

Networks of Common Asset Holdings: Aggregation and Measures of Vulnerability*

Anton Braverman[†] and Andreea Minca[‡]

August 6, 2018

Abstract

This paper quantifies the interrelations induced by common asset holdings among financial institutions. A network representation emerges, where nodes represent portfolios and edge weights aggregate the common asset holdings and the liquidity of these holdings. As a building block, we introduce a simple model of order imbalance that estimates price impacts due to liquidity shocks. In our model, asset prices are set by a competitive risk-neutral market maker and the arrival rates for the buyers and sellers depend on the common asset holdings. We illustrate the relevance of our aggregation method and the resulting network representation using data on mutual fund asset holdings. We compare three related measures of vulnerability in the network and demonstrate a strong dependence between mutual fund returns and these measures.

Keywords: Financial networks, Mutual Funds, Contagion, Liquidity, Flow Imbalance, Network Vulnerability.

JEL: C81, C31, C44

*For helpful comments we thank Pam Moulton, Gala Nuno (discussant) and Rama Cont. This work was presented at the INFORMS Annual Meeting 2013, the European Central Bank Workshop on Using Big Data for Forecasting and Statistics, 2014, the Newton Institute Workshop "Monitoring Systemic Risk: Data, Models and Metrics". the SIAM Meeting in Financial Mathematics, 2014, the Bank of England Research Seminar 2017. Andreea Minca's research is partly funded by NSF CAREER Award #1653354.

[†]Kellogg School of Management, Northwestern University, anton.braverman@kellogg.northwestern.edu, +1 (847)-491-5508.

[‡]School of Operations Research and Information Engineering, Cornell University, acm299@cornell.edu, +1(607)2559133.

Introduction

It is by now generally accepted that common asset holdings mediate contagion among financial institutions. The contagion mechanism works as follows: two institutions, A and B, have the same asset as a part of their portfolio. Suddenly, A is forced to liquidate the asset due to some exogenous shock. Liquidation has an impact on the common asset price, and therefore on the value of institution B's portfolio. The initial shock may be due to leverage targeting, see e.g. [Adrian and Shin, 2010, Greenwood et al., 2015]; to a bank run, [Gorton and Metrick, 2012]; to large payables related to derivatives [Zawadowski, 2013, Amini et al., 2015]; to investor flows [Coval and Stafford, 2007]. For all these reasons, common asset holdings create a *de facto* network that may transmit financial distress.

The main contribution of this paper is to provide a methodology to capture the linkages due to common asset holdings among financial institutions. We propose a *weighted* network representation. The nodes represent the portfolios of these institutions. The edge weights capture the strength of interrelations due to common asset holdings. On the theoretic side, the question we ask is: *How can we quantify these links?*

The question of weight attribution is essentially a question of aggregation of the actual portfolio holdings and the liquidity characteristics of the common assets. To see why liquidity characteristics are crucial to the model, consider the following example of two institutions whose portfolios each consist of 1000 units of a perfectly liquid stock, i.e, a stock whose price does not change no matter how much of it is traded. Despite the fact that these institutions have common assets, no institution can affect the other by trading the perfectly liquid asset. In this case, the strength of the interrelation is zero. If, on the other hand, the asset were illiquid, then the two funds would be very strongly related.

Portfolio holdings are directly measurable from the data. Estimating asset liquidity, on the other hand, is in itself a challenging problem. The simplest way to incorporate asset liquidity in networks of common asset holdings is using a linear exogenous price impact function. In this case, Kyle's lambda [Kyle, 1985] captures the liquidity characteristics of the stock.

A building block of our network model is a model for non-linear price impact that accounts for the common asset holdings. Our approach is influenced by the market microstructure literature that models temporary liquidity price impacts in relation to imbalances in the order flow. In our model, we consider that the asset has a fundamental value, but, depending on the prevailing supply and demand, it will trade at a discounted value due to a market clearing condition imposed by a specialist. The supply (demand) is due to forced liquidations in the network (or asset purchases in the network, for those institutions with positive liquidity shocks).

The most important implication of our proposed network representation is that institutions' vulnerabilities to their neighbors shocks can be quantified. The vulnerability of a node is to be understood as a measure of how a node will respond to the shocks applied to its neighbors. In periods when neighbors have positive shocks, vulnerable nodes benefit from network effects, while in periods when neighbors have negative shocks, vulnerable nodes' returns are diminished due to

network effects.

In the second part of the paper, we test the relation between vulnerability and returns using data on mutual funds. The application of the network model to mutual funds requires us to specify the source of the initial liquidity shock. As documented in [Coval and Stafford, 2007], investor flows into (out of) mutual funds trigger expansion (reduction) of positions in assets; they refer to this as flow-induced trading. Flow-induced trading occurs because mutual funds do not hold large cash reserves, therefore they must liquidate assets to repay leaving investors or expand positions when they receive inflows. [Coval and Stafford, 2007] find that flow-induced trading creates price pressure, which is precisely how one fund affects others in the network of common asset holdings.

Contribution. We compare three measures of vulnerability in the network of mutual funds' portfolios and demonstrate that these measures are correlated with funds' returns. While the first measure is a very natural measure of vulnerability and resembles measures of a node's strength well known in social networks [Barrat et al., 2004], the second and third measure are flow adjusted refinements. Our main contribution in this paper is the third measure which removes the linear price impact assumption and is shown to outperform the second measure.

The three measures of vulnerability can be described as follows. The first measure, the vulnerability index VI , is a baseline measure. It is generated, for each mutual fund, by summing up all its exposures through common asset holdings to the other mutual funds, all renormalized by the size of the mutual fund. This can be interpreted as the aggregate effect of the network on the mutual fund, under the condition of *uniform liquidity shocks across all funds*. We find that this measure, while not making any specific assumptions on the extent of the initial liquidity shock, predicts mutual fund returns following market-wide events such as the equity market crashes in 2008 and 2011. Moreover, steady increases in the average vulnerability index across network are shown to precede significant drops in the total net assets of mutual funds. We also find that the average vulnerability is exacerbated during periods of crises, and that this is both a liquidity effect and due to increased portfolio similarity.

The second measure, the flow-adjusted vulnerability measure FAV , removes the assumption of uniform liquidity shocks across all funds, and defines a *mutual-fund specific liquidity shock induced by fund flows*. In both the first and second measures, the liquidity characteristics of each stock are captured using a stock specific linear price impact function.

The third measure, FAV^* is a refinement of the flow-adjusted vulnerability measure, in which we propose a non-linear price impact function. The order imbalance depends on the common asset holdings and the fund flows.

The two flow-adjusted vulnerability measures FAV and FAV^* complement the baseline measure. They do not predict returns, since they use concurrent fund flows as triggers of the initial liquidity shocks in the network of common asset holdings. However, contrary to the baseline measure, these flow-adjusted vulnerability measures are applicable at all times and not just during times of market-wide events.

These measures are positively correlated with fund returns throughout all our sample period.

Their explanatory power is maintained after controlling for the concurrent fund flows. When applied to the data, we find that the measure FAV^* , in which the price impact is non-linear outperforms the measure FAV based on a linear price impact: FAV is no longer significant in a horse race.

There are, of course, many other effects that drive asset prices, other than the mechanical reaction to liquidations of institutional investors. Despite its simplicity, our network based framework is shown to have explanatory power. The network-based framework could be extended in various directions, to incorporate a larger part of the financial system beyond mutual funds, for example banks, hedge funds, retail investors. It could also be extended to various triggers of the initial shocks beyond investor redemptions, such as leverage targeting, fire sales, or predatory trading. The model could go beyond first order network effects, integrate more sophisticated liquidation strategies beyond the scaling of existing asset positions, or allow for long-short asset positions. This level of complexity is however beyond the scope of the current paper. We expect that the explanatory power of the proposed measures would only increase.

Relation to the previous literature Capturing interrelations among financial institutions due to common asset holdings has been studied by [Blocher, 2016], who consider common asset holdings by mutual funds, [Greenwood et al., 2015, Caccioli et al., 2014, Duarte and Eisenbach, 2015, Capponi and Larsson, 2015], who consider common asset holdings by banks and [Caccioli et al., 2015], who consider the interplay between contagion in the network of common asset holdings and contagion in the network of interbank loans. Our focus on mutual funds is driven by public data availability. While mutual funds individually may not be a threat for financial stability, our paper is the first to point out the emergence of indirect networks of common asset holdings by mutual funds. Currently, contagion through asset holdings is part of banks' stress tests, but only contagion from other banks is taken into consideration. Contagion among funds/ asset managers and banks is not currently part of the stress tests. We make the point that if we have contagion among mutual funds, it is very likely that there is contagion among banks and funds as well.

There is a significant econometrics literature that aims at estimating network measures of dependence for a panel of time series, see e.g. [Billio et al., 2012, Diebold and Yilmaz, 2013, Barigozzi and Brownlees, 2013, Musmeci et al., 2015, Stavroglou et al., 2017]. Within this literature, in general, no assumptions are needed on the structural reasons that lead to dependence in the time series. Rather, their estimates lead to identifying channels that lead to this dependence.

Our approach is complementary. Once a channel of contagion is recognized, for example common asset holdings, our framework allows us to represent at each point in time, *using the structural details of the portfolio holdings*, the linkages among institutions due to this specific channel of contagion. Based on this representation, we can derive a time series for the vulnerability measures for the components in the system, that can be used for monitoring purposes by regulators, fund managers, etc.

The papers closest to ours are [Blocher, 2016] and [Lou, 2012]. In [Blocher, 2016], the author constructs a network of mutual funds and assigns edge weights between two funds based on the

similarity of the investing strategies of each fund. He then uses this network to partially explain future returns and fund flows. His focus is on demonstrating the impact of second order network neighbors on future returns. He considers the portfolio liquidity estimates based on [Amihud, 2002] as separate factors driving future returns. In particular, in his model, the network representation and the aggregate liquidity of the portfolios are entirely separated. The study in [Blocher, 2016] is based on [Cohen et al., 2005] and uses the same similarity measure to explain future returns.

Although he does not explicitly use a network representation, [Lou, 2012] considers how flow-induced trading affects stock return predictability. He then aggregates these affects across all stocks belonging to a portfolio to give a measure that predicts portfolio returns. He incorporates the individual stock liquidity by using the total number of shares held by mutual funds. This measure of stock liquidity is motivated by [Gompers and Metrick, 2001] who show that mutual funds' holdings are skewed toward liquid stocks. Unlike [Lou, 2012], our focus in the current study is not return predictability *per se*, but to demonstrate that vulnerability in networks of common asset holdings can be measured and that the resulting measures can explain returns.

In their study on bank holdings, [Greenwood et al., 2015] incorporate individual stock liquidity by including a market depth parameter for each stock, which measures the price impact of trading the stock. However, this parameter is later assumed to be identical for all stocks, making the individual stock characteristics irrelevant to the overall network effect. A similar strategy is used in [Duarte and Eisenbach, 2015]. Contrary to [Greenwood et al., 2015], [Duarte and Eisenbach, 2015] incorporate a sophisticated model for investor outflows. [Capponi and Larsson, 2015] analyze theoretically a model similar to [Greenwood et al., 2015].

On the theoretical side, [Caccioli et al., 2014] model default contagion in a network of banks that have constraints on leverage. Their underlying dependence structure is based on common asset holdings, with an exogenous price impact function. This work is in the same spirit as the theoretical literature on random graph applications to financial networks, see e.g. [Amini et al., 2016, Battiston et al., 2012]. These works rely on approximations, in large networks, on the whole scope of various channels of contagion. We are complementary to the theoretical model of [Caccioli et al., 2014] in that we use real world mutual fund data.

In general, our paper is part of the growing literature on financial networks, see e.g. some of the more recent works [Blume et al., 2011, Cabrales et al., 2017, Allen et al., 2012, Cohen-Cole et al., 2011, Elliott et al., 2014, Gai and Kapadia, 2010, Glasserman and Young, 2015, Gourieroux et al., 2012]. For reviews see, e.g. [Babus and Allen, 2009, Amini and Minca, 2013, Kara et al., 2015]. Importantly, most of the literature on financial networks interprets the interrelations as contractual liabilities of various maturities. Contrary to networks of liabilities, the links in networks of common asset holdings are not readily specified as quantities available on balance sheets. The ability to quantify these links is a keystone to understanding the systemic risk due to common asset holdings.

Our network construction has some important side implications for asset pricing. In our model, no exogenous parameters such as asset correlations are required and the network of common asset holdings can be thought of as a partial dependence structure among portfolio returns. Of course,

the network of common asset holdings can also be seen as a dependence structure among the stock returns themselves, and may explain endogenous correlation arising from institutional ownership, see the theoretical model in [Cont and Wagalath, 2013]. In this sense, the correlation of fund returns with the second measure provides empirical support for this theory. A more direct verification, using historically uncorrelated stocks and institutional ownership data, is provided in [Gao et al., 2017].

The fund vulnerability that our measures aim to capture can be considered as the *dual* of the stock frailty introduced in [Greenwood and Thesmar, 2011]: The vulnerability of a fund is obtained by aggregating the price impact of liquidations across all stocks owned by a fund; The frailty of a stock is obtained by aggregating the price impact due to liquidation across all funds owning the stock. As in the study on vulnerable banks [Greenwood et al., 2015], in [Greenwood and Thesmar, 2011] the price impact is linear and captured by a market depth parameter, assumed to be identical across stocks.

In our companion paper [Guo et al., 2016], we investigate the network topology of the common asset holdings networks. We find that this network has a scale-free structure and that funds tend to cluster in a small number of clusters. Moreover we add the network centrality as an additional factor to the Carhart four factors model, and show that it is a significant factor driving returns. The focus in [Guo et al., 2016] is on hedge fund networks, and network effects are demonstrated for that industry as well. In contrast, in the current paper we refine the way we can incorporate stock liquidity in measures of fund vulnerability by departing from the linear price impact which underlines the baseline vulnerability measure.

The rest of the paper is structured as follows: Section 1 introduces the network representation, Section 2 describes the data on mutual fund holdings, Section 3 introduces three measures of vulnerability and tests the dependence of these measures and the mutual fund returns and Section 4 concludes.

1 Model

In this section we discuss a network model for a cross-section of common asset holdings by financial institutions (banks, mutual funds, etc.). First we introduce the notation. Since each institution is allowed to have only one portfolio, we will use the term portfolio to refer to one institution in our network.

Consider the case where $\mathbf{N} = \{1, \dots, N\}$ is a set of portfolios and $\mathbf{K} = \{1, \dots, K\}$ is a set of stocks. Each portfolio owns a subset of \mathbf{K} , and two portfolios may have common holdings.

Let $S = (s_1, \dots, s_K)$ be the vector of stock prices.

We denote the holdings of each portfolio by the matrix

$$B = [\beta_{ki}] \quad i \in \mathbf{N}, k \in \mathbf{K},$$

where β_{ki} represents the number of shares of stock k owned by portfolio i . We do not allow short positions, therefore each β_{ki} must be non-negative.

This assumption may be relaxed when one wishes to consider other financial institutions (such as hedge funds) where short-selling is more prevalent. This could allow for network effects akin the ‘quant’ event of 2007, when the liquidation of a long-short neutral portfolio on the S&P500 did not affect the value of index funds [Khandani and Lo, 2011]. In our terms, these index funds would have zero weights on their linkages to the liquidated long-short portfolio. Since we are working with mutual fund data, in our setting, considering only long positions is not a major restriction.

The value of portfolio i can be written as

$$P_i = \sum_{k=1}^K \beta_{ki} s_k = \beta_i \cdot S, \quad \text{where } \beta_i = (\beta_{1i}, \dots, \beta_{Ki}).$$

Denote the vector of portfolio values by $P = (P_1, \dots, P_K)$.

We can represent the interrelations of financial institutions through common asset holdings as a network whose nodes correspond to the portfolios. There is an edge between two nodes if the corresponding portfolios hold common assets. In order to capture the extent of the relationship between two portfolios, we introduce edge weights. It is intuitive to want two portfolios with large asset commonality to have a strong relationship and portfolios with little asset commonality to have a weak relationship. We therefore define the weight of an edge between two portfolios by answering the following question: what effect will the liquidation of fund i have on fund j ?

1.1 Edge weights under linear price impact

In order to define the edge weights, we must incorporate the *price impact* of trading (In our model, we consider that trading is uninformed and the price impact is exclusively due to fund liquidations). To capture this effect, we define the function $PI_k(x)$ to be the (relative) change in the price of stock k due to x shares of this stock being traded.

In the baseline model, we make the simplifying assumption that the price impact function is linear and of the form:

$$PI_k(x) = \frac{x}{\lambda_k}, \quad (1)$$

where λ_k is such that buying/selling $\frac{\lambda_k}{100}$ stocks will move the price of the asset up/down by 1%. The parameter λ_k captures the *market depth* of stock k , see the seminal paper [Kyle, 1985]. In Section 3.3, we will develop a substitute for the linear price impact based on a model of a market in which arrival rates for the buyers and sellers depend on the common of asset holdings and prices are set by a competitive, risk-neutral market maker.

In order for the price impact to take into account the different characteristics of each stock,

we scale it by the average daily volume traded and multiply it by the stock's volatility as in [Almgren et al., 2005, Amihud, 2002]

$$\lambda_k = \frac{1}{\tilde{\lambda}} \frac{ADP_k}{\sigma_k},$$

where ADP_k is the average daily volume of trades, σ_k is the daily returns standard deviation of stock k and $\tilde{\lambda}$ is an invariant across stocks [Kyle and Obizhaeva, 2016]. Our results however will not depend on the proportionality constant $\tilde{\lambda}$.

We are now ready to define our edge weights. When portfolio i liquidates its shares of asset k , the price of the asset s_k drops by $\frac{\beta_{ki}}{\lambda_k} s_k$. This causes portfolio j 's value to decrease by $\beta_{kj} \frac{\beta_{ki}}{\lambda_k} s_k$. Hence, the total loss experienced by j if portfolio i liquidates can be calculated by summing this quantity across all assets

$$w_{ij} = \sum_{k=1}^K \frac{\beta_{ki}}{\lambda_k} \beta_{kj} s_k. \quad (2)$$

Observe that $w_{ij} = w_{ji}$ (symmetric) and if i and j have no assets in common, then $w_{ij} = 0$. Thus we set the weight of the edge connecting i and j to be w_{ij} .

The first two measures of vulnerability in the network of common asset holdings use directly the links specified above. Our third measure of vulnerability is based on the model introduced in the next section for the impact of order imbalance on a stock's price.

1.2 A non-linear price impact model

In the previous section, we assumed that the change in the price of an asset was proportional to the net supply/demand, with the market depth being the constant term in this linear relationship. In this section develop a non-linear alternative to the price impact, by proposing a simple mechanism to determine how imbalances in supply and demand of an asset affect its price.

Suppose we have a single asset with constant fundamental value p , which is traded continuously by a specialist at a single price \hat{p} . Trading occurs during times $t \in [0, 1]$ and the price \hat{p} is chosen by the specialist at time $t = 0$ and remains constant over time. We assume that buyers and sellers arrive to the market according to independent, time-homogeneous Poisson Processes with rates $r_B(\hat{p})$ and $r_S(\hat{p})$, both dependent on \hat{p} . We assume that as \hat{p} decreases, $r_B(\hat{p})$ increases and $r_S(\hat{p})$ decreases. Intuitively, this relationship means that as the price decreases, more people are willing to buy and less people are willing to sell.

We assume that the specialist knows these arrival rates. This is a reasonable assumption since the specialist sees all order flow and can therefore estimate these rates. We assume that the specialist has deep pockets i.e. that the specialist's risk of running out of inventory or capital during the trading period $[0, 1]$ is negligible. In reality, a specialist will set a price spread as compensation for the liquidity risk he bears; however, in this simplified model the specialist must trade at a single price. Since trading occurs at a single price, the specialist does not earn any trading profits. We therefore assume that the specialist will choose \hat{p} to be a market clearing price. By this we mean that its the price that sets at zero the expected order imbalance during the trading period.

To quantify this condition, we note that the expected number of buyers that arrive by time 1 is given by $r_B(\hat{p})$ and similarly, the expected number of sellers that arrive by time 1 is given by $r_S(\hat{p})$. The goal of the risk neutral specialist is therefore to choose \hat{p} such that its expected trade imbalance is zero

$$(r_B(\hat{p}) - r_S(\hat{p})) = 0. \quad (3)$$

It's clear that the specialist will thus choose \hat{p} so that $r_B(\hat{p}) = r_S(\hat{p})$. We can rewrite \hat{p} as $\hat{p} = p \times d$, where d can be thought of as a price discount applied by the specialist to attract more buyers and deter sellers, or vice versa (in the case when $d > 1$ it's a premium, but we will still refer to it as a discount). We can also rewrite the arrival rates in the following form

$$r_B(\hat{p}) = r_B(d) = \phi_B(d) \times N_B \quad (4)$$

$$r_S(\hat{p}) = r_S(d) = \phi_S(d) \times N_S, \quad (5)$$

where N_B and N_S can be interpreted as the total number of potential buyers and sellers on the market, respectively; The quantities $\phi_B(d)$ and $\phi_S(d)$ are the fraction of the buyers and sellers that the specialist attracts with his choice of discount. The assumption that lower prices (larger discount) will attract buyers and deter sellers is captured by requiring that as d decreases, $\phi_B(d)$ will increase and $\phi_S(d)$ will decrease.

Condition (3) can now be written as

$$\frac{N_B}{N_S} = \frac{\phi_S(d)}{\phi_B(d)} = \phi(d), \quad (6)$$

where $\phi(d)$ is now a monotone increasing function in d . Applying ϕ^{-1} to both sides, we get

$$d = \phi^{-1}\left(\frac{N_B}{N_S}\right) = f\left(\frac{N_B}{N_S}\right). \quad (7)$$

Thus the discount can be written as an increasing function of the ratio ($f = \phi^{-1}$ is guaranteed to be increasing since ϕ is). If we had some information about the number of buyers and sellers in the market (N_B and N_S) and the behavior of f prior to the start of trading, we could estimate the effect this would have on the price. In the empirical section of the paper, this is precisely our goal. By imposing a power law on f and estimating N_B and N_S from expected mutual fund order flows, we will show that this model allows us to explain mutual fund returns due to these price discounts.

2 Data

We use quarterly mutual fund holdings data from the CRSP Mutual Fund database ranging from 01/2003 - 12/2012. The mutual fund database does not suffer from survivorship bias. We only use equity funds (funds with Lipper Asset Code 'EQ') or funds that have at least 50% of their holdings composed of common stock. We focus on U.S. funds by excluding funds with Lipper

Objective Code ‘GL’ or ‘IF’ (global or international fund) and exclude any funds with ‘global’, ‘international’, ‘europe’ or ‘emerging’ in their names. We exclude any funds with a missing total net assets (TNA) value or missing an associated portfolio number.

We calculate monthly portfolio TNA and returns data by aggregating fund data across share classes by using share class TNA as weights, and exclude any portfolios whose total net assets are under 1 million USD. We filter out any holdings that are not long positions in common stock. This is done by excluding any holdings with a coupon rate or maturity date and anything whose share class is not 10 or 11 (representing common stock). The number of shares owned by the portfolio has to be positive, and the market value of the holding has to be non-zero.

Although holdings for most portfolios are reported at the end of each quarter, some portfolios report their holdings before the end of the quarter. To deal with the latter category of portfolios, we assume that their positions do not change from the reported date until the end of the quarter and combine these reported positions along with asset prices from the end of the quarter to construct the portfolio holdings. The database is missing holdings information for many portfolios in Q3 2010, therefore we do not use that quarter’s holdings in our empirical work.

Fund flows are calculated using the formula

$$Flow_t = \frac{TNA_t - (1 + r_t)TNA_{t-1}}{TNA_{t-1}}, \quad (8)$$

where TNA_t is the total net assets of a portfolio in period t and r_t is the return of the portfolio in period t .

To calculate stock market depths ($\lambda_k = \frac{ADP_k}{\lambda \times \sigma_k}$), we use daily stock data from the CRSP US Stock Database. Stock average daily trading volumes and daily returns standard deviations are calculated for each quarter based on that quarter’s volume and returns alone (i.e. not using volumes and returns from previous quarters). The exact value of $\tilde{\lambda}$ plays no role in our data analysis because we will use regression models and results will not depend on this exact value. Therefore, for the purpose of computing λ_k , we set $\tilde{\lambda} = 1$.

3 Three measures of vulnerability

3.1 The baseline vulnerability measure

Two nodes that are connected in the network are referred to as neighbors. During periods of mass asset liquidations (purchases), a portfolio may be subject to losses (gains) due to the trading activity of its neighbors. These losses (gains) are not uniform for each portfolio and will be more extreme for portfolios with more neighbors. In order to measure this endogenous impact on portfolio i ’s value, we define the portfolio vulnerability measure

$$VI_i = \frac{1}{P_i} \sum_{\substack{j=1 \\ j \neq i}}^N w_{ji}. \quad (9)$$

This measure corresponds to the first order effects on node i 's loss, imposed by its neighbors. To intuitively understand this measure, ϵVI_i the fraction by which portfolio i 's value will decrease (increase) if all its neighbors liquidate (expand) their portfolios by a factor of ϵ .

The measure VI_i is similar to the vulnerability measure proposed by [Greenwood et al., 2015] in their study of banking networks with common asset holdings. The difference is that in their case, banks adjust portfolios to satisfy a constant leverage constraint, whereas in the case of mutual funds, portfolio adjustment is triggered by investor flow.

Essentially, these measures are natural measures of a node's vulnerability in the network. In a study of general complex weighted networks, [Barrat et al., 2004] discuss a quantity called the vertex *strength*, defined as the sum of weights of all edges adjacent to a vertex i , which in our mutual fund network is $\sum_{\substack{j=1 \\ j \neq i}}^N w_{ji}$. To obtain VI_i , we simply scale the vertex strength by the value of the portfolio P_i – the total net assets of the portfolio – to account for the scale effect.

To illustrate the use of this first vulnerability measure, consider the following. As a result of significant market events, mutual funds can be faced with mass outflows (inflows) of investors. As documented by [Coval and Stafford, 2007], an outflow (inflow) experienced by a fund will cause the fund to liquidate (expand) its asset positions. The flow-induced liquidations (expansions) of asset positions will cause funds that were not affected by the original market events to experience losses (gains) due to their neighbors' actions.

Indeed, this is what happened when the banking sector was hit by massive losses during the financial crisis of 2007/2008. Portfolio managers with banking sector assets modified their portfolio holdings; in particular, they liquidated assets not belonging to the banking sector to pay leaving investors and in turn, they had an impact on the prices of these assets, see e.g. [Hau and Lai, 2016]. Thus, even portfolios with no financial stock holdings still lost money due to the losses propagating from the initial banking sector shock. If i and j were two such portfolios and $VI_i < VI_j$, we would expect portfolio j to incur a higher loss than i , because it would have been a more vulnerable node in the network (and thus would have been more likely to have been impacted by a neighbor's liquidations). We formalize this idea by introducing the following hypothesis:

Hypothesis: Portfolios with a higher vulnerability have lower returns in periods of mass liquidations.

In order to test this hypothesis, we examine mutual fund data obtained from the CRSP database.

3.1.1 Vulnerability measure and future returns

The most significant drops in cumulative TNA occurred in Q4 2008, shortly after the collapse of the Lehman Brothers, and Q3 2011 during which S&P downgraded the U.S. credit rating. Both quarters experienced crashes in various market indices (Dow Jones, S&P 500). Using holdings data from the end of Q3 2008 and Q2 2011, we calculate portfolio vulnerabilities and then plot portfolio returns for the following quarters of Q4 2008 and Q3 2011 against those vulnerabilities. Figures 1

and 2 show plots of returns vs. vulnerability along with the fitted regression lines. Both quarters show a significant negative relationship between vulnerability and future returns.

[Figure 1 here]

[Figure 2 here]

We point out that these quarters were chosen due to the extreme events that occurred in them. *In many other quarters, the vulnerability measure failed to display any significant relationship with future returns.* This is because the vulnerability measure is only useful conditional on a market wide event occurring that triggers mass liquidations.

3.1.2 Portfolio vulnerabilities over time

We examine the average portfolio vulnerabilities as they evolve over our sample period of 2003-2012. Figure 3 plots the average vulnerability of portfolios alongside the cumulative TNA of all portfolios. Two distinct spikes in portfolio vulnerability are visible in Q4 2008 and Q3-Q4 2011. These spikes coincide with the two largest troughs in the cumulative total net assets (TNA) of portfolios, which occurred after the Lehman Brothers collapse in September 2008 and the downgrading of the U.S. credit rating by S&P in Q3 2011.

[Figure 3 here]

Note that the average portfolio vulnerability was relatively stable before Q1 2007 and then steadily increased after Q2 2007 up to Q4 2008 (with the sharpest increase after the fall of Lehman Brothers, in Q4 2008). The cumulative TNA started to drop only after Q2 2008. This suggests that *increases in average vulnerability precede significant drops in total net assets.*

Moreover, from the spikes in average vulnerability, we note that *portfolio vulnerabilities exacerbate their increase in response to financial distress.* According to our model, there are two factors that can cause this sudden increase in vulnerabilities; portfolios can increase their exposure to other portfolios if everybody starts buying similar assets, or vulnerabilities may increase if the market depths of stocks fall.

To examine the possibility that portfolios may be purchasing similar assets in response to financial distress, we look at stock ownership distributions. Each quarter, for every stock held by portfolios in our sample, we calculate the percentage of portfolios that own that stock (stock ownership). Then we use this data to estimate the distribution, across stocks, of stock ownership. We then calculate the 5% quantile of this distribution (our choice of 5% was based on observing the distributions, choosing a different threshold will not significantly affect our conclusions). This quantile gives the percentage of stocks held by more than 5% of all portfolios. Heavier tails can occur due to increased commonality between portfolio assets, or simply because the number of portfolios is increasing.

Figure 4 displays the quantiles of the distribution, across stocks, of stock ownership along with the number of portfolios in our sample over time. The quantiles increase over time, while the number of portfolios have a linear increase only up to the crisis. In fact the number of portfolios is constant (with a slight drop at the end) over the periods 2008-2010, and 2010-2012. These figures thus suggest that the heavier tails are simply a result of more funds entering the market prior to the crisis, while in the crisis periods these heavier tails are certainly due to increases in portfolio asset commonality. *The portfolio vulnerability spikes are therefore attributed in crisis periods to plunges in the market depth of stocks in conjunction with increases in asset commonality.*

[Figure 4 here]

[Figure 5 here]

To summarize, we find that portfolio vulnerabilities spiked significantly during the periods of mass financial distress in the late quarters of 2008 and 2011. These spikes were caused by increased asset commonalities, in conjunction with decreases in market depths suggesting liquidity shortages. During these periods, portfolios were much more vulnerable to losses from widespread investor outflow.

We now consider the dual of portfolio vulnerability which is the concept of stock frailty in [Greenwood and Thesmar, 2011]: The vulnerability of a fund is obtained by aggregating the price impact of liquidations across all stocks owned by a fund; The frailty of a stock is obtained by aggregating the price impact due to liquidations across all funds owning the stock. Similarly as above, we introduce a measure of stock frailty, which we call stock vulnerability measure. We aggregate stock vulnerabilities across sectors to obtain a measure of the sector's vulnerability, due to institutional ownership.

Figure 6 plots this measure over time, for seven sectors. For each sector, the vulnerability over time follows the same pattern as the average portfolio vulnerability in Figure 3. Consistent with our previous conclusion that vulnerability spikes are attributed to plunges in stock liquidity, the financial sector changes from the least vulnerable sector in 2003 to the most vulnerable sector during the financial crisis. These dual measures of vulnerability of stocks (sectors), rather than funds, may be useful tools for comparing stocks (sectors), at any given time, in terms of their vulnerabilities. Our results suggest a type of stickiness of the vulnerabilities: for all sectors the vulnerabilities change over time, yet the ranking of sectors in terms of vulnerabilities, *for the non-financial* sectors, remained relatively stable over the considered time frame.

[Figure 6 here]

3.2 Flow-adjusted vulnerability

In this section, we refine the vulnerability measure by incorporating fund-specific investor flow. If a fund is expected to expand its existing asset positions, then its actions will benefit its neighbors

and this fund can be considered as a ‘good’ neighbor. Conversely, funds that liquidate assets are ‘bad’ neighbors. The vulnerability measure is limited because it is a simple measure of a portfolio’s centrality and does not distinguish between good and bad neighbors. Thus, to be useful, it implicitly requires the same trading direction throughout the whole network (i.e. everyone liquidates or everyone purchases), so that everyone is either a good neighbor or a bad neighbor. Although this mass trading behavior is more likely to occur during a financial crisis or a large boom in the market, most of the time fund behavior is heterogeneous and VI is no longer as useful. The measures introduced below addresses this limitation.

In the remainder of this section we assume that the price impact is linear, as in the definition of the vulnerability measure.

Given a network of funds at the start of a quarter, we set F_i to be the flow *to be experienced* by fund i over this quarter. The fund’s flow-induced trading during this quarter will then have an impact on asset prices. To capture the effect of a fund on its neighbors, we impose two simplifying assumptions on the trading behavior of a fund:

1. **No new assets** are purchased as a result of inflows, existing positions are expanded instead.
2. **Proportional buying/selling:** when a fund experiences an outflow $F_i < 0$, it liquidates an equal fraction of each asset and when a fund experiences an inflow $F_i > 0$, it expands its positions proportionally.

Our assumption of proportional buying/selling of assets is unlikely to hold in reality. Indeed, [Hau and Lai, 2016] found that a fund experiencing outflows will raise money by liquidating its best performing assets. Yet, they find no evidence for an important interaction among stock liquidity and holding reductions. Attempting to determine exact liquidation policies is an interesting problem in itself; however, it is well outside the scope of this paper. We demonstrate that even under a simple policy such as this, our vulnerability measures are still a significant factor in explaining returns.

At the end of the quarter, the impact of portfolio j ’s flow-induced trading on portfolio i is

$$\frac{1}{(1 + F_i)P_i} \sum_{k=1}^K \frac{F_j \beta_{kj}}{\lambda_k} (1 + F_i) \beta_{ki} s_k = \frac{F_j w_{ji}}{P_i}. \quad (10)$$

This impact is now either positive or negative, all depending on the sign of F_j .

We now define the **flow-adjusted vulnerability (FAV) measure** for fund i as the sum of the impacts of all its neighbors (including itself). It is the percentage change in portfolio i ’s value due to flow-induced trading

$$FAV_i = \frac{1}{P_i} \sum_{j=1}^N F_j w_{ji}. \quad (11)$$

The FAV can be calculated using the portfolio holdings network at the start of each quarter and the investor flows over that quarter. In section 3.4, we show that a portfolio’s FAV is positively correlated with its returns over this quarter. *Unlike the vulnerability measure VI , the FAV can*

be used to explain returns at all times and not only in periods of (near) uniform trading behavior. However, contrary to the *VI*, the *FAV* cannot predict returns, as computing the *FAV* requires knowing the investor flow during a quarter ahead of time. Therefore, the *FAV* for a given quarter can only be computed at the end of the quarter, once the flows (and returns) are already known.

Our measure is similar to the flow induced trading (FIT) measure proposed by [Lou, 2012]. The difference lies in the way stock liquidity is accounted for. The *FIT* uses the total number of stocks owned by all funds as a measure of each stock's liquidity, motivated by [Gompers and Metrick, 2001] who show that mutual funds' holdings are skewed toward liquid stocks. We use information about a stock's trading volume and volatility.

The main limitation of the *FAV* measure is that the price impact is linear. In the next section, we address this shortcoming by refining the *FAV* using the model discussed in Section 1.2.

3.3 Refining the flow-adjusted vulnerability measure

To replace the assumption of linear price impact for each stock, we shall use the model from Section 1.2. In the model, a stock with fundamental value p is traded continuously during some time period $[0, 1]$ at a discounted price $d \times p$. We let the time period $[0, 1]$ represent one quarter and we let the fundamental price be the price at the start of the quarter. The price at the end of the quarter will therefore be the discounted price $d \times p$. We showed that the discount can be written as

$$d = f\left(\frac{N_B}{N_S}\right), \quad (12)$$

where f is an increasing function and N_B and N_S represent the total number of potential buyers and sellers on the market.

We estimate N_B and N_S using the total number of stock bought and sold by all funds during the quarter. According to our assumptions, the number of shares of a stock that a single fund will trade is equal to its investor flow during the quarter multiplied by the number of shares of that stock it owns. Therefore, for stock k , we set the estimators

$$\hat{N}_{B,k} = \sum_{i=1}^N \beta_{ki} F_i \mathbb{1}_{\{F_i > 0\}}, \quad (13)$$

$$\hat{N}_{S,k} = \sum_{i=1}^N \beta_{ki} F_i \mathbb{1}_{\{F_i < 0\}}. \quad (14)$$

This implies that the impact of trading (captured by d) is simply a function of the ratio of the order inflow to the order outflow.

Next, we need to impose some conditions on the form of f . In particular, we assume that f is a power-law of the form

$$f(x) = x^\alpha. \quad (15)$$

For the purposes of performing regression analysis, we set $\alpha = 1/3$; however, our results do not depend qualitatively on this particular choice of α .

We can now write the return of the stock at the end of a quarter as

$$return_k = \frac{d_k \times s_k - s_k}{s_k} = d_k - 1 = f \left(\frac{N_{B,k}}{N_{S,k}} \right) - 1 \approx \left(\frac{\hat{N}_{B,k}}{\hat{N}_{S,k}} \right)^{1/3} - 1. \quad (16)$$

By aggregating these returns across all stocks in a portfolio, we obtain the refinement of the flow-adjusted vulnerability measure

$$FAV_i^* = \frac{1}{P_i} \sum_{k=1}^K \beta_{ki} s_k \left(\left(\frac{\hat{N}_{B,k}}{\hat{N}_{S,k}} \right)^{1/3} - 1 \right).$$

This refined measure can be interpreted as a measure of **vulnerability to order imbalance**.

Note that in the above measure the neighbors (in the network of asset holdings) do not appear explicitly. The effects of their distress on a given portfolio are captured by the price discount factor, which is driven by the imbalance of the supply and demand. This measure directly aggregates the discount (weighted by the relative position) across all stocks belonging to a portfolio to obtain the impact of trading on the value of an entire portfolio. This can be viewed as a refinement of the *FAV* measure, because instead of using a linear price impact function, we use the model we proposed in Section 1.2 to measure how order flow will impact prices.

3.4 Explaining mutual fund performance

We compare the ability of both the *FAV* and *FAV*^{*} measures to explain mutual fund returns. Table 1 displays the results of regressing a fund's quarterly returns against the *FAV* and *FAV*^{*} measures computed for the fund at the start of that quarter. We run a Fama-Macbeth regression: in the first stage we run N regressions, for each mutual fund, to determine the factor loadings. In the second stage we run T (the number of quarters) cross-sectional regressions of returns against the estimated loadings to determine the risk premium for each factor. We use Newey-West corrections of four lags.

As expected, both measures are positively correlated with returns. Furthermore, we find that after accounting for the *FAV*^{*} measure, the *FAV* is no longer significant in explaining returns. These results indicate that the *FAV*^{*} is indeed a refinement of the *FAV*, and suggest that the ratio of order inflows to outflows may be a more suitable measure of price impact than simply considering net order flow.

Computing the *FAV* and *FAV*^{*} measures for a quarter requires using fund flows for that quarter, which are certainly expected to be correlated with the returns for the quarter. Therefore, any relationship between the two vulnerability measures and fund returns may occur simply because we used concurrent flows to compute the measure. The most important point in the validity of the regressions is therefore introducing the concurrent flows as control variables. Unsurprisingly, these

are positively correlated with the returns in the same quarter.

We find that the two measures of vulnerability maintain their explanatory power after controlling for the concurrent flows.

[Table 1 here]

4 Conclusion

In this paper we measure the interrelations due to common asset holdings. The aggregation of asset holdings in a network structure depends on the model of distress propagation, i.e, the sequence: the initial shock, the liquidations in response to the initial shock, the effect of liquidations on other participants. We construct a model for the price impact of trading, in which demand and supply depend on the asset holdings and fund flows.

The network representation is useful to derive measures of vulnerability of funds to the shocks experienced by their neighbors in the network. We find, using mutual fund data, that the vulnerability index is useful in predicting returns in periods of mass liquidations. In such periods, we can identify vulnerable funds based on asset holdings and the liquidity characteristics of the stocks.

The vulnerability index VI is only useful during periods of mass liquidations. However, our flow-adjusted measure of vulnerability to order imbalance FAV , is shown to be correlated with returns throughout all our sample period.

We have shown how to measure vulnerabilities in a network of mutual funds. Our model only considers mutual funds and any imbalance comes from mutual fund selling pressure. In the companion paper [Guo et al., 2016], we show that similar network effects due to common asset holdings arise for the hedge fund industry. In addition, we analyze the topology of the hedge fund network and demonstrate the existence of large clusters and heavy tails in the degree distribution. As mentioned in the introduction, there are several available studies of networks in banks' portfolios, where liquidations are either driven by repo investor outflows or leverage targeting. It may be the case that other network reasons are important (say, interlocking boards between firms). However, we believe that networks of common asset holdings play a primary role in returns and this fact is valid for several industries. In past years (not included in the study), we have witnessed unprecedented growth in passive investing and we conjecture that this would exacerbate the network effects due to common asset holdings. The fact remains that all available studies and stress tests consider only one industry: mutual funds, hedge funds, banks. Having accumulated evidence of network effects due to common asset holdings in several industries, the critical next step is a cross-industry study of network-based vulnerabilities to yield more precise indicators.

In this paper we focused on vulnerabilities. In the other direction, the network of common asset holdings is useful to identify funds that are likely to have large network externalities on the other funds, as well as clusters of funds which collectively pose systemic risk. An important direction that emerges out of the current study is to understand how funds should optimally allocate their funds in order to manage their vulnerability.

References

[Adrian and Shin, 2010] Adrian, T. and Shin, H. S. (2010). Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418–437.

[Allen et al., 2012] Allen, F., Babus, A., and Carletti, E. (2012). Asset commonality, debt maturity and systemic risk. *Journal of Financial Economics*, 104(3):519–534.

[Almgren et al., 2005] Almgren, R., Thum, C., Hauptmann, E., and Li, H. (2005). Direct estimation of equity market impact. *Risk*, pages 57–62.

[Amihud, 2002] Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56.

[Amini et al., 2016] Amini, H., Cont, R., and Minca, A. (2016). Resilience to contagion in financial networks. *Mathematical finance*, 26(2):329–365.

[Amini et al., 2015] Amini, H., Filipović, D., and Minca, A. (2015). Systemic risk and central clearing counterparty design. *Available at SSRN*.

[Amini and Minca, 2013] Amini, H. and Minca, A. (2013). Mathematical modeling of systemic risk. In Kranakis, E., editor, *Advances in Network Analysis and its Applications*, volume 18 of *Mathematics in Industry*, pages 3–26. Springer Berlin Heidelberg.

[Babus and Allen, 2009] Babus, A. and Allen, F. (2009). Networks in finance. In Kleindorfer, P. and Wind, J., editors, *Network-based Strategies and Competencies*, pages 367–382. Wharton School Publishing.

[Barigozzi and Brownlees, 2013] Barigozzi, M. and Brownlees, C. (2013). Nets: network estimation for time series. *Available at SSRN*.

[Barrat et al., 2004] Barrat, A., Barthélémy, M., Pastor-Satorras, R., and Vespignani, A. (2004). The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences of the United States of America*, 101(11):3747–3752.

[Battiston et al., 2012] Battiston, S., Delli Gatti, D., Gallegati, M., Greenwald, B., and Stiglitz, J. E. (2012). Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk. *Journal of Economic Dynamics and Control*, 36(8):1121–1141.

[Billio et al., 2012] Billio, M., Getmansky, M., Lo, A. W., and Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535–559.

[Blocher, 2016] Blocher, J. (2016). Network externalities in mutual funds. *Journal of Financial Markets*, 30:1–26.

[Blume et al., 2011] Blume, L., Easley, D., Kleinberg, J., Kleinberg, R., and Tardos, E. (2011). Network formation in the presence of contagious risk. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 1–10. ACM.

[Cabrales et al., 2017] Cabrales, A., Gottardi, P., and Vega-Redondo, F. (2017). Risk sharing and contagion in networks. *The Review of Financial Studies*, 30(9):3086–3127.

[Caccioli et al., 2015] Caccioli, F., Farmer, J. D., Foti, N., and Rockmore, D. (2015). Overlapping portfolios, contagion, and financial stability. *Journal of Economic Dynamics and Control*, 51:50–63.

[Caccioli et al., 2014] Caccioli, F., Shrestha, M., Moore, C., and Farmer, J. D. (2014). Stability analysis of financial contagion due to overlapping portfolios. *Journal of Banking & Finance*, 46:233–245.

[Capponi and Larsson, 2015] Capponi, A. and Larsson, M. (2015). Price contagion through balance sheet linkages. *The Review of Asset Pricing Studies*, 5(2):227–253.

[Cohen et al., 2005] Cohen, R. B., Coval, J. D., and Pástor, L. (2005). Judging fund managers by the company they keep. *The Journal of Finance*, 60(3):1057–1096.

[Cohen-Cole et al., 2011] Cohen-Cole, E., Patacchini, E., and Zenou, Y. (2011). Systemic risk and network formation in the interbank market. *CEPR Discussion Paper No. DP8332*.

[Cont and Wagalath, 2013] Cont, R. and Wagalath, L. (2013). Running for the exit: Distressed selling and endogenous correlation in financial markets. *Mathematical Finance*, 23(4):718–741.

[Coval and Stafford, 2007] Coval, J. and Stafford, E. (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics*, 86(2):479 – 512.

[Diebold and Yilmaz, 2013] Diebold, F. X. and Yilmaz, K. (2013). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, *forthcoming*.

[Duarte and Eisenbach, 2015] Duarte, F. and Eisenbach, T. M. (2015). Fire-sale spillovers and systemic risk. *FRB of New York Staff Report*, (645).

[Elliott et al., 2014] Elliott, M., Golub, B., and Jackson, M. O. (2014). Financial networks and contagion. *American Economic Review*, 104(10):3115–53.

[Gai and Kapadia, 2010] Gai, P. and Kapadia, S. (2010). Contagion in Financial Networks. *Proceedings of the Royal Society A*, 466(2120):2401–2423.

[Gao et al., 2017] Gao, G. P., Moulton, P. C., and Ng, D. T. (2017). Institutional ownership and return predictability across economically unrelated stocks. *Journal of Financial Intermediation*, 31:45 – 63.

[Glasserman and Young, 2015] Glasserman, P. and Young, H. P. (2015). How likely is contagion in financial networks? *Journal of Banking & Finance*, 50:383–399.

[Gompers and Metrick, 2001] Gompers, P. A. and Metrick, A. (2001). Institutional investors and equity prices. *The Quarterly Journal of Economics*, 116(1):229–259.

[Gorton and Metrick, 2012] Gorton, G. and Metrick, A. (2012). Securitized banking and the run on repo. *Journal of Financial Economics*, 104(3):425–451.

[Gourieroux et al., 2012] Gourieroux, C., Héam, J.-C., and Monfort, A. (2012). Bilateral exposures and systemic solvency risk. *Canadian Journal of Economics/Revue canadienne d'économique*, 45(4):1273–1309.

[Greenwood et al., 2015] Greenwood, R., Landier, A., and Thesmar, D. (2015). Vulnerable banks. *Journal of Financial Economics*, 115(3):471–485.

[Greenwood and Thesmar, 2011] Greenwood, R. and Thesmar, D. (2011). Stock price fragility. *Journal of Financial Economics*, 102(3):471–490.

[Guo et al., 2016] Guo, W., Minca, A., and Wang, L. (2016). The topology of overlapping portfolio networks. *Statistics & Risk Modeling*, 33(3-4):139–155.

[Hau and Lai, 2016] Hau, H. and Lai, S. (2016). The role of equity funds in the financial crisis propagation. *Review of Finance*, 21(1):77–108.

[Kara et al., 2015] Kara, G. I., Tian, M. H., and Yellen, M. (2015). Taxonomy of studies on interconnectedness. *Available at SSRN*.

[Khandani and Lo, 2011] Khandani, A. E. and Lo, A. W. (2011). What happened to the quants in august 2007? evidence from factors and transactions data. *Journal of Financial Markets*, 14(1):1–46.

[Kyle, 1985] Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, pages 1315–1335.

[Kyle and Obizhaeva, 2016] Kyle, A. S. and Obizhaeva, A. A. (2016). Market microstructure invariance: Empirical hypotheses. *Econometrica*, 84(4):1345–1404.

[Lou, 2012] Lou, D. (2012). A flow-based explanation for return predictability. *Review of Financial Studies*, 25(12):3457–3489.

[Musmeci et al., 2015] Musmeci, N., Di Matteo, T., and Aste, T. (2015). Risk diversification: a study of persistence with a filtered correlation-network approach. *Journal of Network Theory in Finance*, 1(1):1–22.

[Stavroglou et al., 2017] Stavroglou, S., Pantelous, A. A., Soramaki, K., and Zuev, K. (2017). Causality networks of financial assets. *Journal of Network Theory in Finance*, 3(2):17–67.

[Zawadowski, 2013] Zawadowski, A. (2013). Entangled financial systems. *Review of Financial Studies*, 26(5):1291–1323.

Table 1: Fama-Macbeth regressions of future fund returns with Newey-West corrections of four lags. The dependent variable is fund return over quarter t , the FAV^* and FAV measures are computed using the asset holdings at the start of quarter t and the fund flows over quarter t . We control for the log total net assets owned by a fund at the end of the previous quarter, the log total number of shares owned by the fund at the end of the previous quarter and the return of the fund in the previous quarter. We also control for fund flows over quarter t to ensure that the explanatory power of the two measures is not a simple consequence of the correlation of returns with concurrent flows. When computing the values of FAV^* and FAV , we exclude funds whose inflows are greater than 100%.

	(1)	(2)	(3)
FAV_t^*		0.0683*** (4.85)	0.0705*** (4.36)
FAV_t	0.463*** (4.38)		-0.00706 (-0.08)
$\log(TNA_{t-1})$	0.00101 (0.69)	0.000635 (0.41)	0.000623 (0.39)
$\log(Shares_{t-1})$	-0.000800 (-0.57)	-0.000352 (-0.24)	-0.000333 (-0.22)
$flow_t$	0.0230*** (4.55)	0.0162** (3.49)	0.0155** (3.34)
$return_{t-1}$	0.000742 (0.02)	-0.0233 (-0.51)	-0.0251 (-0.56)
Constant	0.0227 (1.48)	0.0205 (1.29)	0.0203 (1.23)
Sample Size	87257	87257	87257
R^2	0.1746	0.2032	0.2105

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

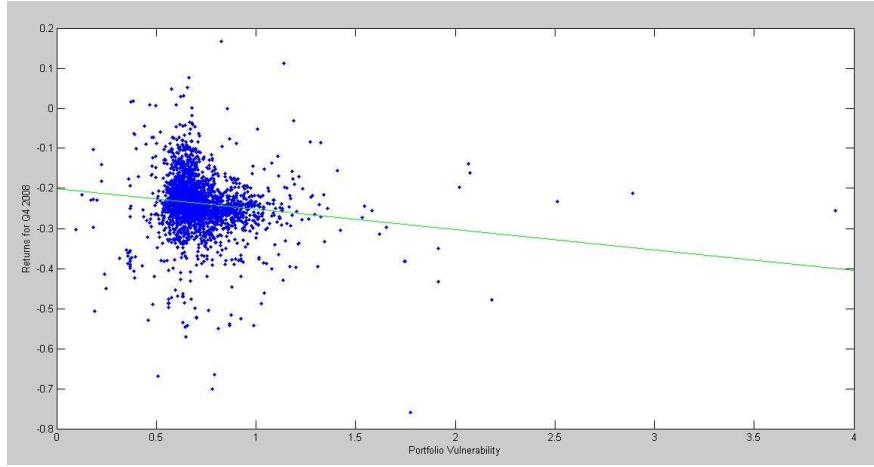


Figure 1: Portfolio returns for Q4 2008 plotted against portfolio vulnerability calculated at the start of the quarter.

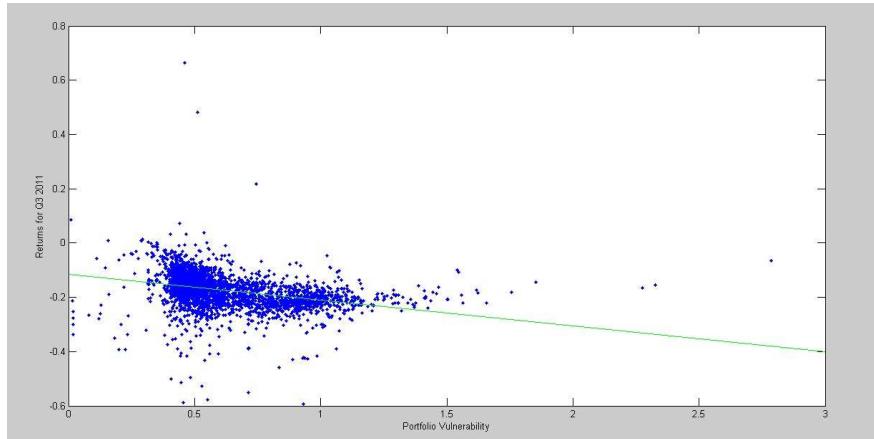


Figure 2: Portfolio returns for Q3 2011 plotted against portfolio vulnerability calculated at the start of the quarter.

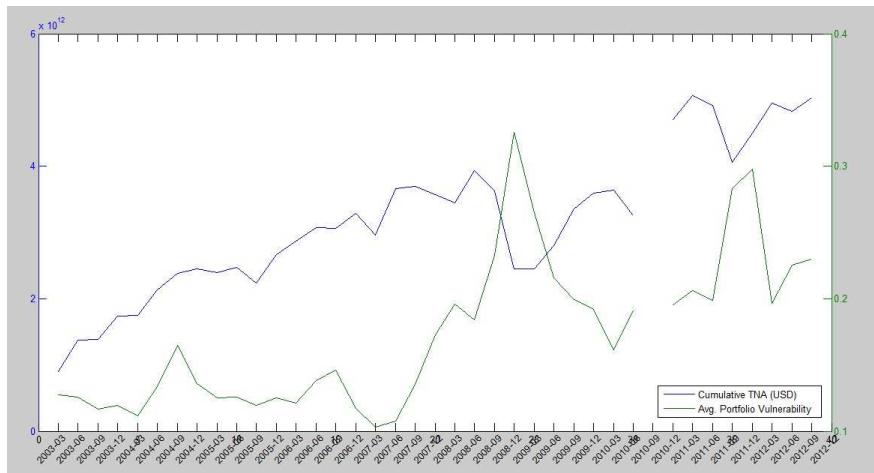


Figure 3: Average portfolio vulnerability VI plotted alongside the cumulative TNA of all portfolios. Data for 09/2010 is missing.

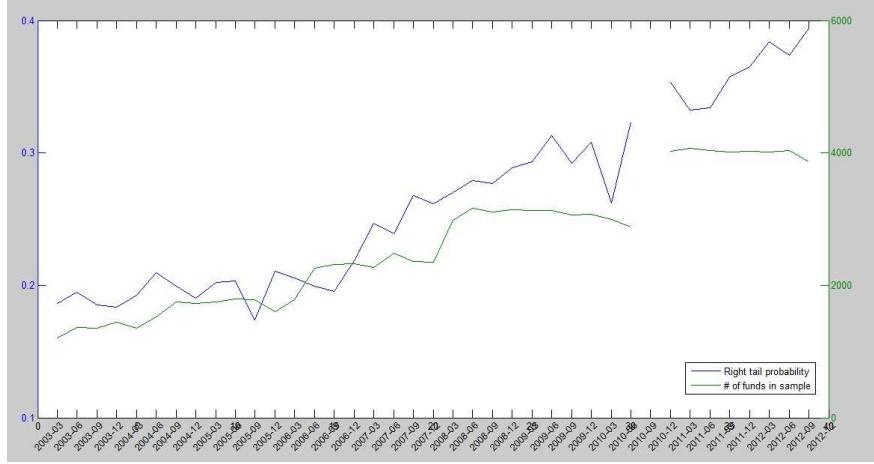


Figure 4: Time series of the 5% quantile of the distribution of stock ownership and the number of portfolios used in the sample. Data for 09/2010 and 12/2012 is missing.

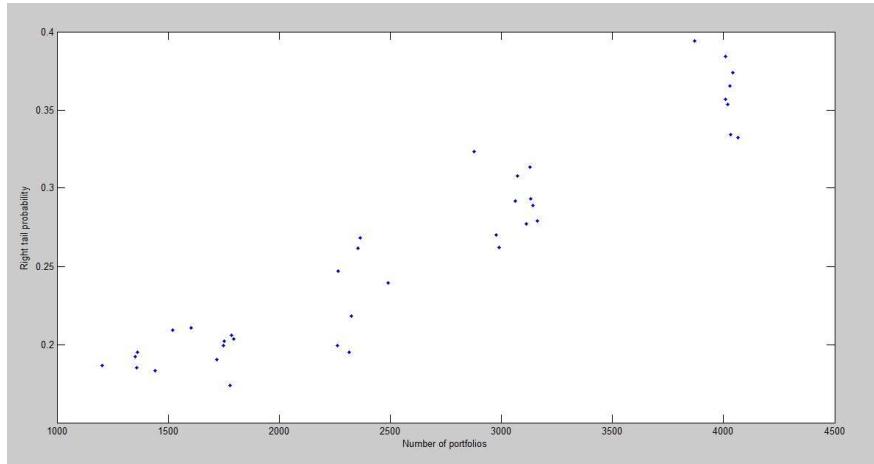


Figure 5: Stock ownership tail probabilities against the sample size. A clear linear relationship is evident. Data for 09/2010 and 12/2012 is missing.

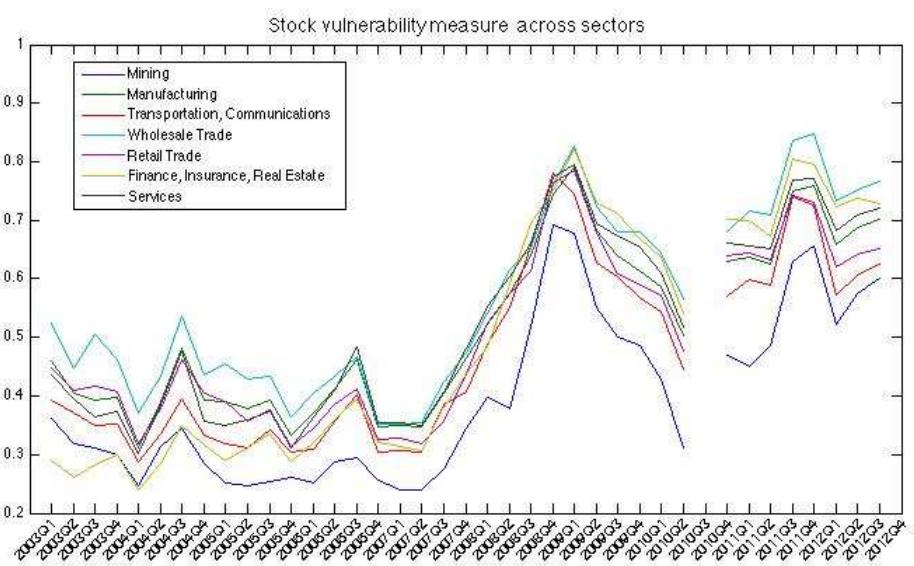


Figure 6: Stock vulnerability measure across sectors. Data for 09/2010 and 12/2012 is missing.