time and price priority alone. Even if two traders were to attempt to time a transaction as to ensure a match, the high volume of transactions on these markets makes this nearly impossible to complete. As such, we have a strong claim that individuals cannot choose their network partners and, thus, no selection effects should be present. In other

[9] We discuss above that the unweighted mean of returns at the network level might not always be one, given that some traders earn large profits. The reallocation can impact to a small degree this unweighted return but cannot impact the weighted network returns, which must always be equal to one.

Table 6
Estimation with weighted network effects and network fixed effects.

This table extends the weighted network model to include network fixed effects to control for group-specific unobserved factors. It shows results from the Standard and Poor's (S&P) 500 futures market and the Dow futures market. The difference between this table and Table 4 lies in the construction of returns. The results in Table 4 used the individual-level gross returns. Here we use the deviation in returns from the average return at the network level in each time period. For each type of network density and each market, we report the range of maximum likelihood estimation results across 21 trading days. The first row shows the estimates of the parameter θ, the network effect coefficient. t-Statistics are reported below coefficient estimates. We include the adjusted R-squared value from each specification and the average multiplier. We denote significance of coefficients at the 10%, 5%, and 1% levels with ***, **, and *, respectively.

	Sparse networks 250 trades per time period		Moderately dense networks Five hundred trades per time period		Dense networks One thousand trades per time period	
	Low	High	Low	High	Low	High
S&P 500 e-mini futures						
Network effect coefficient (θ)	0 94***	0.96***	0.95***	0.98***	0.97***	0.99***
t-Statistic	964.33	2917.94	544.24	718.19	438.62	666.50
Constant	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.74	0.77	0.73	0.77	0.73	0.77
Average multiplier (ϕ)	15.61	28.57	19.61	41.62	29.41	71.36
DOW futures						
Networkeffect coefficient (θ)	0.82***	0.88***	0.84***	0.92***	0.90***	0.95***
t-Statistic	385.74	448.79	307.21	411.65	237.46	310.21
Constant	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.71	0.78	0.70	0.79	0.71	0.78
Average multiplier (ϕ)	5.56	8.33	6.25	12.98	10.20	19.59

words, network topology is exogenous here. The possible presence of unobserved correlated effects instead arises from the fact that agents in the same group tend to behave similarly because they face a common environment or common shocks. These are typically unobserved factors. For example, traders with similar training, who sit in similar rooms or use trading screens that show similar types of data, could be influenced in their trading patterns in ways that generate correlations in returns. While we believe this to be very unlikely, we can control for these unobserved effects by reestimating our model after taking deviations in returns with respect to the group-specific means, i.e., from the average returns of (direct) trading partners. That is, if agents in a given empirically observed network have some similarity that leads them to earn higher returns as a group, we average out this group-level effect and look only for the presence of spillovers. Our primary specification already largely nets out market-level returns by virtue of the fact that aggregate market-level returns are one. In this case, we also control for group-level unobserved heterogeneity. In sum, there is little reason to believe that in an electronically matched market one would observe any effect of this sort.

Results are in Table 6 and illustrate very small differences from those in Table 4.

These results are useful, also, for another reason. The market that we are discussing is zero-sum. Benefits to a given individual are necessarily reflected in losses to another. As a result, complementarities in returns must necessarily be reflected in losses elsewhere in the network. We handle this issue by estimating our results in deviations from average-level returns for an individual's own network. In deviations, complementarities no longer are reflected elsewhere in the network structure and we can consequently use our results to evaluate the impact of a shock to the system. The particular context of analysis and our approach thus enable us to uncover a causal relation between network structure and profitability.

6. A policy experiment

One of the advantages of this approach is that it provides a mechanism via which policy makers and regulators can understand the impacts of their choices on the risk in the system. As a leading example, the August 2010 passage of the Dodd–Frank Wall Street Reform and Consumer Protection Act included a call for the evaluation of position limits in futures markets. The impact of such limits has been fiercely debated.[10]

In this section, we construct a counterfactual study that explores the consequences of this policy using our framework. Our exercise runs as follows. We set an arbitrary transaction limit for a given period of time. Given the restriction, we re-estimate our Eq. (4) assuming that any traders who, in the data, transact a greater number than this amount, transacted only the fixed maximum. Specifically, we restrict to C the number of contracts that can be

[10] An industry lawsuit in US District Court protesting the implementation of the Commodity Futures Trading Commission (CFTC) was decided in favor of industry in September 2012. The court sent the proposed regulations back to the CFTC for reworking.

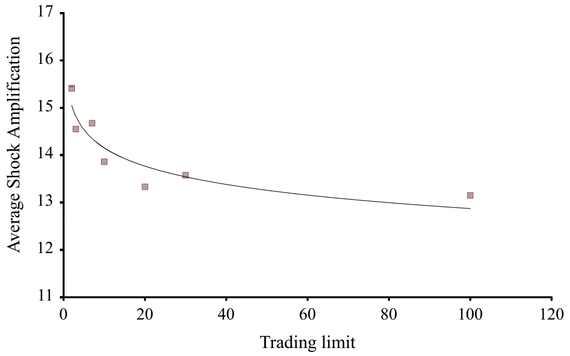

Fig. 8. This figure shows the results of a simulation in which traders face trading limits. Each simulation result is an estimate of the average multiplier. The vertical axis shows this average shock amplification estimate. The horizontal axis shows the maximum trading limit in the simulation. Limits on the horizontal axis indicate maximum trading volume during a pre-specified time period. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

purchased in one-tenth of a trading day, thus setting artificial bounds on the weights of our matrix W. This does not change the network structure other than the weights of the links. We consider $C=\{2,3,7,10,20,30,100\}$ for the Dow futures market and investigate the consequences of such limits given that high returns are associated with high degrees of centrality irrespective of network complexity. Fig. 8 shows the estimated shock amplification ($\phi=1/(1-\theta)$) for each of the values of C.

The simulation has two policy interpretations. First, the figure shows that tighter trading limits lead to higher values of ϕ (marked with red squares). These values can be interpreted as a measure of the transmission of shocks in a trading system. However, some additional detail is warranted. In our context, we measure the size of the pass-through to the system as a whole following an idiosyncratic shock. This is conceptually distinct from increases in the frequency of shocks (which we do not address). What we observe from this exercise is that the size of shock propagation increases as trading limits become tighter. In the case explored here, a move from no position limits to a strict one would increase the multiplier in the system, ϕ, from approximately 13 to 16.

Second, we can also infer from the exercise that tighter limits distribute the impact of the shock across a wider range of market participants. That is, while a shock in the constrained world could be widely distributed, an equivalent shock in the unconstrained world to a large trader could pass to only a small number of counterparties. This phenomenon arises because in our experiment we do not simulate new links between traders. The mechanism by which the transmission of shocks increases is to decrease the centrality of the network; that is, the limits downplay the importance of the traders who had previously exceeded the limit and been central.

Effectively, this highlights that the policy comes with a distinct trade-off. On the one hand, in our simulation, it has the potential benefit of dispersing adverse shocks to a wider range of market participants. On the other hand, the limits also appear to generate larger aggregate consequences from each shock. The \$5.00 loss could now be magnified to \$6.00 or \$7.00. The trade-off between the two determines the aggregate impact of the policy, and its final impact undoubtedly is market-specific.

7. Literature review

With the financial crisis, and increasing concerns about financial integration and stability as a leading example, a large number of theoretical papers have begun to exploit the network of mutual exposures among institutions to explain financial contagion and spillovers. Allen and Babus (2009) survey the growing literature and Allen, Babus, and Carletti (2012) provide an example of how interconnections lead to shock propagation. From an empirical point of view, however, little agreement exists in the literature on how to estimate the propagation of financial distress.

We contribute to this strand of the financial connections literature by providing an empirical approach able to capture the pathways of spillovers in a market with a single asset. By providing details on the spread of risk and the sources of profitability at this level of disaggregation, this paper increases understanding of the transmission of shocks in a trading system and in the development of policy.

As well, because we measure trader-level returns, we point to the literature that discusses the investment performance of individuals across portfolios, the price of individual or groups of assets, etc.[11] Another literature exists on the profitability of financial intermediaries, including specialists and trading desks.[12] By studying an individual asset across all traders, we can isolate the importance of financial interconnections. We contribute by suggesting that the profitability of trading is influenced by the particular market role, as described by the position in the network.

The closest paper to this one, to our knowledge, is Billio, Getmansky, Lo, and Pellizon (2012). It is one of the first using network measures to discuss the importance of financial interconnections. The authors creatively exploit aggregate data to infer links, but they do not have information on direct links between financial agents. They derive the connections from the Granger-caused correlations between returns in different financial sectors. Using this information, they measure network centrality using eigenvector centrality. Eigenvector centrality is a measure

[11] In studies that include portfolio management concerns, the heterogeneity in returns has been attributed to costs differences (Anand, Irvine, Puckett, and Venkataraman, 2010; Perold, 1988). Often the differences are explained by managerial ability to maintain the persistence in returns over time. For mutual funds, Kacperczyk and Seru (2007), Bollen and Busse (2005), and Busse and Irvine (2006) show that mutual funds maintain relative performance beyond expenses or momentum over multiple time periods.

[12] Reiss and Werner (1998) suggest that interdealer trade occurs between the dealers with the most extreme inventory imbalances. Sofianos (1995) disaggregates gross trading revenues into spread and positioning revenues and argues that, on average, about one-third of spread revenues go to offset positioning losses. Hasbrouck and Sofianos (1993) find that specialists are capable of rapidly adjusting their positions toward time-varying targets, and the decomposition of specialist trading profits by trading horizon shows that the principal source of these profits is over the short term.

that has many of the benefits of Bonacich centrality in that it calculates the connections to each node and weighs nodes with more connections more heavily in the calculation. The key distinction between this measure and the Bonacich one is that eigenvector centrality considers all connections in the network equally. That is, a counterparty's connection five, ten, or 15 links distant is as important under this measure as the connection to the counterparty itself. Bonacich centrality allows for different weights, which we estimate. In the application presented by Billio, Getmansky, Lo, and Pellizon (2012), the difference is immaterial as they do not study second- or higher-degree connections. The context we discuss here requires a varying treatment of connections based on their distance and, as a result, suggests an alternative approach.

8. Concluding remarks

Our analysis explains a conjectured, but to date unproven, feature of financial markets: returns from trading are correlated with the position agents occupy in a trading network. Using our network-based empirical strategy on two highly liquid financial markets, we are able to explain a large portion of the individual-level variation in returns. This finding has potentially large salience.

Most important, one of our results is that individual-level shocks are greatly amplified and spread in these markets. A one-unit change in individual-level returns can be amplified even 50 times. This implies very rapid propagation of shocks and little ability to avoid contagion. The estimate of network effects with financial data has a nice interpretation as a measure of risk magnification and spread. In fact, network effects can capture the propagation and amplification of financial shocks.

Because these results are a function of the network structure, they point policy makers in the direction of potential interventions. The rapid spread and amplification derive from the network structure. Adjusting the structure can impact the speed of spillovers. This points toward interventions in the matching algorithm, potentially during times of anticipated crisis. The most direct antecedent of high spillovers is the presence of market makers, both HFT and broker-dealers. The trading strategies of agents whose principal function is market intermediation have the impact of creating a market structure that is highly sensitive to shocks. Altering the matching algorithm to reduce the incentive of market makers to race to the center of the network would minimize the impact of shocks.

At one extreme, one could eliminate the impact of the race to the center by concentrating trading into hourly or twice-hourly auctions instead of continuous trading. By clustering trading into periodic auctions, the market itself would take on the matching and liquidity functions of market makers. Periodic auctions are already used in some markets as a way to ensure efficient price discovery in the absence of market makers. Whether the loss of liquidity in the 30–60 min between auctions has tangible costs to market participants is a topic for further research.

An alternate policy intervention that could minimize the scale of shock amplification would be to maintain continuous auctions but limit transaction speed of market participants. Speed limitations would alter the network structure by preventing the aggregation of high-speed traders at the center of the network. This aggregation occurs because many of these traders have similar strategies, and the speed of their transactions results in many of their trades being with each other. Instead, traders with fundamental liquidity needs would end up trading with each other more frequently, reducing the centralization of the network and with it the size of shock amplification.

Our policy simulation experiment also discusses the potential impact of trading limits on shock amplification and the trade-offs that emerge as a result.

While each of these policy interventions could have trade-offs, this paper emphasizes that the network structure itself can have a tangible impact on profits of participants and the amplification of shocks. Each of these can be influenced by the regulatory and operational structure of the market and, as a result, is an area for policy makers to consider going forward.

A long literature in sociology and economics would suggest that network patterns are important in non-market interactions, based on a variety of plausible mechanisms. These include social stigma, information sharing, peer pressure, and more. The difficulty in translating the methodologies developed in the social science to financial markets, particularly electronic ones, is that there is little basis to believe that any of the mechanisms are at work. Orders are matched at random by a computer based on time and price priority, leaving little room for considering social impact even if traders had a motivation to do so. Thus, our conclusions are statements about the empirical importance of the networks that emerge as a result of equilibrium order strategies.

Appendix A. The spatial autoregressive model and network centrality

For ease of interpretation, let us write Eq. (4) in matrix notation and derive the reduced form. The following derivations are helpful in understanding why Eq. (4) captures recursively the network effects at any degree of separation and the link with a particular network centrality measure, Bonacich centrality (Bonacich, 1987).

Eq. (4) can be written as

$$\mathbf{r} = \theta G \mathbf{r} + \beta \mathbf{x} + \epsilon, \tag{7}$$

where $\mathbf{r}$ is an $N \times 1$ vector of outcomes of N agents, $\mathbf{x}$ is an $N \times M$ matrix of M variables that could influence agent behavior but are not related to networks, G is the N-square matrix that keeps track of the direct links between agents, and ϵ is an $N \times 1$ vector of error terms, which are uncorrelated with the regressors.

Given a small-enough value of $\theta \geq 0$, one can define the matrix

$$[I - \theta G]^{-1} = \sum_{p=0}^{+\infty} \theta^p G^p \tag{8}$$

The p-th power of the matrix G collects the total number of paths, both direct and indirect, in the network starting at node i and ending at node j. The parameter θ is a decay factor that scales down the relative weight of longer paths; i.e., paths of length p are weighted by θ^p. It turns out that an exact strict upper bound for the scalar θ is given by the inverse of the largest eigenvalue of G (Debreu and Herstein, 1953).

In a row-normalized matrix, such as the one used in Eq. (4) to represent average returns, the largest eigenvalue is one. If $|\theta|<1$, Eq. (8) is well defined, that is, the infinite sum converges. The condition $|\theta|<1$ captures the idea that connections further away are less influential than direct contacts and guarantees that the matrix $[I-\theta G]^{-1}$ is able to capture all the effects that stem from a given network topology; that is, the cascades of effects stemming from direct and indirect connections.

If $|\theta|>1$, the process is explosive. In a financial network context, it is equivalent to a complete financial collapse. While interesting in its own right, we do not analyze this case here. We focus on how, even in the absence of a complete financial collapse, a small shock can cascade causing large, measurable and quantifiable damage. Therefore we consider $|\theta|<1$.

If one solves for $\mathbf{r}$ in Eq. (7), the result is a reduced form equation:

$$\mathbf{r}=[I-\theta G]^{-1}\beta\mathbf{x}+[I-\theta G]^{-1}\epsilon \tag{9}$$

Definition 1 (*Bonacich, 1987*). Consider a network g with adjacency N-square matrix G and a scalar θ such that $\mathbf{M}(g,\theta)=[I-\theta G]^{-1}$ is well defined and non-negative. Let $\mathbf{1}$ denotes the N-dimensional vector of ones. The vector of centralities of parameter θ in g is

$$\mathbf{b}(g,\theta)=[I-\theta G]^{-1}\cdot\mathbf{1}. \tag{10}$$

The centrality of node i is thus $b_i(g,\theta)=\sum_{j=1}^{n}m_{ij}(g,\theta)$ and counts the total number of paths in g starting from i. It is the sum of all loops $m_{ii}(g,\theta)$ starting from i and ending at i and all outer paths $\sum_{j\neq i}m_{ij}(g,\theta)$ that connect i to every other player $j\neq i$; that is

$$b_i(g,\theta)=m_{ii}(g,\theta)+\sum_{j\neq i}m_{ij}(g,\theta). \tag{11}$$

By definition, $m_{ii}(g,\theta)\geq 1$ and, thus, $b_i(g,\theta)\geq 1$, with equality when $\theta=0$.

Therefore, once one has on hand an estimate of θ, the distribution of Bonacich centralities can be derived for all the agents in the network.

References

Allen, F., Babus, A., 2009. Networks in finance. In: Kleindorfer, P., Wind, J. (Eds.), The Network Challenge, Wharton School Publishing, Philadelphia, PA, pp. 367–382.

Allen, F., Babus, A., Carletti, E., 2012. Financial connections and systemic risk. Journal of Financial Economics 104 (3), 519–534.

Allen, F., Gale, D., 2000. Financial contagion. Journal of Political Economy 108, 1–33.

Anand, A., Irvine, P., Puckett, A., Venkataraman, K., 2010. Performance of institutional trading desks: an analysis of persistence in trading costs. Review of Economic Studies 25 (2), 557–598.

Anselin, L., 1988. Spatial Econometrics: Methods and Models. Kluwer Academic Publishers, Boston, MA.

Borgatti, S.P., 2003. The key player problem. In: Breiger, R., Carley, K., Pattison, P. (Eds.), Dynamic Social Network Modeling and Analysis, National Academy of Sciences Press, New York, pp. 241–252.

Billio, M., Getmansky, M., Lo, A.W., Pellizon, L., 2012. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. Journal of Financial Economics 104, 535–559.

Bollen, N., Busse, J., 2005. Short-term persistence in mutual fund performance. Review of Financial Studies 18, 569–597.

Bonacich, P., 1987. Power and centrality: a family of measures. American Journal of Sociology 92, 1170–1182.

Bonacich, P., 2007. Some unique properties of eigenvector centrality. Social Networks 29, 555–564.

Bramoullé, Y., Djebbari, H., Fortin, B., 2009. Identification of peer effects through social networks. Journal of Econometrics 150, 41–55.

Brunnermeier, M., 2009. Deciphering the liquidity and credit crunch 2007–2008. Journal of Economic Perspectives 23, 77–100.

Brunnermeier, M., Pederson, L., 2009. Market liquidity and funding liquidity. Review of Financial Studies 22, 2201–2238.

Busse, J., Irvine, P., 2006. Bayesian alphas and mutual fund performance. Journal of Finance 61, 2251–2288.

Cohen-Cole, E., 2006. Multiple groups identification in the linear-in-means model. Economics Letters 92, 753–758.

Debreu, G., Herstein, I.N., 1953. Non-negative square matrices. Econometrica 21, 597–607.

Hasbrouck, J., Sofianos, G., 1993. The trades of market makers: an analysis of NYSE specialists. Journal of Finance 48, 1565–1594.

Kacperczyk, M., Seru, A., 2007. Fund manager use of public information: new evidence of managerial skill. Journal of Finance 62, 485–528.

Kirilenko, A., Kyle, A., Samadi, M., Tuzun, T., 2011. The flash crash: the impact of high frequency trading on an electronic market. Unpublished working paper.

Lee, L.F., Liu, X., Lin, X., 2010. Specification and estimation of social interaction models with network structures. Econometrics Journal 13, 145–176.

Liu, X., Lee, L.F., 2010. GMM estimation of social interaction models with centrality. Journal of Econometrics 159, 99–115.

Manski, C., 1993. Identifications of endogenous social effects: the reflection problem. Review of Economic Studies 60, 531–542.

Newman, M., 2004. Analysis of weighted networks. Physical Review E 70, 056131.

Perold, A., 1988. The implementation shortfall: paper versus reality. Journal of Portfolio Management 14, 4–9.

Reiss, P., Werner, I., 1998. Does risk sharing motivate inter-dealer trading? Journal of Finance 53, 1657–1704.

Sofianos, G., 1995. Specialist gross trading revenues at the New York stock exchange, Unpublished working paper 95-01. New York Stock Exchange, New York, NY.

Wasserman, S., Faust, K., 1994. Social Network Analysis: Methods and Applications. Cambridge University Press, Cambridge UK.