Journal of Banking and Finance 155 (2023) 106989

Contents lists available at ScienceDirect

Journal of Banking and Finance

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

Assessing and mitigating fire sales risk under partial information

Raymond Ka-Kay Pang, Luitgard Anna Maria Veraart *

London School of Economics and Political Science, Department of Mathematics, Houghton Street, London, WC2A 2AE, United Kingdom

ARTICLE INFO

JEL classification:
G21
G28
G32
G33
G20

Keywords:
Systemic risk
Fire sales
Stress testing
Financial networks
Matrix reconstruction
Policy interventions

ABSTRACT

We consider the problem of assessing and mitigating fire sales risk for banks under partial information. Using data from the European Banking Authority's stress tests, we consider the matrix of asset holdings of different banks. We first analyse fire sales risk under both full and partial information using different matrix reconstruction methods. We then investigate how well some policy interventions aimed at mitigating fire sales risk perform if they are applied based on only partial information. We find that even under partial information, using suitable network reconstruction methods to decide on policy interventions can significantly mitigate risk from fire sales. Furthermore, we show that some interventions based on reconstructed networks significantly outperform ad hoc methods that decide on interventions only based on the size of an institution and do not account for overlapping portfolios.

1. Introduction

Fire sales pose a key threat to financial stability since they can significantly amplify initial losses. They are one of the main channels of systemic risk. During the 2007-2008 Global Financial Crisis (GFC) amplification mechanisms played a major role. It was estimated that $ 300bn of subprime mortgage-related losses were turned into over $ 2.5 trillion of potential write-downs in the global banking sector within a year (Brazier, 2017). Therefore it is a key concern for financial regulators to identify any potential channels of systemic risk and find tools and mechanisms to mitigate their impact. In this paper, we will focus on fire sales.

While significant progress on modelling fire sales has been made, the models proposed usually assume that the asset holdings of the financial institutions are observable, see e.g., Shleifer and Vishny (2010); Cont and Wagalath (2013); Greenwood et al. (2015); Cont and Schaanning (2017, 2019). In practice, however, often only partial information about the asset holdings is available. Usually, regulators have only detailed information on the banks that they regulate and not beyond.

In this paper, we show how one can both assess and mitigate fire sales risk under partial information. We consider a matrix X, where each element X_{nk} represents the amount of asset k that bank n holds (in EUR). We are interested in situations in which these individual positions are not observable but the corresponding column and row sums of X representing the total market capitalisation of asset k and the total assets of bank n are observable. To conduct stress testing under partial information, we use matrix reconstruction methods by (Upper and Worms, 2004; Anand et al., 2015; Cimini et al., 2015; Gandy and Veraart, 2017, 2019) to reconstruct the asset holding matrix X from the observed row and column sums.

Our paper makes two main contributions: First, it conducts a horse race between different network reconstruction methods and compares their performance in quantifying fire sales risk in financial stress tests under partial information. Second, we show that there are clear benefits of using suitable network reconstruction techniques not just for quantifying fire sales risk but also for mitigating it. In particular, we show that policy interventions based on suitable network reconstruction methods can significantly outperform ad hoc policy interventions that do not account for the interconnectedness of financial institutions. We identify which network reconstruction methods are best suited to use for policy interventions to mitigate fire sales risk. To the best of our knowledge, our analysis is the first that considers the mitigation of fire sales risk under partial information.

---

* Corresponding author.
E-mail address: l.veraart@lse.ac.uk (L.A.M. Veraart).

https://doi.org/10.1016/j.jbankfin.2023.106989
Received 25 April 2021; Accepted 31 August 2023

We consider the modelling framework and risk measures developed by Greenwood et al. (2015) to assess fire sales risk. Their key assumption is that banks aim to maintain their target leverage, i.e., the ratio of debt to equity.[1] Empirical evidence for this behaviour has been provided by Adrian and Shin (2010). We conduct a stress testing exercise using empirical data from the European Banking Authority (EBA) that they used in their 2011 and 2016 EBA stress tests. The EBA has collected full information on the asset holdings of the banks that participated in these stress tests. This allows us to compare results under partial information to results under full information.

We find that all matrix reconstruction methods considered are able to reproduce the general trend, namely, that fire sales risk was lower in 2016 than in 2011. We show how the performance of the different network reconstruction methods applied to quantifying fire sales risk depends on the stress scenarios. Overall, we find that reconstruction methods attempting to approximate the distribution of the underlying network such as Cimini et al. (2015); Gandy and Veraart (2017, 2019) are better suited to assess fire sales risk from partial information than optimisation-based network reconstruction methods such as the ones by Upper and Worms (2004); Anand et al. (2015).

Next, we consider several policy interventions to reduce fire sales risk. In particular, we are interested in evaluating whether it is possible to conduct policy interventions at an early stage of a fire sales cascade to mitigate losses. At such an early stage, it is unlikely that full information on the underlying asset holdings is available. We therefore again conduct two types of analyses: First, we assume that full information on the underlying asset holding network is available and use this full information to decide on policy interventions. Second, we base all policy intervention decisions on asset holding networks that were reconstructed from partial information and we compare the outcome of the system to the outcome under full information. We compare ad hoc strategies that do not use network reconstruction to strategies that use network reconstruction to decide on policy interventions. We find that there are clear benefits of using network reconstruction over ad hoc methods. In particular, we find that the Bayesian approach to network reconstruction by Gandy and Veraart (2017, 2019) is particularly successful when used for policy interventions to mitigate fire sales risk.

Another contribution of our analysis is to show theoretically how the fire sales risk measures introduced by Greenwood et al. (2015) depend on the individual entries of the asset holdings matrix and how much they only depend on the aggregate information, i.e., the row and column sums of the asset holding matrix. We provide theoretical results that show that for some stress scenarios only very limited information on the underlying asset holdings matrix is required to either assess the related fire sales risk or to make a meaningful intervention to mitigate these risks. In particular, we show that determining the initial spread of losses via connected portfolios requires far less information than determining who is eventually negatively affected by fire sales. This explains why policies aimed at mitigating the initial round of fire-sales losses can still be successfully applied under partial information.

The structure of the paper is as follows. In Section 2, we describe the modelling framework by Greenwood et al. (2015) for quantifying fire sales. In Section 3, we provide theoretical results on how the fire sales measure depended on the asset holding matrices. Furthermore, we provide some background on the network reconstruction methods used. In addition, we describe the stress testing data by the European Banking Authority (EBA) that we use in our empirical analysis. We compare the performance of different matrix reconstruction methods in replicating fire sales measures by Greenwood et al. (2015) for the EBA data. In Section 4, we present our results on the performance of policy interventions based on both full and partial information. Finally, Section 5 concludes.

1.1. Related literature

Our analysis is based on the framework for quantifying fire sales risk by Greenwood et al. (2015). In contrast to Greenwood et al. (2015) who assume full knowledge of the asset holdings, we assume that only partial information of the asset holdings is available. Hence, we conduct a two-step analysis. In the first step, we reconstruct the network of asset holdings from partial information. In a second step, we apply the measures by Greenwood et al. (2015) to the reconstructed networks and compare the results to those obtained under full information.

We will consider a range of reconstruction methods to obtain an estimate of the asset holding matrix based on partial information. The goal of all these methods is to reconstruct the individual entries of the matrix from given row and column sums. We will consider the network reconstruction methods proposed by Upper and Worms (2004), Anand et al. (2015), Cimini et al. (2015), and Gandy and Veraart (2017, 2019) in our analysis and compare their performance. We provide more details on them in Appendix A.

Several papers have compared the performance of different matrix reconstruction methods. Gandy and Veraart (2019) have compared the Bayesian method by Gandy and Veraart (2017) (and some extensions) to the approaches by Cimini et al. (2015) and Upper and Worms (2004) using network data of Credit Default Swaps exposures where the reference entity was a UK institution. They found that the Bayesian method outperformed alternative reconstruction methods using a wide range of matrix comparison measures. Anand et al. (2018) compares a wide range of methods (not including the Bayesian approach by Gandy and Veraart (2017, 2019)) by applying them to data from 25 different markets from 13 jurisdictions. They find that it depends on the characteristics of the networks which method works best for its reconstruction. Among the probabilistic methods, they found that the method by Cimini et al. (2015) worked best. Lebacher et al. (2019) compare several network reconstruction methods including, e.g., entropy, Bayesian and gravity (a regularised entropy method with a penalising factor) type reconstruction methods using SWIFT data. Their paper finds that the performance of the reconstruction methods is dependent on the type of network being reconstructed, arriving at a similar conclusion as in Anand et al. (2018).

The papers Di Gangi et al. (2018), Squartini et al. (2017) and Ramadiah et al. (2020) conduct a similar analysis as we do - in the sense that they apply network reconstruction methods for assessing fire sales. Di Gangi et al. (2018) focuses on variations of the maximum entropy method for the network reconstruction and uses the Greenwood et al. (2015) measures to quantify fire sales risk using data from the USA. Squartini et al. (2017) apply the model by Cimini et al. (2015) to reconstruct bipartite networks of asset holdings. They use data on security holdings by the European Central Bank and use the fire sales measure by Greenwood et al. (2015) to consider a relative systemicness index in evaluating the reconstructed matrices. Ramadiah et al. (2020) evaluate a range of reconstruction methods with systemic risk indicators for fire sales risk. They use data from bank-firm interactions in Japan and analyse the effect of aggregation on the performance of reconstructed matrices.

Our paper deviates from the existing literature by first analysing the Bayesian reconstruction methods by Gandy and Veraart (2017, 2019) in the context of fire sales risk and by comparing them to other approaches. But second and most importantly, we also analyse the effect of different policy interventions under both full and partial information.

---

[1] To see the effect of leverage targeting, consider a bank with a stylised balance sheet whose asset side consists of securities and whose liabilities side consists of debt and equity. So, if market stress leads to a decrease in the value of securities, i.e., the assets, then on the liabilities side of the balance sheet the value of the equity decreases, and hence the leverage increases. To move back to the target leverage a bank can now sell assets to pay off some of its debt. Hence, a decrease in asset values can trigger (fire) sales. If large quantities of assets are liquidated, this creates a price impact, i.e., their prices decrease. This forces other institutions to re-evaluate their portfolios, which might lead to further rounds of deleveraging and price impacts.

Table 1
Balance sheet of bank $n \in \mathcal{N}$ at time t.

Assets	Liabilities
assets α_{nt}	debt d_{nt} equity e_{nt}

So far policy interventions have only been studied under full information. Shleifer and Vishny (2010) considers the effect of credit easing on fire sales risk in comparison with other policies. Capponi and Larsson (2015) builds on the systemic measures from Greenwood et al. (2015) and Duarte and Eisenbach (2021) propose a liquidation strategy to reduce the systemic risk of the network. Greenwood et al. (2015) consider a wide selection of policy interventions to mitigate fire sales risk. But all these papers have not considered such interventions in only partially observed financial networks which is what we do here.

2. Fire sales in financial networks

We now describe the modelling framework for stress testing and assessing fire sales risk by Greenwood et al. (2015).

2.1. The financial market

The financial market consists of $N \in \mathbb{N}$ banks and $K \in \mathbb{N}$ assets, the set of banks is denoted by $\mathcal{N} = \{1, \dots, N\}$ and the set of assets is denoted by $\mathcal{S} = \{1, \dots, K\}$. The main model considers two periods with time indices $t = 1, 2$.[2]

We denote by $X = (X_{nk})_{n\in\mathcal{N},k\in\mathcal{S}} \in [0,\infty)^{N\times K}$ the asset holdings matrix at time $t = 1$, where X_{nk} represents the amount of asset $k \in \{1 \dots, K\}$ that bank $n \in \{1, \dots, N\}$ holds in million EUR. Furthermore, we consider the row and column sums of X given by

$$\alpha_{n1} = \sum_{k=1}^{K} X_{nk} \quad \forall n \in \mathcal{N}, \qquad c_k = \sum_{n=1}^{N} X_{nk} \quad \forall k \in \mathcal{S}, \tag{1}$$

and refer to α_{n1} as the total assets of bank n at time $t = 1$ and to c_k as the total capitalisation of asset k. (Strictly speaking, c_k is the total capitalisation of asset k among the nodes $\mathcal{N}$, but since we do not consider other financial institutions beyond those in $\mathcal{N}$ we will not make this distinction.)

We also define the matrix of portfolio weights denoted by $M = (m_{nk})_{n\in\mathcal{N},k\in\mathcal{S}} \in \mathbb{R}^{N\times K}$, where $m_{nk} = X_{nk}/\alpha_{n1}$ i.e., m_{nk} describes the weight of asset k within the total asset portfolio of bank n. In particular, for all $n \in \mathcal{N}$, $\sum_{k=1}^{K} m_{nk} = 1$.

We consider a stylised balance sheet, see Table 1, in which for each bank $n \in \mathcal{N}$ its time t debt is denoted by d_{nt} and its time t equity is denoted by e_{nt}. Then, its total assets at time t are given by $\alpha_{nt} = e_{nt} + d_{nt}$ and the time t leverage of bank n is given by

$$b_{nt} = \frac{d_{nt}}{e_{nt}} = \frac{\alpha_{nt} - e_{nt}}{e_{nt}}. \tag{2}$$

2.2. The stress test and fire sale mechanism

As part of a stress testing exercise, Greenwood et al. (2015) assume that at time $t = 1$ there is a negative shock to (some of the) assets. We denote by $F_1 = (f_{11}, \dots, f_{K1})^\top$, with $f_{k1} \le 0$ for all $k \in \mathcal{S}$, the shock vector which is a vector of non-positive net asset returns. The unlevered return on the portfolios of the N banks, denoted by $R_1 = (R_{11}, \dots, R_{N1})^\top \in \mathbb{R}^N$, is then given by $R_1 = MF_1$. In particular, $R_{n1} = \sum_{k=1}^{K} m_{nk} f_{k1}$ for all $n \in \mathcal{N}$.

Greenwood et al. (2015) assume that, in response to such a negative shock, banks will sell assets to return to their target (original) leverage b_{n1}.[3]

Furthermore, they assume that banks sell assets proportionally to their existing holdings determined by the matrix of portfolio weights M.[4] In addition, they assume that the sale of assets will have a linear price impact, modelled by K parameters $l_1, \dots, l_K \in [0,\infty)$. In particular, the parameter $-l_k \le 0$ models the negative shock to asset $k \in \mathcal{S}$ per million EUR of asset k sold.[5] In particular, the fire sale in one specific asset does not affect prices in any other assets. This price impact represents a second shock to the market which could in principle lead to further deleveraging, but we do not consider later rounds of deleveraging in this paper.

2.3. Measuring fire sales risk

We now define the measures for quantifying fire sales risk proposed by Greenwood et al. (2015) using slightly different notation in some places.

The **aggregate vulnerability**, denoted by $\mathcal{AV}$, is the total banks' equity lost due to deleveraging following an initial shock F_1 divided by the total equity in the system before the shock. Mathematically it is defined as

$$\mathcal{AV} = \sum_{n=1}^{N} \mathcal{SYS}(n), \tag{3}$$

where $\mathcal{SYS}(n) \in [0,\infty)$ denotes the systemicness of bank $n \in \mathcal{N}$. Hence, the *systemicness* of a bank quantifies the effect that an individual bank $n \in \mathcal{N}$ has on the aggregate vulnerability.

For each bank $n \in \mathcal{N}$, the **systemicness** $\mathcal{SYS}(n) \in [0,\infty)$ measures the contributed relative equity loss from an individual bank n (relative to the total equity of banks in the network). The systemicness of a bank $n \in \mathcal{N}$ is defined as

$$\mathcal{SYS}(n) = \gamma_{n1} \frac{\alpha_{n1}}{\sum_{v=1}^{N} e_{v1}} b_{n1}(-R_{n1}), \tag{4}$$

$$\gamma_{n1} = \sum_{k=1}^{K} \left(\sum_{p=1}^{N} \alpha_{p1} m_{pk} \right) l_k m_{nk} = \sum_{k=1}^{K} c_k l_k m_{nk}. \tag{5}$$

From this representation, one can see that systemicness is a product of four factors which all have an economic interpretation. In particular, as discussed in Greenwood et al. (2015), the systemicness of a bank is larger if: γ_{n1}, referred to as “connectedness” in Greenwood et al. (2015), is larger meaning that it is more connected in the sense that it holds assets with a large market capitalisation c_k, or whose sale has a large price impact l_k; the size measured by $\frac{\alpha_{n1}}{\sum_{v=1}^{N} e_{v1}}$ is larger since banks with larger total assets will liquidate more assets in a fire sale; the leverage b_{n1} is larger since the leverage amplifies the volume of assets sold in order to maintain the target leverage; $(-R_{n1})$ is larger, i.e., it is hit by a larger shock.

---

[2] An extension to more than two periods has been discussed in Greenwood et al. (2015) as well, but we will not consider these extensions here. Multiple rounds of deleveraging have also been considered in Cont and Schaanning (2017) and Huang et al. (2013). To account for higher order effects, fire sales have also been modelled directly as fixed point problems, see e.g., Cifuentes et al. (2005) and Amini et al. (2016).

[3] For a bank directly affected by the shock it holds that $\alpha_{n1} R_{n1} < 0$. Hence, its leverage increases from $b_{n1} = d_{n1}/e_{n1}$ to $d_{n1}/(e_{n1} + \alpha_{n1} R_{n1})$. It therefore sells $y_{n1} = -\alpha_{n1} b_{n1} R_{n1} > 0$ assets, if it has enough assets still available, i.e., if $-\alpha_{n1} b_{n1} R_{n1} < \alpha_{n1}(1 + R_{n1})$. It then pays back parts of its debt leading to a new leverage of $(d_{n1} - y_{n1})/(e_{n1} + \alpha_{n1} R_{n1}) = b_{n1}$, which is indeed the original (target) leverage. (Since a bank can never sell more assets than it has, to be precise one would need to set $y_{n1} = \left(\min\{-\alpha_{n1} b_{n1} R_{n1}, \alpha_{n1}(1 + R_{n1})\}\right)^+$. In our case studies, this cap was never reached, therefore we ignore it in the following to keep the notation simpler.

[4] This means, that bank n sells $-m_{nk} y_{n1} \ge 0$ of asset k (in million EUR). In our empirical analysis, all nodes had enough assets left to sell according to this rule.

[5] The total amount of asset k sold is $\phi_{k1} = -\sum_{n=1}^{N} m_{nk} \alpha_{n1} b_{n1} R_{n1}$ and the resulting time 2 shock to asset k is then $f_{2k} = -l_k \phi_{k1}$.

The **direct vulnerability** of bank $n \in \mathcal{N}$, denoted by $\mathcal{DV}(n) \in [0,\infty)$, is the fraction of its equity lost directly due to the initial shock F_1. It is given by

$$\mathcal{DV}(n) = \frac{\alpha_{n1}(-R_{n1})}{e_{n1}}. \tag{6}$$

The **indirect vulnerability** for a bank $n \in \mathcal{N}$, denoted by $\mathcal{IV}(n) \in [0,\infty)$, measures the fraction of its equity that is lost due to the deleveraging of the banks. It is defined as follows

$$\mathcal{IV}(n) = \frac{\alpha_{n1}}{e_{n1}} \sum_{k=1}^{K} \left[l_k m_{nk} \left((-1) \sum_{j=1}^{N} m_{jk} \alpha_{j1} b_{j1} R_{j1} \right) \right]. \tag{7}$$

The intuition behind the formula (7) is as follows. The first factor α_{n1}/e_{n1} measures the effect of the leverage of a bank. Higher leverage will result in higher indirect vulnerability. The second term of interest is $l_k m_{nk}$ and as shown in Greenwood et al. (2015) it can be interpreted as an illiquidity-weighted exposure measure to asset k. Finally, the total volume of asset k sold is $-\sum_{j=1}^{N} m_{jk}\alpha_{j1}b_{j1}R_{j1} \geq 0$. Hence, nodes that hold a large amount of illiquid assets that are sold in large quantities have a high indirect vulnerability.

Although the systemicness and the indirect vulnerability share common factors, a bank n may have a high $\mathcal{IV}(n)$ and a low $\mathcal{SYS}(n)$ or vice versa, as noted in Greenwood et al. (2015).

2.4. Observability and choice of model parameters

The financial market is characterised by a matrix of asset holdings X and the equity of the institutions e_{n1}, $n \in \mathcal{N}$. The equities e_{n1} are in principle observable from balance sheet data. The asset holding matrix is not necessarily fully observable.

Definition 2.1 *(Full and partial information)*. We refer to a stress test as being *under full information* if the asset holding matrix X is fully observed, i.e., if for all $n \in \mathcal{N}, k \in \mathcal{S}$, the individual entries X_{nk} are known.
We refer to a stress test as being *under partial information* if only the row and column sums of the asset holding matrix X given in (1) are known, but for all $n \in \mathcal{N}, k \in \mathcal{S}$, the individual entries X_{nk} are unknown.
In both situations, we assume that the equity e_{n1} is known for all $n \in \mathcal{N}$.

In practice, detailed information on individual asset holdings X_{nk} is often not available, in particular when one considers financial institutions operating in different jurisdictions. The row and column sums of the asset holding matrix are more widely available (from balance sheet and market data).

To conduct the stress testing exercise we will need to specify the shocks f_{k1}, $k \in K$ and the price impact parameters l_k, $k \in \mathcal{S}$. We will consider different choices for the shocks f_k, $k \in \mathcal{S}$.

Definition 2.2 *($\tilde{K}$-asset shock and all asset shock)*. We refer to a situation in which only $0 < \tilde{K} \leq K$ assets are shocked as a *$\tilde{K}$-asset shock*. We denote the indices of shocked assets by $\mathcal{I}^{\tilde{K}} \subseteq \mathcal{S}$. Then, $|\mathcal{I}^{\tilde{K}}| = \tilde{K}$. Furthermore, $f_i < 0$ for all $i \in \mathcal{I}^{\tilde{K}}$ and $f_i = 0$ for all $i \in \mathcal{S} \setminus \mathcal{I}^{\tilde{K}}$.

If all assets are shocked equally, i.e., $f_1 = \ldots = f_K = f < 0$, we will refer to such a shock as an *all asset shock*.

Note that according to our definition, every all asset shock is a K-asset shock, but not every K-asset shock is an all asset shock since an all asset shock has the additional feature that all assets are shocked equally.

We will distinguish between two different parametric choices for the price impact parameters l_k.

Definition 2.3 *(Constant and capitalisation-dependent price impact)*. We will refer to a price impact given by $l_1 = \ldots = l_K = l \in [0,\infty)$, as a *constant price impact*. We will refer to a price impact given by $l_k = \frac{\rho}{c_k}$ $\forall k \in \mathcal{S}$, where $\rho > 0$ and $c_k > 0$ is the capitalisation of asset k defined in (1), as the *capitalisation-dependent price impact*. If, $c_k = 0$, we set $l_k = 0$.

A constant price impact was assumed in the empirical analysis in Greenwood et al. (2015). Since it has been argued that a constant price impact can overestimate the losses for liquid assets and underestimate the losses for illiquid assets (Cont and Schaanning, 2017, p. 19), we additionally consider a capitalisation-dependent price impact which assumes that the sale of assets with larger market capitalisation leads to a smaller price impact. This implicitly assumes that assets with a larger market capitalisation are more liquid and therefore cause a smaller price impact when sold.[6] One could adjust the definition of a capitalisation dependent price impact if one wanted to allow for the existence of external investors in the model.[7]

3. Assessing fire sales risk under full and partial information

Next, we conduct stress testing to analyse fire sales risk under both full and partial information.

3.1. Dependence of fire sales risks on the asset holding matrix

We start by presenting our theoretical results on the dependence of the systemicness, the aggregate vulnerability, and the direct vulnerability on the asset holding matrix X. We assume that the row and column sums of X, given in (1), are known and we determine how much these three risk measures ($\mathcal{AV}$, $\mathcal{SYS}(n)$ and $\mathcal{DV}(n)$) depend on the individual entries of X (beyond the information contained in the row and column sums). We show that there are several situations in which some of these risk measures do not depend on the individual entries of the asset holdings matrix X at all but only on its row and column sums.

Proposition 3.1. *Let $X \in [0,\infty)^{N \times K}$ be an asset holdings matrix. Suppose that its row and column sums $\alpha_{n1}, n \in \mathcal{N}$ and $c_k, k \in \mathcal{S}$, defined in (1), are known.*

1. *for each $n \in \mathcal{N}$, the systemicness $\mathcal{SYS}(n)$ can depend on the individual entries of X only via γ_{n1} and R_{n1};*
2. *the aggregate vulnerability $\mathcal{AV}$ can depend on the individual entries of X only via γ_{n1} and R_{n1}, where $n \in \mathcal{N}$;*
3. *the direct vulnerability $\mathcal{DV}_1(n)$ of an institution n can depend on the individual entries of X only via R_{n1};*
4. *For a constant price impact, γ_{n1} can depend on the individual entries of X only via its nth row; for a capitalisation dependent price impact, γ_{n1} does not depend on X.*
5. *R_{n1} can depend on the individual entries of X only via its nth row; furthermore,*
    (a) *for an all asset shock, R_{n1} and hence $\mathcal{DV}(n)$ do not depend on the individual entries of X.*

---

[6] In our empirical analysis we will set $\rho = -\log(0.1)$. This parametric assumption is inspired by models such as Cifuentes et al. (2005), which use an exponential function to describe the inverse demand function that maps the quantities being sold to a price. In our case, we do not consider the quantities being sold, but the total market capitalisation and therefore our definition is slightly different from the classical characterisation in terms of an inverse demand function. One could also consider other inverse demand functions in this setting, see e.g., Bichuch and Feinstein (2022).

[7] These external investors would hold (some of) the assets but would not engage in the leverage targeting/fire sales mechanisms. The total market capitalisation of an asset k would then consist of $c_k + c_k^{(e)}$, where c_k is the market capitalisation of asset k among the nodes in $\mathcal{N}$ and $c_k^{(e)}$ is the market capitalisation of asset k held by external investors. Then a capitalisation-dependent price impact could be defined as $l_k^e = \frac{\rho}{c_k + c_k^{(e)}}$ for a fixed $\rho > 0$. We will not consider this generalisation in the following.

(b) *for a $\tilde{K}$-asset shock, R_{n1} and hence $\mathcal{DV}(n)$ only depend on the columns with indices in $\mathcal{I}^{\tilde{K}}$ within the nth row, but not on the full nth row of X.*

Corollary 3.2. *In addition to the assumptions of Proposition 3.1, let the price impact be capitalisation-dependent and let $n \in \mathcal{N}$. Then,*

1. *for an all asset shock, the systemicness $\mathcal{SYS}(n)$, the direct vulnerability $\mathcal{DV}(n)$ and the aggregate vulnerability $\mathcal{AV}$ do not depend on the individual entries of X.*
2. *for a $\tilde{K}$-asset shock, both the systemicness $\mathcal{SYS}(n)$ and the direct vulnerability $\mathcal{DV}(n)$ only depend on the columns with indices in $\mathcal{I}^{\tilde{K}}$ of the nth row, but not on the full nth row of X. Furthermore, the aggregate vulnerability $\mathcal{AV}$ only depends on X via its columns with indices in $\mathcal{I}^{\tilde{K}}$.*

The proofs of Proposition 3.1 and Corollary 3.2 are in Appendix B. Note that the indirect vulnerability often depends on the individual entries of X. In the following, we will focus our analysis on situations in which the risk measures do indeed depend on (parts of) the underlying matrix X and assess the effect of using different matrix reconstruction methods to estimate the individual entries of the matrix X.

Remark 3.3 *(Notation).* To make it clear which matrix is used to compute the corresponding risk measures, we will sometimes write $\mathcal{SYS}^X(n)$, $\mathcal{AV}^X$, $\mathcal{DV}^X(n)$, $\mathcal{IV}^X(n)$ to show explicitly that these risk measures are computed based on the matrix X. In our analysis we will later allow the matrix X to be random, i.e., each element of X is a random variable, in which case then the corresponding systemic risk measures also become random variables.
We will also later use the notation $\gamma_{n1} = \gamma_{n1}(X)$ and $R_{n1} = R_{n1}(X)$, where $n \in \mathcal{N}$, to indicate that both γ_{n1} and R_{n1} can depend on X.

We have seen that the asset holding matrix X enters the different risk measures only via R_{n1} and γ_{n1}, where $n \in \mathcal{N}$. It plays a different role in these two quantities.

In R_{n1}, the asset holding matrix directly influences the magnitude of the original shock. In particular, the loss of bank n's equity after the initial shock is given by

$$\alpha_{n1} R_{n1} = \alpha_{n1} \sum_{k=1}^{K} m_{nk} f_{k1} = \alpha_{n1} \sum_{k=1}^{K} \frac{X_{nk}}{\alpha_{n1}} f_{k1} = \sum_{k=1}^{K} X_{nk} f_{k1} \leq 0. \tag{8}$$

For an all asset shock, equation (8) simplifies further to $\alpha_{n1} R_{n1} = f \sum_{k=1}^{K} X_{nk} = f \alpha_{n1}$. Hence, for an all asset shock, one only needs to know the row sum α_{n1} rather than the individual entries of X to determine the magnitude of the shock.

In γ_{n1}, the asset holding matrix enters through the rule of how assets are sold following a stress. Greenwood et al. (2015) assume that stressed banks sell assets according to the proportion of their original portfolio positions, i.e., if bank n sells a total of y_{n1} assets (in EUR), it sells $m_{nk} y_{n1} = \frac{X_{nk}}{\alpha_{n1}} y_{n1}$ of asset k (in EUR).

One can generalise the selling rule by assuming that banks no longer sell according to the matrix m but according to a matrix $\mu \in [0,1]^{N\times K}$, where $\sum_{k=1}^{K} \mu_{nk} = 1$ for all $n \in \mathcal{N}$.[8] Then, for a general selling rule characterised by the matrix μ we use

$$\gamma_{n1}^{(\mu)} = \sum_{k=1}^{K} c_k l_k \mu_{nk},$$

rather than $\gamma_{n1} = \sum_{k=1}^{K} c_k l_k \frac{X_{nk}}{\alpha_{n1}}$ in the formulae for the different fire sale risk measures. Such a formulation allows us to capture situations in which the actual selling rule that banks use under stress is unknown, which would often be the case in practice.

We can then compute bounds on the influence of the selling rule by considering for each $n \in \mathcal{N}$

$$\max_{\mu\in[0,1]^{N\times K}} \gamma_{n1}^{(\mu)} \qquad \text{subject to } \sum_{k=1}^{K} \mu_{nk} = 1, \tag{10}$$

or the corresponding minimisation problem. These linear optimisation problems can be solved analytically. The maximum of the objective function is $\max_{k\in\mathcal{S}}\{c_k l_k\}$ which is attained by setting $\mu_{nk} = 1$ at the index k where the maximum $\max_{k\in\mathcal{S}}\{c_k l_k\}$ is attained and $\mu_{nk} = 0$ for all remaining indices k. In particular, the optimal solution does not depend on n. The corresponding result in which max is replaced by min holds for the minimum. The optimal strategy corresponding to the maximisation problem selects the asset with the highest capitalisation-weighted price impact.[9]

These results hold for general price impact parameters $l_1, \ldots, l_K$. For a capitalisation-dependent price impact, however, the connectedness simplifies to

$$\gamma_{n1}^{(\mu)} = \sum_{k=1}^{K} c_k l_k \mu_{nk} = \sum_{k=1}^{K} c_k \frac{\rho}{c_k} \mu_{nk} = \rho$$

and hence it does not depend on the selling rule μ.

Furthermore, we find that if $\gamma_{n1}^{(\mu)} = \tilde{\gamma}$ for all $n \in \mathcal{N}$, i.e., if all banks n have the same connectedness, then the network effect arising from $\gamma_{n1}^{(\mu)}$ only becomes a scaling factor in the aggregate vulnerability, in particular,

$$\mathcal{AV} = \sum_{n=1}^{N} \gamma_{n1} \frac{\alpha_{n1}}{\sum_{v=1}^{N} e_{v1}} b_{n1}(-R_{n1}) = \frac{-\tilde{\gamma}}{\sum_{v=1}^{N} e_{v1}} \sum_{n=1}^{N} b_{n1} \sum_{k=1}^{K} X_{nk} f_{k1}. \tag{11}$$

This situation arises, as discussed for a capitalisation-dependent price impact. It also arises for all selling strategies μ that are not bank specific, i.e., for which $\mu_{nk} = \tilde{\mu}_k$ for all $n \in \mathcal{N}$ and for $\tilde{\mu}_1, \ldots, \tilde{\mu}_K \in [0,1]$ with $\sum_{k=1}^{K} \tilde{\mu}_k = 1$.[10]

These considerations show that in order to have a more involved interaction between the two network effects $\gamma_{n1}^{(\mu)}$ and R_{n1}, one needs to consider shocks that are not an all asset shock, selling strategies that vary between banks (as e.g., the strategy considered in Greenwood et al. (2015)) and a price impact that is not capitalisation-dependent.

In the following, we will therefore focus on the selling strategy assumed by Greenwood et al. (2015), where $\mu = m$ and we will not consider other strategies any further. To analyse the different effects of R_{n1} and γ_{n1}, we will include an all asset shock in our analysis, to isolate the effect of γ_{n1} for fixed (meaning that they do not depend on the individual entries of the asset holding matrix) R_{n1}. Similarly, we will also include a capitalisation-dependent price impact to isolate the effect of R_{n1} for fixed γ_{n1}. Furthermore, we will consider shocks that affect only

---

[8] The general selling rule does not exclude short-selling. If one wanted to exclude short-selling for the general selling rule one would need to require that $\mu_{nk} y_{n1} \leq X_{nk}(1 + f_{k1})$, where the left-hand side is the total amount of asset k sold by bank n, and the right-hand side is the amount of asset k that bank n has after the shock. This implies the following additional condition on μ, namely

$$\mu_{nk} \leq \frac{X_{nk}(1 + f_{k1})}{y_{n1}} \quad \forall k \in \mathcal{S}, \tag{9}$$

for all $n \in \mathcal{N}$ with $y_{n1} > 0$. (We set $\mu_{nk} = 0$ for all $n \in \mathcal{N}$ with $y_{n1} = 0$ and $\forall k \in \mathcal{S}$.)

[9] The solution to the optimisation problem (10) and the corresponding minimisation problem are useful as upper and lower bounds on the potential influence of the selling rule. It is possible, that banks do not hold the amount of assets that needs to be sold according to these optimal strategies. To exclude short-selling one would need to include the additional condition (9) in the optimisation problems.

[10] An example of such a selling strategy would be to sell according to equal proportions $\mu_{nk} = 1/K$ for all $n \in \mathcal{N}$ and for all $k \in \mathcal{S}$.

some assets and use a constant-price impact to study the interaction between γ_{n1} and R_{n1}.

3.2. Reconstructing matrices

We consider five existing methods for reconstructing the asset holding matrix X from the given row and column sums. We briefly summarise them below. More details can be found in Appendix A.

We consider two optimisation-based methods: the Entropy method by Upper and Worms (2004) and the minimum density method by Anand et al. (2015) (MinDen). The optimisation-based matrix reconstruction methods consist of a suitably chosen objective function that is optimised over the set of matrices that satisfies the given constraints on the row and column sums. The result of the reconstruction problem is one matrix that satisfies the constraints. Other possible characteristics of this matrix depend on the chosen objective function. For the Entropy method, the resulting matrix is usually complete, i.e., all entries are non-zero as long as all the row and column sums are non-zero. This means that the banks then have a fully diversified portfolio since they hold positions in each asset. For the MinDen method, the resulting matrix is usually very sparse, i.e., most of the entries are equal to zero. This means that the banks have more diverse positions.[11] For the Entropy method, the matrix that solves the corresponding optimisation problem is available analytically. Therefore, it is possible to characterise all fire sales measures, i.e., the direct and indirect vulnerability, the systemicness, and the aggregate vulnerability, derived from the reconstructed matrix using the Entropy method analytically. We provide the corresponding formulae in Proposition A.1.

We also consider three probabilistic methods: the statistical physics method by Cimini et al. (2015) (StatPhys) (and extended to bipartite networks by Squartini et al. (2017)) and the Bayesian approach by Gandy and Veraart (2017), where we assume two different priors within the Bayesian framework, an Erdős-Rényi-type prior (BayeER) and an empirical fitness type prior (BayeEF) as in Gandy and Veraart (2019). All probabilistic models assume that the matrix of asset holdings is random, i.e., all its elements are random variables. They provide methodologies to generate samples from the distribution of this random asset holding matrix. Therefore, the result of a network reconstruction method using any of the probabilistic methods is a sample of matrices and not just one matrix. All three probabilistic models are calibrated to match the (true) density of the network. For all three probabilistic methods, our analysis uses a sample size of 10,000.[12] The StatPhys method is related to the Entropy method and it is possible to compute the expectation of the fire sale measures applied to the random matrix that is used in the StatPhys method analytically. We provided the details in Appendix A.2.1.

We will now illustrate how the choice of the reconstruction methods affects the risk measures to assess fire sales risk. First, we consider a toy example.

Example 3.4 *(Toy example: assessing fire sale losses on reconstructed matrices)*. We consider the asset holdings matrix $X \in [0,\infty)^{3\times 3}$ reported in Table 2, and two reconstructions of X from its row and column sums using the Entropy method and the MinDen method, respectively. According to the true matrix X each institution holds two assets and each asset is held by two institutions. The Entropy method distributes the weights evenly across the different cells of the matrix resulting in a complete network, meaning all institutions hold all assets, whereas the MinDen method finds the sparsest possible solution in which each institution only holds one asset and this asset is not held by anyone else.

The total assets are $\alpha_{n1} = 4$ for all $n \in \{1,2,3\}$. We assume that all three banks have the same equity, namely $e_{n1} = 1$ which results in a leverage of $b_{n1} = 3$ for all $n \in \{1,2,3\}$. We consider a constant price impact of $l_k = 10^{-2}$ for all assets $k \in \{1,2,3\}$. We consider a 1-asset shock given by $F_1 = (-0.15, 0, 0)^\top$ which only affects the first asset directly.

We report the $\mathcal{AV}$, $\mathcal{SYS}(n)$, $\mathcal{DV}(n)$ and $\mathcal{IV}(n)$ for the true and the two reconstructed matrices for all three institutions $n \in \{1,2,3\}$ in Table 2. We find that under this 1-asset shock, the systemicness, the direct and the indirect vulnerability can change significantly with respect to the network that is used as input.

In particular, we find that the Entropy method underestimates the systemicness, the direct vulnerability, and the indirect vulnerability of banks 1 and 2 (which have the highest systemicness and direct and indirect vulnerability in this example), and overestimates the systemicness, the direct and indirect vulnerability of bank 3 (which is the lowest among all banks in this example).

The MinDen method only attributes a positive systemicness, and direct and indirect vulnerability to bank 1 (and significantly overestimates the true values), and otherwise provides estimates of zero for all three measures for banks 2 and 3.

For the aggregate vulnerability, however, we see that it is correctly estimated by both the Entropy and the MinDen method in this example.

3.3. Data

In our empirical analysis, we consider data[13] collected by the European Banking Authority (EBA)[14] for their stress tests of EU banks in 2011 and 2016. The data consists of balance sheets of some of the largest banks in the EU. The data include $N = 90$ banks in 2011 and $N = 51$ banks in 2016 covering the EU countries (which includes the UK in these years).

We aggregate the asset classes such that all asset classes are consistent across both years. There are $K = 36$ asset classes which include corporate, retail, 30 EEA sovereign loans, US, Japan, Latin America, and other sovereign loans (an aggregated class of remaining sovereign loans). Hence the asset holding matrix is a 90×36 matrix in 2011 and a 51×36 matrix in 2016. All other assets which are not recorded in both years are not included in the asset holdings matrix.[15] We assume that all assets are marketable and can be liquidated, i.e., we apply the framework by Greenwood et al. (2015) directly to the full asset holding matrix as in the empirical case study provided in Greenwood et al. (2015).[16]

Table 3 provides some descriptive statistics for the EBA data used in our empirical analysis. We see that the network densities (defined as $\frac{1}{NK}\sum_{n=1}^{N}\sum_{k=1}^{K}\mathbb{1}_{\{X_{nk}^{\text{year}}>0\}}$, where X^{year} represents either the observed

---

[11] For further discussion and results on the relationship between fire sales risk and diversification versus diversity in asset portfolios we refer to Capponi and Weber (2021) and Detering et al. (2022).

[12] For the Bayesian approach (in which the distribution of interest is approximated using a Gibbs sampler) we choose thinning and burn-in parameters as 10% of the total number of samples as in Gandy and Veraart (2017). We use the R-package *systemicrisk* available at https://CRAN.R-project.org/package=systemicrisk that implements the methods by Gandy and Veraart (2017, 2019).

[13] The data are publicly available from https://eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing/2011 and https://eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing/2016.

[14] The EBA is an independent EU authority whose objective is to maintain financial stability in the EU. It has been established to develop consistent prudential regulation and supervision of the EU's banking sector. Together with the European Systemic Risk Board (ESRB), the EBA conducts stress testing of the EU banking sector to assess its resilience to adverse shocks.

[15] This is done for consistency purposes so that we can apply the same initial shock to both datasets.

[16] In practice, only a part of a bank's assets can be liquidated, see e.g., Cont and Schaanning (2017) for discussion and more details on this. It would be possible to restrict the modelling framework so that only a subset of the available assets are marketable. Related work that considers both marketable and non-marketable assets includes Braouezec and Wagalath (2019), Feinstein (2020), Banerjee and Feinstein (2021).

Table 2

The three matrices on the left in the first row represent the true asset holding matrix X and the two reconstructed matrices using the Entropy and MinDen methods, respectively. The table on the right in the first row shows the systemicness $\mathcal{SYS}(n)$ and $\mathcal{AV}$, and the tables in the second row show the direct vulnerability $\mathcal{DV}(n)$ and the indirect vulnerability $\mathcal{IV}(n)$ for each bank n corresponding to the shock F_1, specified in Example 3.4, and applied to the true and the two reconstructed matrices.

$$\begin{array}{ccc} \text{True } X & \text{Entropy} & \text{MinDen} \\ \begin{pmatrix} 2 & 2 & 0 \\ 2 & 0 & 2 \\ 0 & 2 & 2 \end{pmatrix} & \begin{pmatrix} \frac{4}{3} & \frac{4}{3} & \frac{4}{3} \\ \frac{4}{3} & \frac{4}{3} & \frac{4}{3} \\ \frac{4}{3} & \frac{4}{3} & \frac{4}{3} \end{pmatrix} & \begin{pmatrix} 4 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 4 \end{pmatrix} \end{array}$$

	True X	Entropy	MinDen
$\mathcal{SYS}(1)$	1.2%	0.8%	2.4%
$\mathcal{SYS}(2)$	1.2%	0.8%	0.0%
$\mathcal{SYS}(3)$	0.0%	0.8%	0.0%
$\mathcal{AV}$	2.4%	2.4%	2.4%

	True X	Entropy	MinDen		True X	Entropy	MinDen
$\mathcal{DV}(1)$	30.0%	20.0%	60.0%	$\mathcal{IV}(1)$	2.7%	2.4%	7.2%
$\mathcal{DV}(2)$	30.0%	20.0%	0.0%	$\mathcal{IV}(2)$	2.7%	2.4%	0.0%
$\mathcal{DV}(3)$	0.0%	20.0%	0.0%	$\mathcal{IV}(3)$	1.8%	2.4%	0.0%

Table 3

Summary statistics for the EBA data.

Year	Number of banks	Network density	Leverage			Total assets (EUR)		
			Min	Mean	Max	Min	Mean	Max
	-	-	-	-	-	(bn)	(tn)	(tn)
2011	90	0.44	3.56	33.60	540.68	0.33	0.19	1.21
2016	51	0.48	7.48	19.68	43.15	2.93	0.28	1.39

assets holding matrix in 2011 or in 2016) are almost the same in both years (0.44 in 2011 and 0.48 in 2016). With almost half of the entries being positive in the asset holding matrix, we expect that there is indeed scope for serious contagion effects if fire sales are triggered.

The leverage values that we report in this table correspond to $b_{n1} = d_{n1}/e_{n1}$, where the equity e_{n1} was set to be equal to the common equity Tier 1 capital reported for each bank $n \in \mathcal{N}$ and the debt d_{n1} was set to be equal to $\alpha_{n1} - e_{n1}$. We see that leverages were generally significantly higher in 2011 than in 2016. In the empirical analysis we will cap the leverage at 30 as in Greenwood et al. (2015), this means that we set $b_{n1} = \min\{30, \text{observed leverage of bank } n \text{ at time } 1\}$ to avoid having to deal with a small number of banks which have very high leverages (such as, e.g., maximum leverage of 540.68 in 2011).

Table 3 also shows the range of banks' total assets α_{n1}, $n \in \mathcal{N}$, with the lowest total assets at 329 million EUR in 2011, compared with the largest bank at 1.39tn EUR in 2016. The assets which form the largest total capitalisation in both networks are corporate and retail, comprising 82% of the total assets in 2011 and 80% in 2016. The highest sovereign loans in 2011 are German (2.89%), other sovereign (1.98%), and Italian (1.93%) compared with French (2.68%), German (2.64%), and US (2.13%) sovereign assets in 2016. Asset classes with a large capitalisation which are held by a large number of banks are German, French, Spanish, UK, and Italian sovereign assets.

We provide heatmaps and further discussions on the structure of the empirically observed asset holding matrix and the performance of its reconstruction using different methods for the 2016 EBA data in Appendix C.1.

3.4. Empirical results

Next, we consider the fire sales risk measures for the EBA data for three different stress scenarios and compare the results obtained by using the fully observed matrix of asset holdings to the results derived based on reconstructed asset holding matrices.

The three stress scenarios are as follows:

GIIPS shock: This is a 5-asset shock. We consider a 5% shock to the sovereign loans of Greece, Italy, Ireland, Portugal and Spain (GIIPS), which were countries highly impacted by the 2008 financial crisis. This shock has also been considered in Greenwood et al. (2015). Mathematically, this corresponds to setting $f_k = -0.05$ for each index $k \in \mathcal{S}$ that corresponds to a GIIPS asset and setting $f_k = 0$ for all remaining $k \in \mathcal{S}$ in the initial shock vector.

Bad Brexit shock: This is a 1-asset shock. We consider an economic shock of 10% to UK sovereign loans as a possible scenario for negative consequences arising from Brexit. Mathematically this is captured by setting $f_k = -0.1$ for the index $k \in \mathcal{S}$ that corresponds to the UK asset class and setting $f_k = 0$ for all remaining $|K| - 1$ assets.

All asset shock: We consider a 0.1% shock to all assets. This corresponds to setting $f_k = -0.001$ for all $k \in \mathcal{S}$ in the initial shock vector. This type of shock is widespread and affects all assets in the same way.

For the optimisation-based reconstruction methods, i.e., for MinDen and Entropy, we can apply the different fire sales measures directly to the reconstructed matrix returned by these methods since we only evaluate one matrix. As mentioned before, for the Entropy method we have analytical expressions for all fire sale measures, see Proposition A.1. Since, there are three bank specific measures, namely the systemicness $\mathcal{SYS}(n)$, the direct vulnerability $\mathcal{DV}(n)$, and the indirect vulnerability $\mathcal{IV}(n)$, we consider the average of these measures across all institutions. For example, instead of N measures for the direct vulnerability of the individual institutions, we report the average direct vulnerability over all institutions given by $\frac{1}{N}\sum_{n=1}^{N}\mathcal{DV}(n)$. Table 4 reports the results and $\mathcal{DV}11$ then corresponds to the average direct vulnerability across the N banks in 2011 and $\mathcal{DV}16$ represents the average direct vulnerability in 2016. The same notation is used for the indirect vulnerability. Since the aggregate vulnerability is a measure for the whole system we can report it directly. Note that the aggregate vulnerability is equal to N times the average systemicness and therefore we do not report the average systemicness.[17]

---

[17] We also considered other measures for bank specific quantities. We investigated the L_1-error between the fire sales measures of reconstructed matrices and the true matrix, e.g., for systemicness we consider $\sum_{n=1}^{N}|\mathcal{SYS}^{\text{True}}(n) - \mathcal{SYS}^{X}(n)|$ for the Entropy and the MinDen method, and for the sampling methods, we average the L_1-error across all samples, i.e., we consider $\frac{1}{d}\sum_{v=1}^{d}\sum_{n=1}^{N}|\mathcal{SYS}^{\text{True}}(n) - \mathcal{SYS}^{X^{(v)}}(n)|$. For the three different bank-specific fire sale measures we find a larger deviation for systemicness and indirect vulnerability compared to the direct vulnerability. Generally, the MinDen reconstruc-

For the probabilistic reconstruction measures, i.e., for StatPhys, BayeER, and BayeEF, we obtain a sample of networks denoted by $X^{(1)}, \dots, X^{(d)}$, i.e., the sample size is $d = 10{,}000$. For each reconstructed network $X^{(v)}$, we compute the direct vulnerability of each bank and consider the mean direct vulnerability of all banks for this given network $X^{(v)}$. We then average this mean direct vulnerability over the full sample. More precisely, in the two rows corresponding to $\mathcal{DV}11$, we show the following two numbers (sample mean of the average direct vulnerability and corresponding standard deviation in *italic*):

$$\bar{\mu}_d^{DV} = \mu_d\left(\frac{1}{N}\sum_{n=1}^{N}\mathcal{DV}11^{X^{(v)}}(n)\right) = \frac{1}{d}\sum_{v=1}^{d}\left(\frac{1}{N}\sum_{n=1}^{N}\mathcal{DV}11^{X^{(v)}}(n)\right),$$

$$\sigma_d\left(\frac{1}{N}\sum_{n=1}^{N}\mathcal{DV}11^{X^{(v)}}(n)\right) = \sqrt{\frac{1}{d-1}\sum_{v=1}^{d}\left(\frac{1}{N}\sum_{n=1}^{N}\mathcal{DV}11^{X^{(v)}}(n) - \bar{\mu}_d^{DV}\right)^2}$$

for the probabilistic methods using the 2011 data. The same methodology is applied to the reporting of indirect vulnerability. The aggregate vulnerability in the table corresponds to the average aggregate vulnerability computed over the d elements of the sample, i.e.,

$$\bar{\mu}_d^{AV} = \mu_d\left(\mathcal{AV}11^{X^{(v)}}\right) = \frac{1}{d}\sum_{v=1}^{d}\mathcal{AV}11^{X^{(v)}},$$

$$\sigma_d\left(\mathcal{AV}11^{X^{(v)}}\right) = \sqrt{\frac{1}{d-1}\sum_{v=1}^{d}\left(\mathcal{AV}11^{X^{(v)}} - \bar{\mu}_d^{AV}\right)^2},$$

where again the second quantity is the corresponding standard deviation.

For the StatPhys method, we have derived analytical expressions for the expected fire sale measures in Appendix A.2.1. In the following, we report the Monte Carlo estimates such as $\bar{\mu}_d^{AV}$ of these expectations which are very close to the analytical results.

Table 4 reports the averaged fire sale risk measures. Values highlighted in bold indicate the best-performing reconstruction method per row and per price impact. The direct vulnerabilities are reported only for the capitalisation-dependent price impact since they coincide with those for the constant price impact, see (6). For the all asset shock, several risk measures do not depend on the individual entries of X (see Corollary 3.2) which we indicate by writing *true* in the corresponding entry in the table to highlight, that this value is identical to the value in the column *True*.

We first consider the fire sales measures based on the fully observed matrix (the columns named *True* in Table 4). We observe that all but one fire sales measure corresponding to the 2011 data are consistently higher than the corresponding measures for the 2016 data. This means that overall the fire sales risk has decreased from 2011 to 2016. This observation holds true under both constant and capitalisation-dependent price impact. The only exception, where we observe an increase from 2011 to 2016, is the aggregate vulnerability under a constant price impact for the Bad Brexit shock. A high contributing factor for the general tendency for the fire sales risk measures to decrease from 2011 to 2016 is the difference between the leverages and the target leverages (which is observed leverage capped at 30) which both decrease substantially from 2011 to 2016. In 2011, several banks held low levels of equity compared with 2016 which resulted in high leverages in 2011. The average leverage of banks also decreased from 2011 to 2016, because of the banks included in each dataset. The higher leverage in 2011 was largely driven by a small number of banks, which were not all part of the 2016 data. For example, Greek banks were not part of the 2016 data but included a bank with the second-highest leverage. Therefore, the change in capital requirements and the types of banks included in the EBA stress tests contributed to a decreased leverage.

Furthermore, we consider the effect of the two different choices for the price impact on the fire-sales measures under complete information. First, note that the direct vulnerability $\mathcal{DV}(n)$ does not depend on the price impact, i.e., it does not depend on the parameters $l_1, \dots, l_K$, and hence the direct vulnerabilities corresponding to different price impacts coincide. In contrast, the other fire-sales measures do depend on the price impact. In the 2011 data, we find that the constant price impact results in higher risks associated with fire sales than the capitalisation-dependent price impact, but for the 2016 data this is not necessarily the case. Overall, the key features of the stress tests remain consistent for both choices of price impact, namely that risks from fire sales in 2016 were smaller than in 2011.

Next, we consider the fire sales measures obtained by using five different matrix reconstruction methods. The entry shown in bold represents the best-performing method for this particular row indicating the fire sales measure and a given price impact. Overall, we see there is no clear winner in the sense that one of the methods would consistently outperform all other methods across all fire sales measures and for different types of shocks and price impacts.

So, we will look more specifically at the different fire sales measures. To get an overview, we compare the performance of three model classes: the MinDen, Entropy & StatPhys, BayeER & BayeEF. We provide more details in Appendix A.2.1 on how the Entropy and the StatPhys are indeed related. It is no coincidence that the estimates derived from using the Entropy or the StatPhys methods are very similar. We start with the aggregate vulnerability since this is the only measure that provides a holistic view of the whole network. We find that the Entropy, StatPhys, BayeER, and BayeEF methods provide estimates for the aggregate vulnerability that is often rather close to the true value. Indeed, out of the 10 cases (corresponding to 4 cases for the capitalisation-dependent price impact and 6 cases for the constant price impact) for which we compute an aggregate vulnerability the Entropy & StatPhys method is the best-performing method in 5 cases and the BayeER & BayeEF are the best-performing methods in the remaining 5 cases. The MinDen performs best in only one scenario for estimating aggregate vulnerability. In most cases, it overestimates or underestimates the aggregate vulnerability.

For the bank specific measures that we have just averaged over all banks in the network, i.e., direct and indirect vulnerability we find the following. For the indirect vulnerability, we have 12 different cases (3 stress scenarios × 2 price impacts × 2 years). The MinDen performed based in 3 cases, the Entropy & StatPhys method performed best in 5 cases and the BayeER & BayeEF performed best in 4 cases. For the direct vulnerability, we observed 4 cases (2 stress scenarios × 2 years; note that the price impact does not affect the direct vulnerability). The MinDen performed best in 3 cases and BayeER & BayeEF performed best in 1 case.[18]

Hence, to summarise, the best method for estimating the direct vulnerability across all types of shocks considered here is the MinDen, for the indirect vulnerability the performance of the three classes of methods are very similar, and for estimating the aggregate vulnerability the best methods are the Entropy & StatPhys method and the BayeER & BayeEF.

When distinguished by the type of shock, we find that the Entropy, StatPhys, BayeER, and BayeEF methods tend to be the preferred methods for the GIIPS shock (which affects 5 columns of the asset holding matrix), whereas the MinDen method seems to be the preferred method for the Bad Brexit shock which only affects one column of the asset

---

tion method results in the largest L_1 error due to the sparsity of the reconstructed matrix and the concentration of asset losses from a few banks. Entropy and StatPhys have the greatest similarity, and the performance of the Bayesian approach depends on the measures considered. Overall, the results are similar to those reported in Table 4.

[18] The fact that the direct vulnerabilities for the Entropy and the StatPhys methods coincide in expectation, is not a coincidence, but it follows from the theoretical results derived in Appendix A.2.1.

Table 4
The table presents average fire sales risk measures (averaged over the banks and additionally averaged over the samples) for the 2011 and 2016 EBA data for three different shock scenarios for the true and reconstructed matrices. All numbers are given in percent.

	Capitalisation-dependent Price Impact ($l_k = \rho / c_k \ \forall k$)						Constant Price Impact ($l_k = 5 \times 10^{-13} \ \forall k$)					
	True	MinDen	Entropy	StatPhys	BayeER	BayeEF	True	MinDen	Entropy	StatPhys	BayeER	BayeEF
GIIPS (%)												
$\mathcal{DV}11$	15.58	3.23	7.81	7.81	20.21	**17.68**	same results as for capitalisation-dependent price impact					
	-	-	-	*(0.53)*	*(3.06)*	*(3.07)*						
$\mathcal{IV}11$	460.79	296.96	416.80	**417.58**	612.49	544.98	506.76	304.77	**523.63**	523.68	274.71	325.00
	-	-	-	*(10.67)*	*(38.58)*	*(30.92)*	-	-	-	*(13.36)*	*(16.91)*	*(18.00)*
$\mathcal{AV}11$	291.70	238.77	288.49	288.96	**292.38**	294.24	357.49	178.63	362.43	**362.43**	275.36	293.40
	-	-	-	*(6.15)*	*(5.26)*	*(5.12)*	-	-	-	*(7.41)*	*(11.41)*	*(10.23)*
$\mathcal{DV}16$	5.64	**5.33**	4.42	4.42	8.24	7.57	same results as for capitalisation-dependent price impact					
	-	-	-	*(0.18)*	*(0.78)*	*(0.82)*						
$\mathcal{IV}16$	187.19	254.15	**214.86**	215.22	240.26	228.39	151.35	51.69	221.46	221.47	139.97	**153.33**
	-	-	-	*(5.87)*	*(10.34)*	*(8.73)*	-	-	-	*(5.90)*	*(7.91)*	*(7.99)*
$\mathcal{AV}16$	189.29	244.47	220.83	221.20	**209.70**	210.34	174.81	52.08	227.62	227.63	**179.05**	186.32
	-	-	-	*(5.79)*	*(5.08)*	*(5.05)*	-	-	-	*(5.75)*	*(9.37)*	*(8.69)*
Bad Brexit (%)												
$\mathcal{DV}11$	1.47	**0.85**	3.01	3.01	8.39	6.46	same results as for capitalisation-dependent price impact					
	-	-	-	*(0.47)*	*(2.55)*	*(2.43)*						
$\mathcal{IV}11$	120.19	**140.82**	160.85	161.21	240.11	207.43	155.05	183.76	202.08	202.17	103.70	**126.37**
	-	-	-	*(7.25)*	*(23.28)*	*(17.30)*	-	-	-	*(9.13)*	*(8.82)*	*(9.63)*
$\mathcal{AV}11$	90.23	144.07	**111.34**	111.55	112.09	113.62	109.02	135.01	139.87	139.92	**104.02**	114.58
	-	-	-	*(4.74)*	*(4.03)*	*(3.87)*	-	-	-	*(5.92)*	*(8.02)*	*(7.40)*
$\mathcal{DV}16$	1.59	**1.92**	2.87	2.87	5.50	4.82	same results as for capitalisation-dependent price impact					
	-	-	-	*(0.21)*	*(0.90)*	*(0.90)*						
$\mathcal{IV}16$	130.07	**126.60**	139.24	139.53	156.38	147.43	136.35	105.51	**143.53**	143.59	90.33	99.77
	-	-	-	*(5.92)*	*(10.07)*	*(08.30)*	-	-	-	*(6.08)*	*(7.25)*	*(7.46)*
$\mathcal{AV}16$	149.58	**152.24**	143.11	143.41	135.77	136.34	159.82	106.67	147.51	**147.58**	115.61	121.39
	-	-	-	*(6.00)*	*(5.24)*	*(5.31)*	-	-	-	*(6.14)*	*(9.27)*	*(8.87)*
All Asset (%)												
$\mathcal{DV}11$	3.46	true	true	3.46	true	true	same results as for capitalisation-dependent price impact					
	-	-	-	*(0.05)*	-	-						
$\mathcal{IV}11$	185.51	185.95	184.72	185.07	**185.60**	185.89	228.63	255.63	**232.07**	232.38	153.11	172.48
	-	-	-	*(3.24)*	*(0.68)*	*(0.60)*	-	-	-	*(4.32)*	*(7.81)*	*(7.99)*
$\mathcal{AV}11$	127.86	true	true	128.07	true	true	160.87	161.00	160.63	**160.82**	160.24	159.99
	-	-	-	*(1.32)*	-	-	-	-	-	*(1.68)*	*(0.30)*	*(0.30)*
$\mathcal{DV}16$	2.07	true	true	2.07	true	true	same results as for capitalisation-dependent price impact					
	-	-	-	*(0.02)*	-	-						
$\mathcal{IV}16$	100.62	99.58	100.47	**100.64**	99.33	99.66	92.06	103.63	103.56	103.66	82.60	**87.37**
	-	-	-	*(1.45)*	*(0.23)*	*(0.20)*	-	-	-	*(1.54)*	*(1.45)*	*(1.69)*
$\mathcal{AV}16$	103.26	true	true	103.43	true	true	106.59	105.72	106.43	**106.55**	107.69	107.59
	-	-	-	*(1.28)*	-	-	-	-	-	*(1.34)*	*(0.25)*	*(0.25)*
Bold	-	6	2	2	3	1	-	0	3	4	2	3

holding matrix. For the All Asset shock, we know from Corollary 3.2 that under a capitalisation-dependent price impact $\mathcal{SYS}(n)$, $\mathcal{AV}$, $\mathcal{DV}(n)$ do not depend on the individual entries of X. Hence, if one was interested in such a situation, there is no need to reconstruct the network. If we ignore this result and compute the corresponding risk measures from the reconstructed networks, then we indeed recover the true values exactly (indicated by *true* in the entry in the table) for all methods except for the StatPhys method. The matrices reconstructed using the StatPhys method do not individually satisfy the constraints on the row and column sums but only in expectation and therefore they do not reproduce the true value when plugged into the general formula for the risk measures that have not been simplified in line with the results of Corollary 3.2. To indicate this effect, we have reported the values computed from the matrices returned by the StatPhys method in the table. As discussed this is no contradiction to Corollary 3.2.

Finally, we investigate what the best and worst aggregate vulnerabilities are that are consistent with the given row and column sums of the asset holding matrix. Hence, we consider an additional (optimisation-based) network reconstruction method that maximises (or minimises) the aggregate vulnerability over all (non-negative) matrices that satisfy the row and column constraints. We provide more details on this in Appendix C.2. Table 5 shows the results.

We find that in general there is quite a large difference between the minimum and the maximum aggregate vulnerabilities. Furthermore, the true aggregate vulnerability, i.e., the aggregate vulnerability derived from the true network, is quite centred between the minimum and the maximum aggregate vulnerability. Given this wide range of possible aggregate vulnerabilities, the aggregate vulnerabilities obtained from the different reconstruction methods are remarkably close to the true aggregate vulnerabilities.

Furthermore, we find that the difference between the minimum and maximum aggregate vulnerabilities is generally larger for a constant price impact compared to a capitalisation-dependent price impact. This is in line with our theoretical results (Proposition 3.1 and Corollary 3.2). Also in line with these theoretical results, is the fact that the GIIPS shock scenario in which 5 assets are shocked, has a larger difference between the minimum and maximum aggregate vulnerability than the Bad Brexit shock in which only one asset is shocked.

As discussed before, for an all asset shock and a capitalisation-dependent price impact, the aggregate vulnerability only depends on the asset holding matrix via its row and column sums and therefore

Table 5
True aggregate vulnerability and minimum and maximum of aggregate vulnerabilities derived from asset holding matrices that satisfy the given row and column sums. Two different price impacts and three shock scenarios are considered.

	Capitalisation-dependent ($l_k = \rho / c_k\ \forall k$)			Constant ($l_k = 5 \times 10^{-13}\ \forall k$)		
GIIPS (%)						
	True	Min	Max	True	Min	Max
$\mathcal{AV}11$	291.70	172.18	373.30	357.49	6.26	494.91
$\mathcal{AV}16$	189.29	131.88	294.12	174.81	4.95	311.53
Bad Brexit (%)						
	True	Min	Max	True	Min	Max
$\mathcal{AV}11$	90.23	49.16	144.07	109.02	1.00	216.65
$\mathcal{AV}16$	149.58	60.08	195.02	159.82	1.59	220.78
All Asset (%)						
	True	Min	Max	True	Min	Max
$\mathcal{AV}11$	127.86	true	true	160.87	150.19	171.73
$\mathcal{AV}16$	103.26	true	true	106.59	99.93	113.69

the minimum and maximum aggregate vulnerabilities coincide with the true aggregate vulnerability. For an all asset shock with constant price impact, the individual entries of the asset holding matrix do matter, but we find the range of aggregate vulnerabilities to be rather small.

3.5. Sensitivity analysis and robustness checks

We have seen so far that the sampling-based reconstruction methods (StatPhys, BayeER and BayeEF) provide superior results for several measures of fire sale risk. One might think that this result is purely driven by the fact that these sampling-based methods can be (and in Table 4 have been) calibrated to the true density of the network which was not the case for the MinDen and the Entropy method. In the following, we show that this result remains robust even if the underlying assumption on the network density is changed.

Fig. 1 shows how the aggregate vulnerability computed using the StatPhys, BayeER and BayeEF network reconstruction methods depends on the choice of the target density of the network. It shows the aggregate vulnerabilities as a function of the network density for three sampling-based reconstruction methods (StatPhys (top), BayeER (middle), BayeEF (bottom)). The aggregate vulnerabilities are computed as the mean over a sample of 10,000 reconstructed networks. Additionally, we show the range (labelled “Range”) of the aggregate vulnerabilities from this sample and the minimum and maximum (labelled “Optim Range”) of aggregate vulnerabilities obtained by minimising or maximising the aggregate vulnerability over all asset holding matrices consistent with the row and column sums. The horizontal line (labelled “True”) shows the aggregate vulnerability computed based on the true network, and the dashed horizontal line (labelled “Recon”) shows the reconstructed aggregate vulnerability using the true density for the reconstruction.

We find that even if the networks are reconstructed using a target density that does not coincide with the true density of the network, the aggregate vulnerabilities remain close to the true aggregate vulnerabilities.

We find that the range of the aggregate vulnerabilities computed using the StatPhys method exceeds the range of aggregate vulnerabilities derived by solving the optimisation problem that maximises or minimises the aggregate vulnerability for small densities. This is not a mistake, but a consequence of the StatPhys method not satisfying the row and column constraints exactly, but only in expectation.

We provide more empirical results (for a constant price impact) and further discussions on the sensitivity of the results with respect to additional information (such as the density of the network) in Appendix C.3.1. We also discuss there how additional information can be included in the MinDen and Entropy method.

Furthermore, we conduct a robustness check with respect to the main inputs into the network reconstruction: the row and column sums of the asset holding matrix. To do this we add a noise term to the row and column sums and reconstruct the networks based on the perturbed row and column sums. We find that the reconstructed direct, indirect, and aggregate vulnerabilities remain reasonably close to the true quantities in most cases even for noisy observation. The MinDen method seems to be the most sensitive with respect to the parameter inputs compared to the other reconstruction methods. Hence, the Entropy method and the sampling-based methods seem to be more robust in our case studies and one might therefore use those rather than the MinDen method if there is uncertainty about the input parameters. We report the detailed empirical results in Appendix C.3.2 (see specifically Table C.7).

Finally, we analyse the sensitivity of our results with respect to the selling rule μ for a constant price impact.[19] As discussed before, for a capitalisation-dependent price impact the selling rule does not matter. We solve the maximisation problem (10) and the corresponding minimisation problem that determines upper and lower bounds on γ_{n1}. Fig. 2 reports the results for the 2011 and 2016 data. It shows boxplots (and violin plots, i.e., the corresponding densities) of $\gamma_{11}, \dots, \gamma_{N1}$ corresponding to the proportional selling rule by Greenwood et al. (2015) together with the upper and lower bound from the optimisation problem and the constant $\gamma^{\text{Entropy}} = \gamma_{n1}(X^{\text{Entropy}})$ (see Proposition A.1)) derived from using the Entropy reconstruction method. We see that the γ_{n1}, $n \in \mathcal{N}$ that correspond to the proportional selling rule by Greenwood et al. (2015) are rather similar for most of the banks and overall rather close to the upper bound. Only for 2016, we find a small number of banks whose parameters γ_{n1} are close to the lower bound. We also find that the estimate $\gamma^{\text{Entropy}} = \gamma_{n1}(X^{\text{Entropy}})$, $n \in \mathcal{N}$ that is obtained from using the Entropy reconstruction method (indicated by the dotted line in Fig. 2) is close to the median of the true γ_{n1}, $n \in \mathcal{N}$. Furthermore, we show in Appendix A.2.1 that the expected connectivity using the StatPhys method coincides with the connectivity derived using the Entropy method, formally $\gamma^{\text{Entropy}} = \mathbb{E}[\gamma_{n1}(X^{\text{StatPhys}})]$ for all $n \in \mathcal{N}$.

4. Assessing the effect of policy interventions under full and partial information

We now investigate how fire sales risk can be mitigated through policy interventions. In contrast to Greenwood et al. (2015), we investigate how well fire sales risk can be mitigated if a policymaker decides on an intervention without the full knowledge of the asset holding network. We focus on two types of interventions: leverage caps and capital injections.

4.1. Leverage caps

Greenwood et al. (2015) have analysed a range of policy interventions to mitigate the effects of fire sales. In particular, for a GIIPS shock, they consider the effect of debt renationalisation, an introduction of Eurobonds, ring-fencing risky assets, merging exposed banks with unexposed ones and leverage caps. They find that “capping leverage is the only policy that delivers a sizeable reduction in $\mathcal{AV}$ [aggregate vulnerability]”, (Greenwood et al., 2015, p. 481). We will therefore focus on a leverage cap first.

The leverage cap policy in Greenwood et al. (2015) can be defined as follows.

Definition 4.1 *(Policy intervention: leverage cap).* For each bank $n \in \mathcal{N}$ with leverage d_{n1}/e_{n1}, the leverage cap policy sets the target leverage to $b_{n1} = \min\left\{B, \frac{d_{n1}}{e_{n1}}\right\}$ for a constant $B > 0$. It is assumed that all banks

---

[19] In the fire sale literature, a wide range of liquidation strategies has been considered. Some are exogenous and some are the result of an optimisation problem, see e.g., Caballero and Simsek (2013), Feinstein (2017), Braouezec and Wagalath (2019), Banerjee and Feinstein (2021).

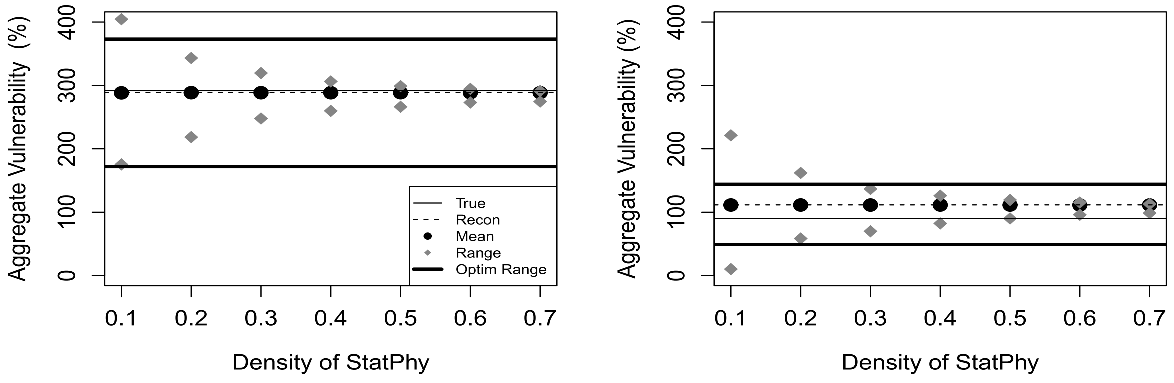

(a) StatPhys under GIIPS 2011. (b) StatPhys under Bad Brexit 2011.

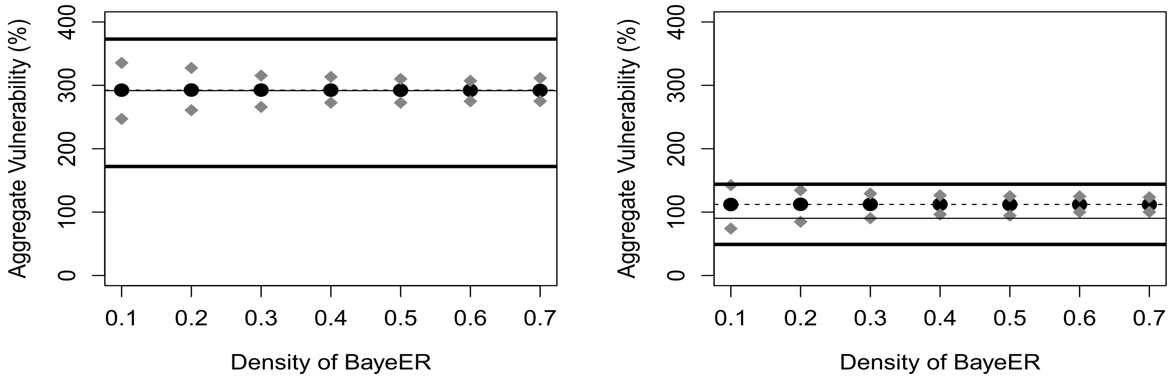

(c) BayeER under GIIPS 2011. (d) BayeER under Bad Brexit 2011.

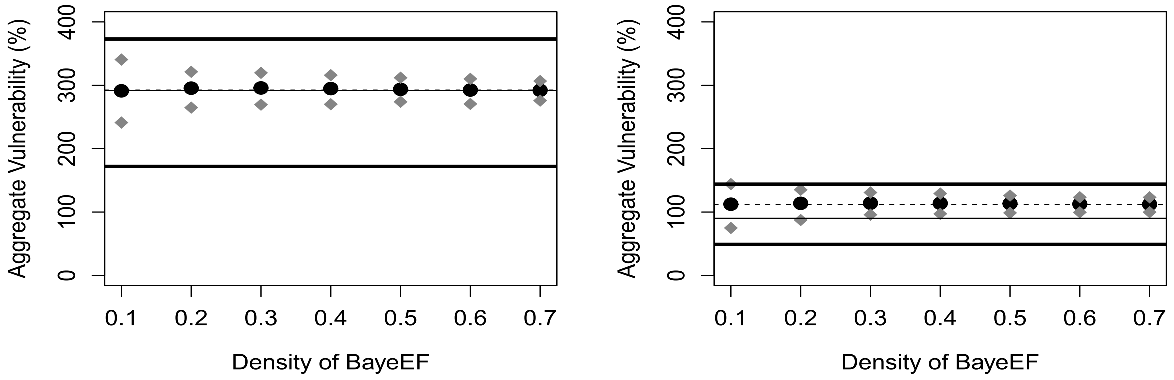

(e) BayeEF under GIIPS 2011. (f) BayeEF under Bad Brexit 2011.

Fig. 1. Aggregate vulnerabilities as a function of the network density for three sampling-based reconstruction methods (StatPhys (top), BayeER (middle), BayeEF (bottom)). The results are for the 2011 data and a capitalisation-dependent price impact.

$n \in \mathcal{N}$ for which $d_{n1}/e_{n1} > B$, are able to raise equity to reach the new lower leverage of B without changing the size of their balance sheet.

In our empirical analysis, we set the leverage cap to be $B = 15$. Recall that all banks in the sample had leverage of at most 30 (after an initial cap had been applied).

A reduction in target leverage automatically reduces the need to fire-sell assets as shown in Greenwood et al. (2015) and hence such a strategy reduces the aggregate vulnerability.

The potential problem with such a strategy is that banks might need to raise a significant amount of equity to satisfy a leverage cap that significantly reduces fire sales risk, “The cost of the policy is large [...], and the action is drastic” (Greenwood et al., 2015, p. 481).

To illustrate the cost of this strategy, we present Fig. 3. For each bank n, it shows the change in the $\mathcal{SYS}(n)$ between a target leverage at 30 and at 15, relative to the raised equity for each individual bank. In this example, we consider the capitalisation-dependent price impact ($\rho = -\log(0.1)$) and an all asset shock ($f_k = f = -0.001, \forall k \in \mathcal{S}$),

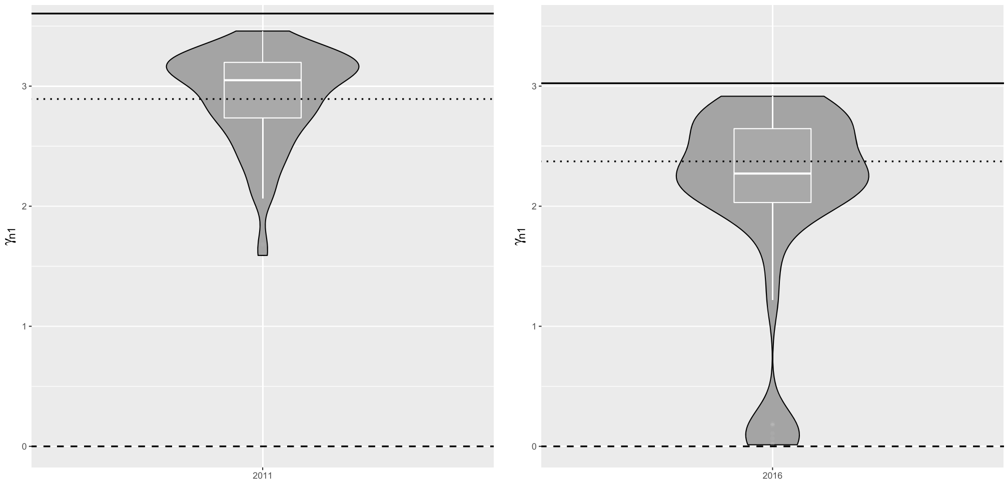

(a) 2011 data (b) 2016 data

Fig. 2. The box and violin plots of $(\gamma_{n1})_{n\in\mathcal{N}}$ are based on the full information and a constant price impact. There are three horizontal lines: The solid line is at $\max_{\mu\in[0,1]^{N\times K},\sum_{k=1}^{K}\mu_{nk}=1}\gamma_{n1}^{(\mu)}$ (which is the same for all n), the dashed line is at $\min_{\mu\in[0,1]^{N\times K},\sum_{k=1}^{K}\mu_{nk}=1}\gamma_{n1}^{(\mu)}$ (which again is the same for all n), and the dotted line is at $\gamma^{\text{Entropy}} = \gamma_{n1}(X^{\text{Entropy}})$.

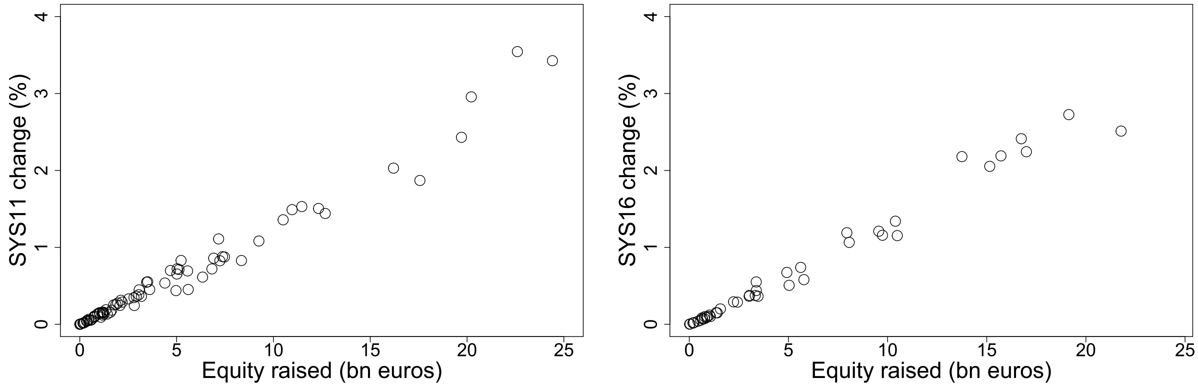

(a) Changes in $\mathcal{SYS}(n)$ for $n \in \mathcal{N}$ in 2011. (b) Changes in $\mathcal{SYS}(n)$ for $n \in \mathcal{N}$ in 2016.

Fig. 3. The scatter plots represent the changes in the $\mathcal{SYS}$ between capping target leverage at 30 and at 15, relative to the total amount of equity raised for 2011 (left) and 2016 (right). The displayed points are for banks whose target leverage is higher than the leverage cap of 15.

see Section 3.4. We observe a linear relationship between the change in $\mathcal{SYS}(n)$ values and the equity raised. For both plots (i.e., for both years 2011 and 2016), there is a cluster of banks where a smaller increase in equity results in small decreases in the contributed equity loss. The effect of a leverage cap only becomes distinctive between banks when larger equity values are considered, in this case past 10bn EUR. We also see that in 2011, a larger amount of equity needs to be raised by some banks to reach the same target leverage compared to 2016.

Overall, we find that banks would need to raise 357.2 billion EUR in 2011 and 233.85 billion EUR in 2016 to satisfy a leverage cap of 15.

Remark 4.2 *(Leverage cap under partial information).* The leverage cap policy does not depend on the asset holding matrix X. Therefore, implementing a leverage cap is equally successful with or without full information on the asset holding matrix.

4.2. Capital injections

Since reducing fire sales externalities via leverage caps is expensive, Greenwood et al. (2015) considered optimal equity injection as an alternative and the most cost-effective way to reduce aggregate vulnerability. The assumption is that a regulator has a fixed amount of cash $I > 0$ that can be distributed among the N banks. We first define a capital injection policy.

Definition 4.3 *(Policy intervention: capital injection).* Let $0 < I \le \sum_{n=1}^{N} d_{n1}$ be the total amount of cash that a regulator is willing to invest in banks' equity at time 1. Then, a capital injection policy is characterised by a vector $\underline{i} = (i_1, \dots, i_n)^\top$, where $0 \le i_n \le d_{n1}$ $\forall n \in \mathcal{N}$ and $\sum_{n=1}^{N} i_n = I$. Each bank n uses its capital injection i_n to repay parts of its debt, leading to a new leverage after the capital injection of

$$b_{n1}^{*} = \frac{d_{n1} - i_n}{e_{n1} + i_n}. \tag{12}$$

We assume that capital injections would occur in time period $t = 1$. One could extend the analysis to multiple rounds of deleveraging and multiple rounds of capital injections.

Now, the goal is to find an optimal capital injection policy, i.e., an optimal choice of $\underline{i} = (i_1, \dots, i_N)^\top$. Greenwood et al. (2015) considered the objective to minimise the systemicness of each bank under a GIIPS shock subject to some budget constraints. We consider the aggregate vulnerability as the objective function, which is just the sum of the systemicness of each bank. This allows us to consider the key system-wide measure developed in Greenwood et al. (2015) not just for measuring fire sale risk but also for mitigating it. We define an optimal capital injection as follows.

Definition 4.4 *(Policy intervention: optimal capital injection).* Let $0 < I \leq \sum_{n=1}^{N} d_{n1}$. Let $\mathcal{AVI} : [0, d_{11}] \times \ldots \times [0, d_{N1}] \to [0, \infty)$ be given by

$$\mathcal{AVI}(\underline{i}; X) = \sum_{n=1}^{N} \gamma_{n1}(X)(-R_{n1}(X)) \frac{\alpha_{n1}}{\sum_{v=1}^{N} e_{v1}} \frac{d_{n1} - i_n}{e_{n1} + i_n},$$

where X denotes the asset holding matrix. Consider the optimisation problem

$$\min_{\underline{i}=(i_{11},\ldots,i_{N1})^\top} \mathcal{AVI}(\underline{i}; X),$$

subject to

$$0 \leq i_n \leq d_{n1} \quad \forall n \in \mathcal{N}, \qquad \sum_{n=1}^{N} i_n = I. \tag{13}$$

We refer to a solution $\underline{i}^{\text{Opt}}(X) = (i_1^{\text{Opt}}(X), \ldots, i_N^{\text{Opt}}(X))^\top$ of (13) as an *optimal capital injection policy*.

Greenwood et al. (2015) find that the optimal capital injections are strongly positively correlated with systemicness, i.e., the optimal i_n^{Opt} are positively correlated with $\mathcal{SYS}(n)$. We therefore also consider a simplified capital injection strategy, in which the injected capital is chosen to be proportional to the systemicness.

Definition 4.5 *(Policy intervention: proportional capital injection).* Let X be an asset holding matrix. We refer to a capital injection $\underline{i}^{\text{Prop}} = (i_1^{\text{Prop}}(X), \ldots, i_N^{\text{Prop}}(X))^\top$, where

$$i_n^{\text{Prop}}(X) = I \frac{\mathcal{SYS}(n)(X)}{\mathcal{AV}(X)} \tag{14}$$

for all $n \in \mathcal{N}$ as a *proportional capital injection policy*.

As a benchmark strategy for capital injections, we consider a “naive” strategy, that allocates capital relative to the total asset holdings of banks. This strategy is independent of the network topology.

Definition 4.6 *(Policy intervention: naive capital injection).* We refer to a capital injection $\underline{i}^{\text{Naive}} = (i_1^{\text{Naive}}, \ldots, i_N^{\text{Naive}})^\top$, where

$$i_n^{\text{Naive}} = I \frac{\alpha_{n1}}{\sum_{v=1}^{N} \alpha_{v1}} \tag{15}$$

for all $n \in \mathcal{N}$ as a *naive capital injection policy*.

Remark 4.7 *(Choice of total capital I).* For our empirical analysis, we assume that the total allocation of capital I is set to 10% of the total equity of the banks, i.e., $I = 0.1 \sum_{n=1}^{N} e_{n1}$. This means that in 2011 we have $I = 70.55$ billion EUR and in 2016 $I = 68.21$ billion EUR.

Finally, to be able to compare capital injections to leverage caps, we consider a capital injection strategy that injects capital such that all institutions have leverage of at most $\tilde{B} > 0$. We define it formally as follows.

Definition 4.8 *(Policy intervention: leverage cap capital injection).* We refer to the capital injection $\underline{i}^{\text{lev}} = (i_1^{\text{lev}}, \ldots, i_N^{\text{lev}})^\top$, where

$$i_n^{\text{lev}} = \max\left\{ \frac{\alpha_{n1}}{1 + \tilde{B}} - e_{n1}, 0 \right\} \tag{16}$$

as the *leverage cap capital injection policy* with leverage cap $\tilde{B} > 0$.

Indeed, for an $n \in \mathcal{N}$ it holds that

$$i_n^{\text{lev}} = \frac{\alpha_{n1}}{1 + \tilde{B}} - e_{n1} > 0 \Leftrightarrow \frac{\alpha_{n1}}{1 + \tilde{B}} > e_{n1} \Leftrightarrow b_{n1} = \frac{\alpha_{n1} - e_{n1}}{e_{n1}} > \tilde{B}.$$

Hence, the leverage cap capital injection policy injects capital in exactly those institutions that exceed the leverage cap $\tilde{B}$. Furthermore, for all $n \in \mathcal{N}$ with $b_{n1} > \tilde{B}$, it follows directly from the definition of i_n^{lev} that $\frac{d_{n1} - i_n^{\text{lev}}}{e_{n1} + i_n^{\text{lev}}} = \tilde{B}$, i.e., those institutions that previously exceeded the leverage cap get a capital injection to reach the leverage of $\tilde{B}$.

Remark 4.9 *(Choice of leverage cap $\tilde{B}$ in leverage cap capital injection).* In our empirical analysis, we determine the leverage cap $\tilde{B}$, by solving

$$\sum_{n=1}^{N} i_n^{\text{lev}} = \sum_{n=1}^{N} \max\left\{ \frac{\alpha_{n1}}{1 + \tilde{B}} - e_{n1}, 0 \right\} = I$$

for $\tilde{B}$ for a given total capital of I that is chosen as in Remark 4.7. Then, the total amount of capital used in the leverage cap capital injection strategies coincides with the total capital used in the other capital injection strategies. This allows us to compare these strategies directly.

We find that injecting a total amount of $I = 70.55$ billion EUR in 2011 corresponds to a leverage cap of $\tilde{B} = 26.77$ for the leverage cap capital injection method in 2011; injecting a total amount of $I = 68.21$ billion EUR in 2016 corresponds to a leverage cap of 21.68 in 2016.

Remark 4.10 *(Capital injection under partial information).* By construction, the naive capital injection policy and the leverage cap capital injection strategy do not depend on the individual entries of the asset holding matrix X. The optimal capital injection policy and the proportional capital injection policy, however, will usually depend on the individual entries of the asset holding matrix X.

4.3. Empirical results on policy interventions

We will now analyse how well leverage caps and the different capital injection strategies work in the 2011 and 2016 data, under both full and partial information. To do so, we compute the relative reduction in aggregate vulnerability between the network without intervention and the network with intervention. We analyse these policies for a GIIPS shock of 5% that we have already considered in the previous section.

4.3.1. Empirical results - leverage cap

First, we consider the intervention of capping the leverage. As already discussed, the leverage cap intervention is independent of the underlying network. Hence, the relative reduction in aggregate vulnerability between the network without a leverage cap and the network with a leverage cap relative to the network without a leverage cap is given by

$$\Delta\mathcal{AV}^{\text{Leverage cap}} = \frac{\mathcal{AV} - \mathcal{AV}^{\text{Leverage cap}}}{\mathcal{AV}}$$

under both full and partial information. Here, $\mathcal{AV}^{\text{Leverage cap}}$ refers to the aggregate vulnerability that is obtained by setting the target leverage to $b_{n1} = \min\{B, \frac{d_{n1}}{e_{n1}}\}$, where we consider two choices of B: $B = 15$ and $B = \tilde{B}$, where $\tilde{B} = 26.77$ in 2011 and $\tilde{B} = 21.68$ in 2016. The choices of $\tilde{B}$ correspond to the leverage caps derived in Remark 4.9, i.e., the amount of equity that needs to be raised to achieve this cap, corresponds to the total amount of capital used in the capital injection strategies. In this sense, the costs of the leverage cap strategy with the target leverage of $b_{n1} = \min\{\tilde{B}, \frac{d_{n1}}{e_{n1}}\}$ coincides with the cost of the capital injection policies.

Table 6 reports the results. When capping the leverage at 15, we see that the leverage cap in 2011 yields a much larger relative reduction in aggregate vulnerability compared to 2016. This is not surprising, since the leverages and the target leverages were generally higher in 2011

Table 6

Relative decrease (in percent) in aggregate vulnerability for the EBA 2011 and 2016 data for different policies (capital injections and leverage caps) and for two different price impacts (capitalisation-dependent and constant). Values in bold indicate which network reconstruction method performed best when used for a given capital injection strategy. Values in a box represent the best capital injection method in a given year and for a given price impact. Entries labelled "−" indicate that these values coincide with the value reported in the column labelled *True*.

		Capitalisation-Dependent Price Impact ($-\log(0.1)/c_k\ \forall k$)						Constant Price Impact ($5 \times 10^{-13}\ \forall k$)					
		True	MinDen	Entropy	StatPhys	BayeER	BayeEF	True	MinDen	Entropy	StatPhys	BayeER	BayeEF
GIIPS (%)													
2011	Naive capital injection	7.29	-	-	-	-	-	7.43	-	-	-	-	-
-	Proportional capital injection	16.11	3.52	7.47	7.28	8.99	**9.88**	15.94	3.66	7.64	7.82	**8.55**	8.31
-	Proportional capital injection (Average)	-	-	-	7.47	**9.48**	9.32	-	-	-	7.63	**8.77**	8.71
-	Optimal capital injection	19.81	2.95	**10.56**	7.54	8.44	9.95	19.51	3.09	**10.52**	7.93	9.16	9.56
-	Optimal capital injection (Average)	-	-	-	10.31	$\boxed{\mathbf{11.98}}$	11.80	-	-	-	10.38	11.27	$\boxed{\mathbf{11.44}}$
-	Leverage Cap ($\tilde{B} = 26.77$)	3.53	-	-	-	-	-	3.41	-	-	-	-	-
-	Leverage Cap ($B = 15$)	36.51	-	-	-	-	-	36.44	-	-	-	-	-
2016	Naive capital injection	8.97	-	-	-	-	-	9.06	-	-	-	-	-
-	Proportional capital injection	23.67	3.10	8.33	8.35	**10.60**	7.98	22.53	3.82	8.47	8.46	$\boxed{\mathbf{10.41}}$	6.85
-	Proportional capital injection (Average)	-	-	-	8.34	**10.17**	10.16	-	-	-	8.47	**9.57**	9.50
-	Optimal capital injection	27.03	3.42	5.57	8.17	10.47	$\boxed{\mathbf{10.66}}$	25.66	3.96	5.65	8.46	**8.53**	8.30
-	Optimal capital injection (Average)	-	-	-	5.99	9.46	**9.56**	-	-	-	6.07	**8.69**	8.61
-	Leverage Cap ($\tilde{B} = 21.68$)	2.96	-	-	-	-	-	2.79	-	-	-	-	-
-	Leverage Cap ($B = 15$)	18.39	-	-	-	-	-	19.07	-	-	-	-	-

than in 2016 and therefore capping the leverage at 15 (from 30) has a much larger effect in 2011 than in 2016.

Capping the leverage at $\tilde{B}$, which in both years is significantly larger than 15, yields smaller relative decreases in aggregate vulnerability than capping at 15, which was to be expected. Of course, less equity needs to be raised to reach a higher cap at $\tilde{B}$ than reaching the lower cap of 15, but then one does not obtain the same benefit from it.

More interesting is the comparison between capping the leverage at $\tilde{B}$ and the capital injection strategies, since these strategies have comparable costs. We find that in both years, even the naive capital injection strategy outperforms the leverage cap strategy at a cap of $\tilde{B}$. More sophisticated capital injection strategies do generally outperform the leverage cap strategy at a cap of $\tilde{B}$ by a larger amount (even if they are used with a network reconstruction method rather than under full information).

4.3.2. *Empirical results - capital injection*

Second, we consider intervention via capital injection. We set the total capital I that is injected in the network to be equal to 10% of the total equity in the given network. We consider different capital injection strategies $\underline{i}$ and compute the relative reduction in aggregate vulnerability in the true financial network corresponding to such a capital injection strategy. In particular, for a given capital injection strategy $\underline{i}$ the corresponding relative reduction in aggregate vulnerability is given by

$$\Delta\mathcal{AV}^{\text{Injection}}(\underline{i}) = \frac{\mathcal{AVI}(\underline{0}; X^{\text{true}}) - \mathcal{AVI}(\underline{i}; X^{\text{true}})}{\mathcal{AVI}(\underline{0}; X^{\text{true}})}, \tag{17}$$

where $\underline{0}$ is the N-dimensional zero vector, and therefore $\mathcal{AVI}(\underline{0}; X^{\text{true}})$ represents the aggregate vulnerability in the fully observed financial network with zero capital injection.

Again, Table 6 shows the results. We first look at the results under full information, i.e., the columns labelled *True*, meaning that the strategy $\underline{i} = \underline{i}(X^{\text{true}})$ is computed based on the fully observed asset holding matrix X. They show the relative reduction in aggregate vulnerability when the capital injection $\underline{i}$ was computed from the fully observed asset holding matrix X^{true}. The optimal capital injection policy performs best throughout which it should do. What is interesting, is that the proportional capital injection strategy still performs only slightly worse than the optimal injection strategy. This implies that injecting capital proportional to the systemicness of the nodes seems to be a good approximation to the optimal strategy derived from solving (13). This is further confirmed by Fig. 4 which shows a scatter plot of the optimal capital injection strategy under full information plotted against the proportional capital injection policy under full information. We see that these two strategies are indeed very similar.

The naive capital injection strategy performs worst throughout with a relative reduction of aggregate vulnerability of around 7% in 2011 and 9% in 2016, respectively, which is significantly lower than e.g., the 19% and 25-27% reductions achieved by using the optimal capital injection policy. Hence, we see that there is a clear benefit of using an optimal strategy (or an approximation of the optimal strategy) in the case of full information.

Next, we consider the potential benefits of the different capital injection strategies if the vector $\underline{i}$ representing the capital injections is determined under partial information by using matrix reconstruction. Any type of capital injection improves aggregate vulnerability. Therefore, even under partial information, capital injections will still reduce the overall aggregate vulnerability. It is not clear, however, how much reduction in aggregate vulnerability can be achieved and this is what we investigate here.

Since the naive capital injection policy is independent of the network, the relative reduction in aggregate vulnerability under full and partial information is the same. Hence, we only consider the optimal and the proportional capital injection strategy under partial information.

For the optimisation-based matrix reconstruction methods MinDen and Entropy, determining the capital injections based on partial information means that we compute $\underline{i}^{\text{Opt}}(\hat{X})$, $\underline{i}^{\text{Prop}}(\hat{X})$, where $\hat{X}$ is the reconstructed matrix that is either derived using the MinDen or the Entropy method. Both optimisation-based reconstruction methods only return one matrix, therefore the corresponding strategies are well defined.[20] Then, we consider the relative reduction in aggregate vulnerabilities $\Delta\mathcal{AV}^{\text{Injection}}(\underline{i}^{\text{Opt}}(\hat{X}))$ and $\Delta\mathcal{AV}^{\text{Injection}}(\underline{i}^{\text{Prop}}(\hat{X}))$ as given in (17).

---

[20] Since the reconstructed matrix obtained from the Entropy method is available in closed form, we can also express $\underline{i}^{\text{Prop}}(X^{\text{Entropy}})$ analytically. In particular, we show in Corollary A.3 that

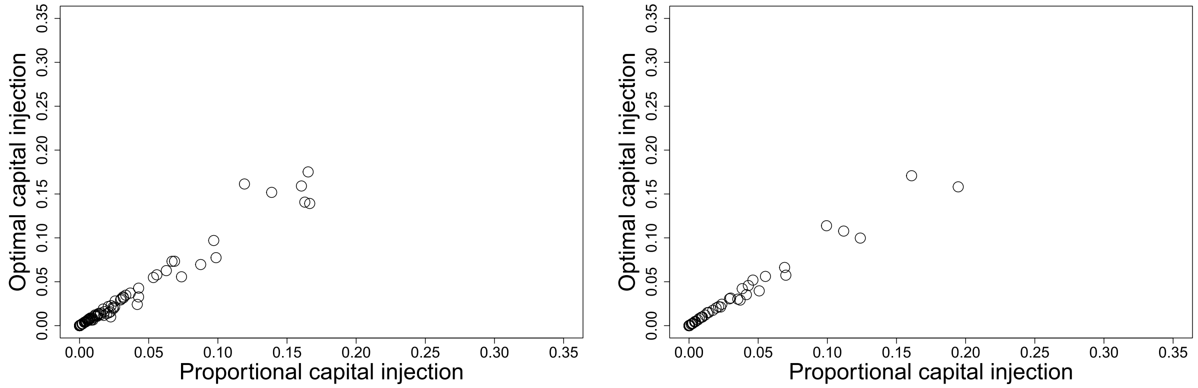

(a) $\mathcal{SYS}$ under capital injections for 2011 data.

(b) $\mathcal{SYS}$ under capital injections for 2016 data.

Fig. 4. The plot of the optimal capital injection strategy against the proportional capital injection strategy under full information for a 5% GIIPS shock and a capitalisation-dependent price impact.

Since the sample-based reconstruction methods StatPhys, BayeER, and BayeEF return not just one reconstructed network but a sample of reconstructed networks, there are different ways how we can compute the proportional and optimal capital injection strategy under partial information. We will consider two approaches: the first approach will just average the strategies derived from the different reconstructed networks. The second approach will choose a strategy associated with the tail of the distribution of aggregate vulnerabilities.

Consider a sample of asset holding matrices $X^{(1)}, \dots, X^{(d)}$ and denote by X^{true} the true asset holding matrix. One possible approach is to compute the proportional or optimal injection strategy on every network $X^{(v)}$ of the sample, i.e., determine $\underline{i}^{\text{Prop}}(X^{(v)})$ and $\underline{i}^{\text{Opt}}(X^{(v)})$ and then consider the sample averages of these strategies given by

$$\underline{i}^{\text{Prop, average}} = \frac{1}{d}\sum_{v=1}^{d} \underline{i}^{\text{Prop}}(X^{(v)}), \qquad \underline{i}^{\text{Opt, average}} = \frac{1}{d}\sum_{v=1}^{d} \underline{i}^{\text{Opt}}(X^{(v)}).$$

We will refer to these strategies as proportional capital injection (average) and optimal capital injection (average) in Table 6.

In addition to these average strategies, we are also interested in the tails of the distribution of aggregate vulnerabilities under capital injection. For the proportional capital injection strategy we determine the injection strategy that corresponds to the 95% percentile of the empirical distribution function of the sample of aggregate vulnerabilities under proportional capital injection, i.e., we determine the index $\tilde{v} \in \{1, \dots, N\}$ such that

$$\mathcal{AVI}(\underline{i}^{\text{Prop}}(X^{(\tilde{v})}); X^{(\tilde{v})}) = \inf\left\{ x \in \mathbb{R} \mid \frac{1}{d}\sum_{v=1}^{d} \mathbb{I}_{\{\mathcal{AVI}(\underline{i}^{\text{Prop}}(X^{(v)}); X^{(v)}) \le x\}} \ge 0.95 \right\}$$

and we denote this index by $v^{(0.95)}$.[21] Hence, this corresponds to one of the highest aggregate vulnerabilities observed in the sample in which proportional capital injection was used. We then compute the relative reduction in aggregate vulnerability that corresponds to the strategy $\underline{i}^{\text{Prop}}(X^{(v^{(0.95)})})$ and report this in Table 6 (in the row *Proportional capital injection*).

For the optimal capital injection policy we consider the optimisation problem which aims to find the capital injection strategy that minimises the 0.95-Percentile of the empirical cumulative distribution function of the aggregate vulnerabilities derived from the d sample networks with capital injection. Formally, we consider

$$\min_{\underline{i}} \quad \inf\left\{ x \in \mathbb{R} \mid \frac{1}{d}\sum_{v=1}^{d} \mathbb{I}_{\{\mathcal{AVI}(\underline{i}; X^{(v)}) \le x\}} \ge 0.95 \right\},$$

subject to

$$0 \le i_n \le d_{n1} \quad \forall n \in \mathcal{N}, \qquad \sum_{n=1}^{N} i_n = I. \tag{18}$$

We then consider the strategy that is a solution to (18) and report the corresponding relative reduction in aggregate vulnerability in Table 6.

When considering the performance of the proportional and optimal capital injection policies under partial information in Table 6 we see that the reduction in relative aggregate vulnerability is significantly lower under partial information than under full information for all types of reconstruction methods. Furthermore, we see that under partial information, sometimes the proportional capital injection policy performs better than the optimal capital injection strategy. For example, for the 2016 data and a constant price impact, the proportional capital injection strategy based on the Entropy method gives a relative reduction of the aggregate vulnerability of 8.47%, compared to a relative reduction of 5.65% achieved by the optimal capital injection strategy. The reason for this is that the optimal capital injection is optimal for the reconstructed asset holding matrix X^{Entropy} and not necessarily optimal for the true matrix X^{true}. When evaluating the performance of the different strategies, however, we use the true matrix X^{true} to compute the relative reduction of aggregate vulnerability (17). Under full information, the optimal capital injection strategy cannot perform worse than the proportional capital injection strategy. Still, we overall find that in our four test cases (2 years and 2 price impacts), the best-performing capital injection strategy (indicated by a box in Table 6) is an optimal capital injection strategy in three cases[22] and a proportional capital injection strategy in only one case.[23]

Among the different reconstruction methods, the MinDen method performs worst in both years and for both choices of price impact. Capital injection strategies that rely on the MinDen method only reduce the relative aggregate vulnerability by around 3%. This level of reduction is therefore much lower than the reduction of around 7-9% that can be

---

$i_n^{\text{Prop}}(X^{\text{Entropy}}) = I\,\frac{\mathcal{SYS}^{X^{\text{Entropy}}}(n)}{\mathcal{AV}^{X^{\text{Entropy}}}} = I\,\frac{\alpha_{n1} b_{n1}}{\sum_{v=1}^{N} \alpha_{v1} b_{v1}}$.

[21] If there is more than one index $\tilde{v}$ satisfying the equation we select one suitable index randomly.

[22] for both price impacts in 2011 and for the capitalisation-dependent price impact in 2016.

[23] for a constant price impact in 2016.

achieved with the naive capital injection strategy that does not even attempt to reconstruct the underlying network.

The other network reconstruction methods, i.e., Entropy, StatPhys, BayeER, and BayeEF perform generally better when used to decide on capital injections. Out of these four methods, the Bayesian methods seem to perform best overall. For the four test cases (2 years and 2 price impacts), the best capital injection strategy (indicated by a box in Table 6) is always one that uses a Bayesian network reconstruction. For the 2011 data, the best capital injection strategy is the optimal capital injection (average) using the BayeER method under the capitalisation-dependent price impact and the BayeEF method under the constant price impact. They reduce the aggregate vulnerability by 11.98% and 11.44% respectively, which is better than the naive strategy which achieves a relative reduction between 7.24 - 7.43%.
For the 2016 data, the optimal capital injection method using the BayeER method is best under the capitalisation-dependent price impact assumption (10.66% relative reduction in aggregate vulnerability compared to 8.97% achieved by the naive capital injection strategy). For the constant price impact assumption, the proportional capital injection strategy using the BayeER method performs best (achieving a relative reduction in aggregate vulnerability of 10.41% compared to 9.06% achieved by the naive strategy).

When fixing the type of capital injection strategy (proportional capital injection, proportional capital injection (average), optimal capital injection, optimal capital injection (average)), and then checking which network reconstruction method performs best in a given year and for a given price impact, then we find that out of the 16 cases, the BayeER performs best in 10 cases, BayeEF performs best in 4 cases and the Entropy method performs best in 2 cases. The StatPhys is never the best-performing method in our examples but still performs reasonably well.

Overall the best-performing capital injection methods using network reconstruction methods reduce the relative aggregate vulnerability in the range between 10 - 11% and are therefore better than the naive capital injection strategy which achieves a reduction between 7 - 9%. In particular, in each of the two years and for both types of price impact we see that all capital injection strategies that use the BayeER network reconstruction always outperform the naive capital injection strategy and the BayeER is the only network reconstruction method considered here for which this is the case.

Hence, we see that using suitable network reconstruction methods to decide on risk mitigation mechanisms in financial networks is indeed beneficial and can achieve better outcomes than using naive intervention strategies.

5. Conclusion

We have investigated how well fire sales risk can be measured and mitigated under partial information. We used the fire sales measures (systemicness, aggregate vulnerability, direct vulnerability and indirect vulnerability) developed by Greenwood et al. (2015) and analysed their dependence on the asset holdings matrix. We then investigated how well these four different measures quantifying risk associated with fire sales can be estimated when the individual entries of the underlying asset holdings matrix are not observable but its row and column sums are. We considered two empirical asset holding matrices, available in the data published by the EBA for their 2011 and 2016 stress tests, and assumed that they were not fully observable. We estimated the asset holding matrix using five different network reconstruction methods available in the literature and found that in general these fire sales measures could be estimated reasonably accurately for a range of shock scenarios.

We then analysed how well risk from fire sales can be mitigated if policy interventions are based on partial information and network reconstruction techniques are used to decide on policies. We considered two policies that were highly effective in the analysis under full information in Greenwood et al. (2015), namely leverage caps and capital injections.

Leverage caps are generally independent of the underlying network and therefore do not require network reconstruction techniques to implement them. In 2011 leverage caps lead to better outcomes than capital injections, but in 2016 when banks' leverages were generally lower, intervention via capital injections leads to better outcomes than leverage caps. Therefore, in the more recent data, capital injections appear more beneficial.

Capital injections can be done using ad hoc methods that do not rely on the asset holding matrix or can be done in a more targeted approach that would account for characteristics of the asset holding matrix. We considered a naive capital injection strategy in which capital is injected in proportion to the size of a bank (measured in terms of the total assets on its balance sheet); no information on the individual entries of the asset holding matrix is needed for this approach. We compare this to capital injection strategies that inject capital in proportion to the systemicness of an institution or in an optimal way (with the objective of reducing the aggregate vulnerability) and these methods then rely on the (reconstructed) asset holding matrix.

We find that it is possible to achieve a significant relative reduction in aggregate vulnerability even under partial information. While the naive capital injection strategy, which does not require network reconstruction, achieves relative reductions in aggregate vulnerability in the range of 7 - 9% in our study, the best-performing capital injection strategies that rely on network reconstruction methods achieved a relative reduction of aggregate vulnerability between 10 - 11%. We found that the Bayesian method (Gandy and Veraart, 2017, 2019) for network reconstruction was the best overall method when used for deciding on capital injections. In particular, we found that any capital injection strategy that we considered that was based on the Bayesian network reconstruction method with an Erdős-Rényi-type prior, always outperformed the naive capital injection strategy.

Hence, we see that network reconstruction techniques are not just useful for measuring risk, but also for managing it. As we have already discussed, the intervention strategies considered here can never do any harm (in the sense that using them cannot increase the aggregate vulnerability of the network). So it was clear that even using them in a non-optimal way can bring potential benefits. What is interesting, however, is to see how much better some of them perform in comparison to naive strategies that do not attempt to reconstruct the underlying network.

CRediT authorship contribution statement

Raymond Ka-Kay Pang: Conceptualization, Data curation, Formal analysis, Methodology, Software, Writing – original draft, Writing – review & editing. **Luitgard Anna Maria Veraart:** Conceptualization, Formal analysis, Methodology, Software, Supervision, Writing – original draft, Writing – review & editing.

Data availability

The data are publicly available. The corresponding website has been provided in the paper.

Acknowledgement

We would like to thank the editor, the associate editor, and the referees for their constructive comments. The first author gratefully acknowledges financial support by a PhD Studentship awarded by the London School of Economics and Political Science.

Appendix A. Background information on matrix reconstruction methods

The matrix reconstruction methods considered can be classified into optimisation-based reconstruction methods, i.e., they determine a matrix that is consistent with given row and column sums by solving a deterministic optimisation problem, and sampling-based reconstruction methods, which assume that the matrix of interest is random and they develop tools to generate a sample from the distribution of the matrix.

A.1. Optimisation-based reconstruction methods

We consider two matrix reconstruction methods, the Entropy method and the MinDen method. Both solve suitable optimisation problems to identify a matrix that is consistent with given row and column sums.

A.1.1. Entropy method

The method that we refer to as the Entropy method in this paper, is also known under several other names, such as iterative proportional fitting procedure, or RAS algorithm, to name a few and has been used in several fields, e.g., in mathematics, economics, computer science etc. To the best of our knowledge it has first been applied to financial networks by Upper and Worms (2004) who used the method to reconstruct a network of interbank liabilities from row and column sums. It has also been considered in the context of reconstructing networks of asset holding matrices in Di Gangi et al. (2018). The Entropy method is an optimisation-based method that minimises the Kullback-Leibler (KL) divergence between a matrix X and a target matrix X^{Entropy}. Applied to our setting it consists of solving the following optimisation problem

$$\begin{aligned} &\min_X \sum_{n=1}^{N}\sum_{k=1}^{K} X_{nk}\log\left(\frac{X_{nk}}{X_{nk}^{\text{Entropy}}}\right), \\ \text{subject to: } &\alpha_{n1} = \sum_{k=1}^{K} X_{nk} \quad \forall n \in \{1,\dots,N\}, \\ &c_k = \sum_{n=1}^{N} X_{nk} \quad \forall k \in \{1,\dots,K\}, \\ &X_{nk} \geq 0 \quad \forall n \in \{1,\dots,N\}, \forall k \in \{1,\dots,K\}, \end{aligned} \tag{A.1}$$

where the initial matrix is defined as

$$X_{nk}^{\text{Entropy}} = \frac{\alpha_{n1}c_k}{A} \quad \forall n \in \{1,\dots,N\}, \forall k \in \{1,\dots,K\}, \tag{A.2}$$

where $A = \sum_{n=1}^{N}\alpha_{n1} = \sum_{k=1}^{K} c_k$.

One can easily check that X^{Entropy} solves this optimisation problem. The reason why this reconstruction problem simplifies so significantly in our situation is that since we consider an asset holdings matrix we only need to require the non-negativity of the matrix and that it satisfies the given row and column sums. We are not in a situation in which the diagonal entries of the matrix that solves the optimisation problem are required to be zero. This additional constraint occurs, for example, in Upper and Worms (2004), in which the network represents interbank lending. Since a bank does not borrow from itself the additional constraint, that the entries on the diagonal are zero, is necessary there.

As one can see from the definition of X^{Entropy}, the reconstructed matrix usually contains only non-zero entries (an entry X_{nk}^{Entropy} in the matrix can only be zero if the corresponding row α_{n1} or column c_k aggregate is zero).

It has been discussed in Di Gangi et al. (2018) how the specific form of X^{Entropy} here can be interpreted as reflecting investors' preference in line with the capital asset pricing model (CAPM) (Sharpe, 1964).

Since the Entropy method provides a closed-form expression for the reconstructed asset holding matrix, all fire sale measures by Greenwood et al. (2015) applied to this reconstructed matrix can be expressed in closed form. Therefore, one immediately obtains the following results.

Proposition A.1. *Suppose the asset holding matrix is estimated using $X^{Entropy}$ given in (A.2). Let $n \in \mathcal{N}$ and $k \in \mathcal{S}$. Then,*

1. *the elements of the portfolio weights matrix are given by $m_{nk}(X^{Entropy}) = \frac{X_{nk}^{Entropy}}{\alpha_{n1}} = \frac{c_k}{A}$;*
2. *the unlevered return is $R^{Entropy} = R_{n1}(X^{Entropy}) = \frac{\sum_{k=1}^{K} c_k f_{k1}}{A}$;*
3. *the connectivity is $\gamma^{Entropy} = \gamma_{n1}(X^{Entropy}) = \sum_{k=1}^{K} c_k l_k \frac{X_{nk}^{Entropy}}{\alpha_{n1}} = \frac{\sum_{k=1}^{K} c_k^2 l_k}{A}$;*
4. *the direct vulnerability is $\mathcal{DV}^{X^{Entropy}}(n) = -\frac{\alpha_{n1}}{e_{n1}} R^{Entropy}$;*
5. *the systemicness is $\mathcal{SYS}^{X^{Entropy}}(n) = \frac{-\gamma^{Entropy} R^{Entropy}}{\sum_{\nu=1}^{N} e_{\nu 1}} \alpha_{n1} b_{n1}$;*
6. *the aggregate vulnerability is $\mathcal{AV}^{X^{Entropy}} = \sum_{n=1}^{N} \mathcal{SYS}^{X^{Entropy}}(n) = \frac{-\gamma^{Entropy} R^{Entropy}}{\sum_{\nu=1}^{N} e_{\nu 1}} \sum_{n=1}^{N} \alpha_{n1} b_{n1}$;*
7. *the indirect vulnerability is $\mathcal{IV}^{X^{Entropy}}(n) = -\gamma^{Entropy} R^{Entropy} \frac{\sum_{\nu=1}^{N} \alpha_{\nu 1} b_{\nu 1}}{A} \times \frac{\alpha_{n1}}{e_{n1}}$.*

Proof of Proposition A.1. 1. The statement follows directly from the definition of X^{Entropy}, since $m_{nk}(X^{\text{Entropy}}) = \frac{X_{nk}^{\text{Entropy}}}{\alpha_{n1}} = \frac{\alpha_{n1}c_k}{\alpha_{n1}A} = \frac{c_k}{A}$.
2.-7. The statements follow directly from part 1. and the definitions of the risk measures. □

Hence, we find that for the Entropy reconstruction method, the two quantities that depend on the network $R_{n1}(X^{\text{Entropy}})$ and $\gamma_{n1}(X^{\text{Entropy}})$ do not depend on n, which means they are not specific to a given institution. For an all asset shock $f_{k1} = f$ for all $k \in \mathcal{S}$, $R^{\text{Entropy}} = f$ in line with Proposition 3.1.

Remark A.2 *(Comparison of systemicness and indirect vulnerability using the Entropy method).* These results show that under the Entropy method, the systemicness is a product of an institution-specific factor $\alpha_{n1}b_{n1}$ (representing total assets times leverage) and a common factor $\frac{-\gamma^{\text{Entropy}} R^{\text{Entropy}}}{\sum_{\nu=1}^{N} e_{\nu 1}}$. The indirect vulnerability also consists of an institution-specific factor $\frac{\alpha_{n1}}{e_{n1}}$ (representing total asset holdings divided by equity) and a common factor $-\gamma^{\text{Entropy}} R^{\text{Entropy}} \frac{\sum_{\nu=1}^{N}\alpha_{\nu 1}b_{\nu 1}}{A}$. Hence, we see that institutions with high total asset holdings times leverage will have a high systemicness, i.e., will play a major role in causing fire sale losses, whereas institutions with large total asset holdings divided by their equity will have a large indirect vulnerability, i.e., they will be susceptible to fire sale losses. So leverage influences systemicness, whereas equity influences indirect vulnerability.

Proposition A.1 also allows us to provide an analytical expression for the proportional capital injection strategy.

Corollary A.3. *Suppose the asset holding matrix is estimated using $X^{Entropy}$ given in (A.2). Let $n \in \mathcal{N}$ and $k \in \mathcal{S}$. Then, the proportional capital injection defined in Definition 4.5, reduces to*

$$i_n^{Prop}(X^{Entropy}) = I\,\frac{\mathcal{SYS}^{X^{Entropy}}(n)}{\mathcal{AV}^{X^{Entropy}}} = I\,\frac{\alpha_{n1}b_{n1}}{\sum_{\nu=1}^{N}\alpha_{\nu 1}b_{\nu 1}}.$$

This means that capital is injected relative to a measure in which the total assets is weighted by leverage. Therefore, $i_n^{\text{Prop}}(X^{\text{Entropy}})$ differs from the naive capital injection strategy $i_n^{\text{Naive}} = I\frac{\alpha_{n1}}{\sum_{\nu=1}^{N}\alpha_{\nu 1}}$ in which only the total asset holdings are considered.

A.1.2. Minimum density method

The minimum density method for network reconstruction was introduced in Anand et al. (2015) in the context of a matrix representing interbank lending. We apply it here to a matrix with a different economic interpretation, namely asset holdings between different banks. It solves an optimisation problem with the objective to find a matrix with the minimum number of edges that is consistent with given row and column sums. The resulting network is therefore usually very sparse. Formally, the optimisation problem in our setting is as follows.

$$\min_{X} \sum_{n=1}^{N} \sum_{k=1}^{K} \mathbb{1}_{\{X_{nk}>0\}},$$

$$\text{subject to: } \alpha_{n1} = \sum_{k=1}^{K} X_{nk} \quad \forall n \in \{1,\dots,N\},$$

$$c_k = \sum_{n=1}^{N} X_{nk} \quad \forall k \in \{1,\dots,K\},$$

$$X_{nk} \geq 0 \quad \forall n \in \{1,\dots,N\}, \forall k \in \{1,\dots,K\}.$$

Anand et al. (2015) provide an algorithm to solve this optimisation problem and also consider generalisations that result in less sparse matrices. We will mainly consider one matrix in our analysis that represents the sparsest solution. As part of our sensitivity analysis we also consider the generalisation by Anand et al. (2015) that constructs less sparse matrices.

A.2. Sampling-based reconstruction methods

We also consider two matrix reconstruction methods that assume that the matrix itself is random and provide methodologies to sample from the appropriate distribution.

A.2.1. Statistical physics method

The method that we refer to as the Statistical Physics method, due to its modelling ideas coming from this area, was developed by Cimini et al. (2015). It was originally proposed to reconstruct a network of interbank lending. It has then been applied to the case of bipartite networks of asset holding networks by Squartini et al. (2017) which is what we do here. Applied to our setting, it is characterised by an $N \times K$-dimensional random matrix X^{StatPhys}, whose individual entries X_{nk}, $n \in \mathcal{N}, k \in \mathcal{S}$ are independent random variables from the following discrete distributions.

$$\mathbb{P}\left(X_{nk}^{\text{StatPhys}} = \frac{\alpha_{n1} c_k}{p_{nk} A}\right) = p_{nk},$$

$$\mathbb{P}(X_{nk}^{\text{StatPhys}} = 0) = 1 - p_{nk},$$

where again $A = \sum_{\nu=1}^{N} \alpha_{\nu 1}$. Furthermore, $p_{nk} = \frac{\phi \alpha_{n1} c_k}{1+\phi \alpha_{n1} c_k} \quad \forall n \in \mathcal{N}, \forall k \in \mathcal{S}$, and $\alpha_{n1} > 0 \, \forall n \in \mathcal{N}$, $c_k > 0 \, \forall k \in \mathcal{S}$, are the given row and column sums, respectively and $\phi > 0$ is a parameter that can be used to calibrate the model.[24] One can check that $p_{nk} \in [0,1] \; \forall n \in \mathcal{N}, \forall k \in \mathcal{S}$. Hence, we see that each entry in the random matrix can only take two possible values - zero or another non-negative value.

It follows directly from the definition that

$$\mathbb{E}\left[\sum_{n=1}^{N} X_{nk}^{\text{StatPhys}}\right] = c_k, \quad \forall k \in \mathcal{S} \quad \text{and} \quad \mathbb{E}\left[\sum_{k=1}^{K} X_{nk}^{\text{StatPhys}}\right] = \alpha_{n1}, \quad \forall n \in \mathcal{N}.$$

This means, that the random matrix X^{StatPhys} satisfies the row and column sums in expectation. If one generates a sample of matrices from this probability distribution, then the individual matrices in the sample will usually not satisfy the row and column sums.

To calibrate the model to a given target density $\delta^{\text{target}} \in (0,1)$ of a network one can use the fact that the expected density of X^{StatPhys} is given by

$$f(\theta) = \frac{1}{NK} \mathbb{E}\left[\sum_{n=1}^{N} \sum_{k=1}^{K} \mathbb{1}_{\{X_{nk}^{\text{StatPhys}}>0\}}\right] = \frac{1}{NK} \sum_{n=1}^{N} \sum_{k=1}^{K} p_{nk}$$

$$= \frac{1}{NK} \sum_{n=1}^{N} \sum_{k=1}^{K} \frac{\phi \alpha_{n1} c_k}{1+\phi \alpha_{n1} c_k}$$

and solve $f(\theta) = \delta^{\text{target}}$ for θ. Note that f is a continuous and non-decreasing function satisfying $f(0) = 0$. Furthermore, $\lim_{\theta \to \infty} f(\theta) = 1$ if $\alpha_{n1} c_k > 0$ for all $n \in \mathcal{N}$ and for all $k \in \mathcal{S}$. Hence, we see that if all row and column sums are non-zero, the model can be calibrated to any target density $\delta^{\text{target}} \in (0,1)$. If some row or column sums are zero, then the underlying network cannot have a density of 1 or a similarly large value and this is indeed reflected by the function f.

Some fire sales measures evaluated using the StatPhys method are related to those evaluated under the Entropy method.

For example, the expected portfolio weights matrix using the StatPhys method satisfies

$$\mathbb{E}[m_{nk}(X^{\text{StatPhys}})] = \frac{1}{\alpha_{n1}} \mathbb{E}[X_{nk}^{\text{StatPhys}}] = \frac{1}{\alpha_{n1}} \frac{\alpha_{n1} c_k}{p_{nk} A} p_{nk} = \frac{c_k}{A}$$

$$= m_{nk}(X^{\text{Entropy}}) \quad \forall n \in \mathcal{N}, \forall k \in \mathcal{S}$$

and hence coincide with the portfolio weights using the Entropy method.

This implies that the quantities that depend on the network are given by

$$\mathbb{E}[R_{n1}(X^{\text{StatPhys}})] = \sum_{k=1}^{K} \mathbb{E}[m_{nk}(X^{\text{StatPhys}})] f_{k1} = R_{n1}(X^{\text{Entropy}}) = R^{\text{Entropy}},$$

$$\mathbb{E}[\gamma_{n1}(X^{\text{StatPhys}})] = \sum_{k=1}^{K} c_k l_k \mathbb{E}[m_{nk}(X^{\text{StatPhys}})] = \gamma_{n1}(X^{\text{Entropy}}) = \gamma^{\text{Entropy}},$$

for all $n \in \mathcal{N}$.

It follows directly that the expected direct vulnerability using the StatPhys method coincides with the direct vulnerability using the Entropy method, formally

$$\mathbb{E}[\mathcal{DV}^{X^{\text{StatPhys}}}(n)] = -\frac{\alpha_{n1}}{e_{n1}} \mathbb{E}[R_{n1}(X^{\text{StatPhys}})] = -\frac{\alpha_{n1}}{e_{n1}} R^{\text{Entropy}} = \mathcal{DV}^{X^{\text{Entropy}}}(n).$$

For the other fire sale measure, however, the expectation of the measure applied to the random matrix X^{StatPhys} does not generally coincide with the measure applied to the deterministic matrix X^{Entropy}. For example, for the systemicness, it follows from direct calculations that

$$\mathbb{E}[\mathcal{SYS}^{X^{\text{StatPhys}}}(n)] = \frac{-\alpha_{n1} b_{n1}}{\sum_{\nu=1}^{N} e_{\nu 1}} \mathbb{E}[\gamma_{n1}(X^{\text{StatPhys}}) R_{n1}(X^{\text{StatPhys}})],$$

where

$$\mathbb{E}[\gamma_{n1}(X^{\text{StatPhys}}) R_{n1}(X^{\text{StatPhys}})] = \gamma^{\text{Entropy}} R^{\text{Entropy}} + \frac{1}{\alpha_{n1}} \frac{\sum_{k=1}^{K} c_k^2 l_k f_{k1}}{\phi A^2}.$$

Since, $\frac{1}{\alpha_{n1}} \frac{\sum_{k=1}^{K} c_k^2 l_k f_{k1}}{\phi A^2} \leq 0$, this implies that the expected systemicness and the expected aggregate vulnerability under the StatPhys method is greater or equal than the corresponding quantities derived using the Entropy method. This is indeed what we find in Table 4.

A.2.2. Bayesian methods

The reconstruction method developed in Gandy and Veraart (2017) takes a Bayesian perspective. Gandy and Veraart (2017) specify a generative model for the network matrix and then condition on the observations, i.e., the row and column sums (and possibly additional known entries of the matrix). Hence, the network reconstruction is achieved through the posterior distribution in the Bayesian setting.

For the generative model several a-priori distributions have been considered in Gandy and Veraart (2017, 2019). We consider two special choices developed in these papers.

The model assumes a generalisation of the Erdős-Rényi random graph model, see Erdős and Rényi (1959), by assuming that directed edges from n to k are generated using independent Bernoulli trials with success probability p_{nk} and weights from an exponential distribution are

---

[24] This has also already been discussed in Gandy and Veraart (2019).

assigned to existing edges. Formally, the a-priori model assumes that for all $n \in \mathcal{N}$ and for all $k \in \mathcal{S}$

$$P(X_{nk} > 0) = p_{nk},$$

$$X_{nk} \mid X_{nk} > 0 \sim \text{Exp}(\lambda_{nk}),$$

where $p = (p_{nk}) \in [0,1]^{N \times K}$, $\lambda = (\lambda_{nk}) \in [0,\infty)^{N \times K}$.

We are then interested in the distribution of the random matrix X conditional on the given row and column sums. Since this distribution is not available in closed form, Gandy and Veraart (2017) have developed an MCMC sampler to generate samples from this distribution.

In the following, we will assume that all parameters of the exponential distributions governing the weights are identical, i.e., $\lambda_{nk} = \tilde{\lambda} \in [0,\infty)$ for all $n \in \mathcal{N}, k \in \mathcal{S}$.

We will now consider two different choices for $p = (p_{nk})$. First, we assume that all a-priori link existence probabilities are identical, i.e., we set $p_{nk} = \tilde{p} \in [0,1]$ for all $n \in \mathcal{N}, k \in \mathcal{S}$. We will refer to the Bayesian model with this a-priori assumption as the BayeER model (where ER stands for Erdős-Rényi). As discussed in Gandy and Veraart (2019), this model can be calibrated to a given network density by choosing appropriate values for $\tilde{p}$ and $\tilde{\lambda}$ and this is what we do in this paper.

Second, we assume that the a-priori link existence probabilities have the same structure as in the StatPhys model. In particular, they are given by $p_{nk} = \frac{\phi \alpha_{n1} c_k}{1+\phi \alpha_{n1} c_k}$ for all $n \in \mathcal{N}$, $k \in \mathcal{S}$. Here again α_{n1} and c_k represent the row and column sums and $\phi > 0$ is a constant used to calibrate the model. We refer to this Bayesian model as the BayeEF model (where EF stands for Empirical Fitness). This is (as the StatPhys model) a fitness model for the underlying network. Fitness network models assume that the link existence probability between a pair of nodes is a function of characteristics of the nodes, so-called fitnesses. In our setting, the row and column sums can be interpreted as fitnesses and the link existence probabilities are indeed functions of the row and column sums. Note, however, that the p_{nk} in the BayeEF model are a-priori link existence probabilities. They do usually not correspond to the posterior link existence probabilities. The StatPhys and the BayeEF are fundamentally different models despite having some similarities in the choice of model inputs. As shown in Gandy and Veraart (2019) also the BayeEF can be calibrated to a given network density and this is what we will do for this second type of Bayesian model as well. The calibration is described in detail in Gandy and Veraart (2019).

Appendix B. Proofs

Proof of Proposition 3.1. 1.-3. From the definition of the systemicness of bank $n \in \mathcal{N}$ in (4), it is clear that $\mathcal{SYS}(n) = \gamma_{n1} \frac{\alpha_{n1}}{\sum_{v=1}^{N} e_{v1}} \times b_{n1}(-R_{n1})$ depends on the network matrix X only via the two factors γ_{n1} and R_{n1}, since all other factors appearing in the formula are aggregate information that is available from the balance sheets of the banks.

This implies that also the aggregate vulnerability, as the sum of all individual systemicnesses, depends on the network matrix X only via γ_{n1} and R_{n1}, where $n \in \mathcal{N}$. We also see directly from the definition, that the direct vulnerability of a bank $n \in \mathcal{N}$ depends on the network matrix X only via the factor R_{n1}.

4. To see that γ_{n1} depends on the individual entries of X only via its nth row, we rewrite γ_{n1} given in (5) as follows

$$\gamma_{n1} = \sum_{k=1}^{K} \left(\sum_{p=1}^{N} \alpha_{p1} m_{pk} \right) l_k m_{nk} = \sum_{k=1}^{K} \left(\sum_{p=1}^{N} \alpha_{p1} \frac{X_{pk}}{\alpha_{p1}} \right) l_k \frac{X_{nk}}{\alpha_{n1}} = \sum_{k=1}^{K} c_k l_k \frac{X_{nk}}{\alpha_{n1}}. \tag{B.1}$$

Hence, we see that γ_{n1} only depends on $X_{n1}, \dots, X_{nK}$. To make the dependence of γ_{n1} on X explicit, we will sometimes write $\gamma_{n1}(X)$.

First, if one assumes a constant price impact, then formula (B.1) reduces to

$$\gamma_{n1} = \frac{l}{\alpha_{n1}} \sum_{k=1}^{K} c_k X_{nk}. \tag{B.2}$$

Indeed, γ_{n1} depends on the individual entries in the nth row of the matrix X since it is proportional to a capitalisation-weighted aggregate of the positions of node n in the K assets. If additionally their market capitalisation was identical, i.e., if $c_1 = \dots c_K = c$ (which would be unlikely in practice), then (B.2) would simplify even further to $\gamma_{n1} = \frac{l}{\alpha_{n1}} c \sum_{k=1}^{K} X_{nk} = lc$, which then no longer depends on the individual entries of X.

Second, if one assumes a capitalisation-dependent price impact, then $c_k l_k = \rho$ for all $k \in \mathcal{S}$ and hence $\gamma_{n1} = \frac{\rho}{\alpha_{n1}} \sum_{k=1}^{K} X_{nk} = \rho \quad \forall n \in \mathcal{N}$, which does not depend on X.

5. We find that

$$R_{n1} = \sum_{k=1}^{K} m_{nk} f_k = \sum_{k=1}^{K} \frac{X_{nk}}{\alpha_{n1}} f_k = \frac{1}{\alpha_{n1}} \sum_{k=1}^{K} X_{nk} f_k, \tag{B.3}$$

which again only depends on the matrix X via its nth row. To make the dependence of R_{n1} on X explicit, we will sometimes write $R_{n1}(X)$.

First, if we consider an all asset shock with $f_1 = \dots f_K = f$, expression (B.3) simplifies to $R_{n1} = \frac{f}{\alpha_{n1}} \sum_{k=1}^{K} X_{nk} = f$ and hence does not depend on the matrix X.

Second, we consider a shock that only affects $\tilde{K} < K$ assets with indices in $\mathcal{I}^{\tilde{K}}$. Then,

$$R_{n1} = \sum_{k=1}^{K} m_{nk} f_k = \sum_{k=1}^{K} \frac{X_{nk}}{\alpha_{n1}} f_k = \frac{1}{\alpha_{n1}} \sum_{k=1}^{K} X_{nk} f_k = \frac{1}{\alpha_{n1}} \sum_{k \in \mathcal{I}^{\tilde{K}}} X_{nk} f_k.$$

Hence, R_{n1} only depends on the columns with indices in $\mathcal{I}^{\tilde{K}}$ within the nth row, but not the full nth row of X.

Since $\mathcal{DV}(n)$ depends on X only via R_{n1} the results for $\mathcal{DV}(n)$ follow directly from the results on R_{n1}. ☐

Proof of Corollary 3.2. 1. Under an all asset shock and a capitalisation-dependent price impact, we know from the proof of Proposition 3.1 that $R_{n1} = f$ and $\gamma_{n1} = \rho$. Hence,

$$\mathcal{SYS}(n) = -f \rho \frac{\alpha_{n1}}{\sum_{v=1}^{N} e_{v1}} b_{n1},$$

$$\mathcal{AV} = \sum_{n=1}^{N} \mathcal{SYS}(n) = \frac{-f\rho}{\sum_{v=1}^{N} e_{v1}} \sum_{n=1}^{N} \alpha_{n1} b_{n1},$$

$$\mathcal{DV}(n) = \frac{-f \alpha_{n1}}{e_{n1}},$$

which do not depend on the individual entries of X.

2. This statement follows directly from Proposition 3.1 and the analytical formulae of γ_{n1} and R_{n1} provided in its proof. ☐

Appendix C. Additional empirical results and sensitivity analysis

C.1. Observed and reconstructed asset holding matrices

To provide some intuition on the empirical asset holding matrix and the performance of different reconstruction methods, we illustrate their performance when applied to the EBA data from 2016. Fig. C.5 shows a heatmap of the true asset holdings matrix X (top left) and five heatmaps corresponding to reconstructed asset holding matrices that only used partial information. For methods that generate a sample of matrices, i.e., the StatPhys method and the Bayesian methods we only show one realisation of a reconstructed asset holding matrix.

Fig. C.5 shows that the matrix obtained using the Entropy method corresponds to a network in which all institutions hold positions in all but one asset. This one asset has a market capitalisation of 0 and corresponds to Liechtenstein sovereign loans.

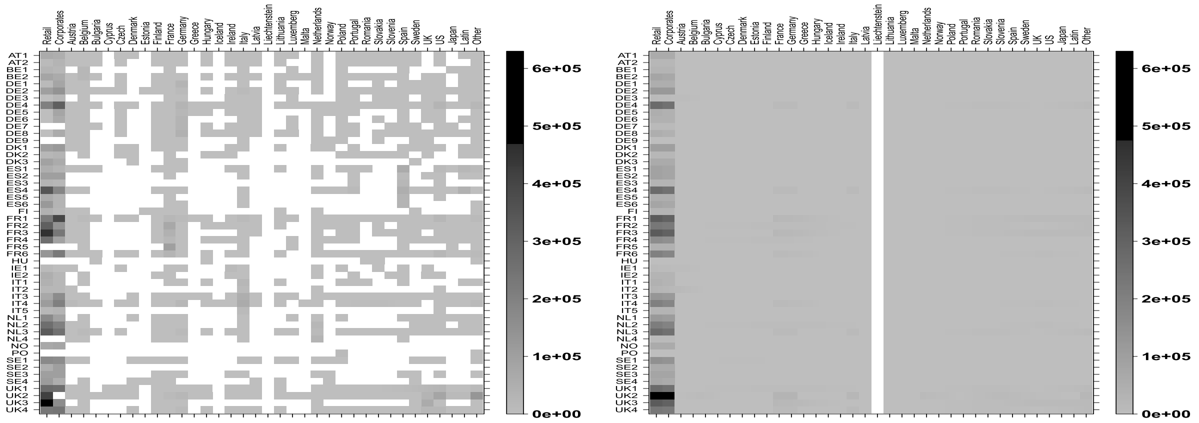

(a) Observed asset holdings matrix. (b) Reconstruction based on Entropy method.

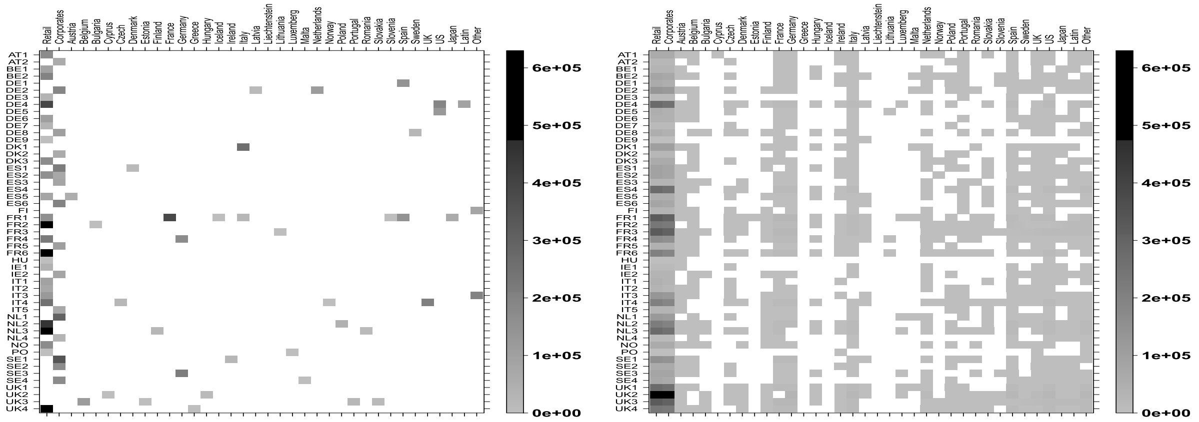

(c) Reconstruction based on MinDen method. (d) One reconstruction based on StatPhys method.

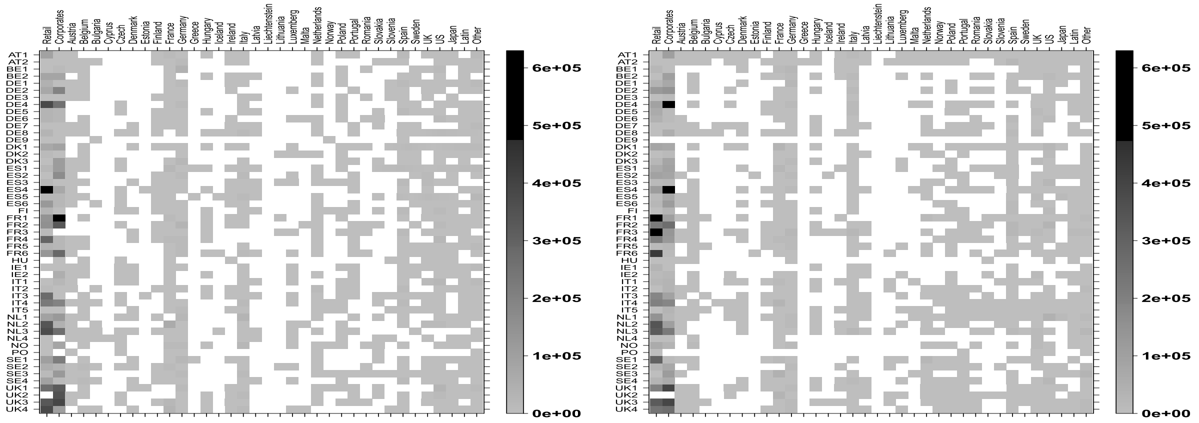

(e) One reconstruction based on BayeER method. (f) One reconstruction based on BayeEF method.

Fig. C.5. Asset holdings matrix for the true matrix (top left) and five reconstructed networks based on different methods for the EBA 2016 data.

For the reconstructed matrix based on the MinDen method, the asset holdings appear scattered where the largest assets holdings are in corporate, retail, German, US, and other sovereign assets. The MinDen matrix does assign zero weights to some of the largest positions observed in the true network i.e., several UK banks hold large positions in retail assets in 2016, but the corresponding entries in the reconstructed matrix based on the MinDen methods are zero.

The sample matrix generated by the StatPhys method shows that according to this reconstruction, all banks invest in the two asset classes corporate and retail (the lower two rows). The weights are consistent with the corresponding two rows in the matrix obtained from using the Entropy method. According to the reconstruction based on the Entropy method and this one sample from the StatPhys method, the bank with label UK1 has the largest holdings in the two asset classes corporate and retail. In contrast, to the Entropy method, the reconstructed matrix based on the StatPhys method is much sparser - it has been calibrated to match the density of the true network. When looking at the samples generated by the Bayesian method we observe that the overall density of the network matches the density of the true network, as was the case for the StatPhys method, since these methods are flexible enough

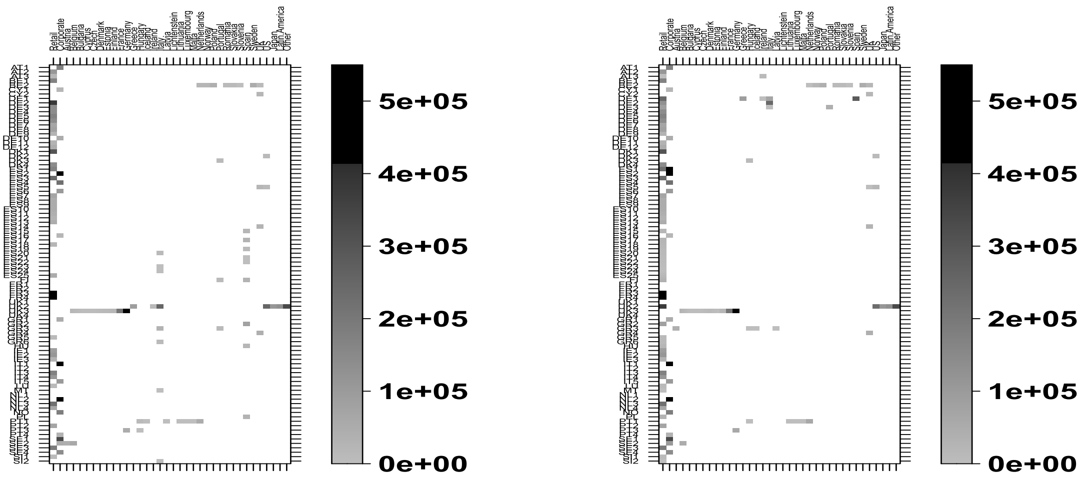

(a) Min AV under the GIIPS scenario. (b) Max AV under the GIIPS scenario.

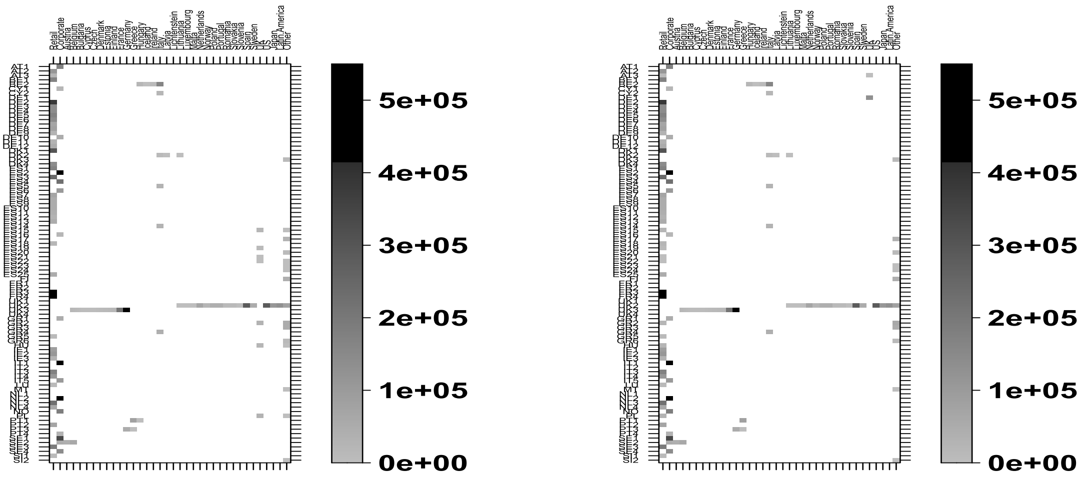

(c) Min AV under the Bad Brexit scenario. (d) Max AV under the Bad Brexit scenario.

Fig. C.6. Examples of asset holdings matrices that maximise or minimise the aggregate vulnerability while respecting the non-negativity and marginal sums of the true matrix. We use the 2011 EBA data for the GIIPS and Bad Brexit scenarios. The aggregate vulnerability are as follows a) 172%, b) 373%, c) 49%, and d) 144%.

that they can easily be calibrated to a given density. Throughout our empirical analysis, we calibrate the StatPhys and the Bayesian methods to the true density of the network unless stated otherwise. The density remains almost the same in both years (0.44 in 2011 and 0.48 in 2016), so our sampling-based methods are calibrated such that almost half the entries of the asset holding matrices are filled.

Furthermore, we see that the reconstructed samples from the Bayesian methods assign weights that are very different from weights obtained by the Entropy or the StatPhys method. In particular, we do observe several high weights and also some zero weights within the lower two rows that represent the holdings in corporate and retail assets. This is not surprising given the greater flexibility of the Bayesian method when it comes to modelling the weight and not just the existence of edges compared to the StatPhys method.

C.2. Computing the maximum and minimum aggregate vulnerability for given row and column sums

In our partial information setting, we assume that only the row and column sums of the asset holding matrix are given. Throughout the paper, we study various network reconstruction methods that reconstruct the asset holding matrix from this partial information. These reconstructed networks can then be plugged into any measure of fire sale risk of interest, such as the aggregate vulnerability, the systemicness, the direct and the indirect vulnerability.

For a fixed measure of fire sale risk, however, one can also try to find an asset holding matrix that maximises (or minimise) this measure over all matrices that satisfy the given constraints on the row and column sums. This is what we do next for the aggregate vulnerability. This optimisation approach will be useful as a benchmark.

We consider the following optimisation problem for finding the maximum aggregate vulnerability:

$$\begin{aligned} &\max_{X} \quad \mathcal{AV}(X), \\ \text{subject to: } &\alpha_{n1} = \sum_{k=1}^{K} X_{nk} \quad \forall n \in \{1,\dots,N\}, \\ &c_k = \sum_{n=1}^{N} X_{nk} \quad \forall k \in \{1,\dots,K\}, \\ &X_{nk} \geq 0 \quad \forall n \in \{1,\dots,N\}, \quad \forall k \in \{1,\dots,K\}. \end{aligned} \tag{C.1}$$

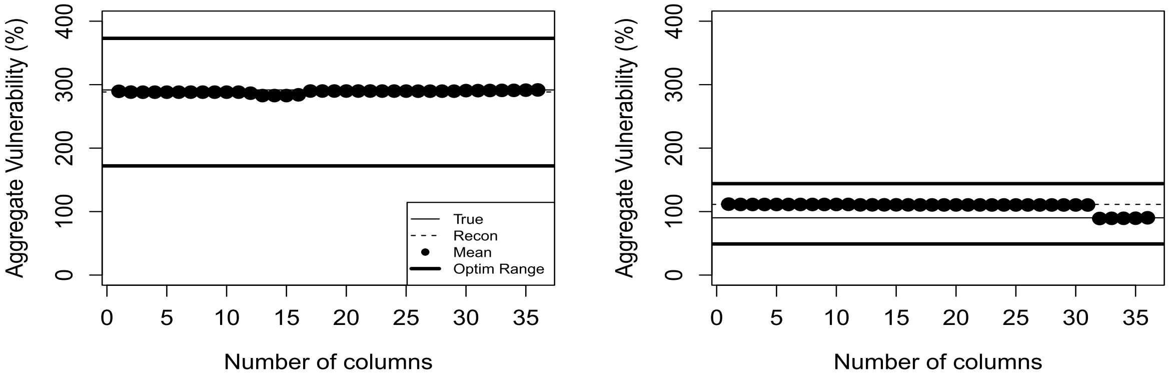

(a) Entropy under GIIPS 2011.

(b) Entropy under Bad Brexit 2011.

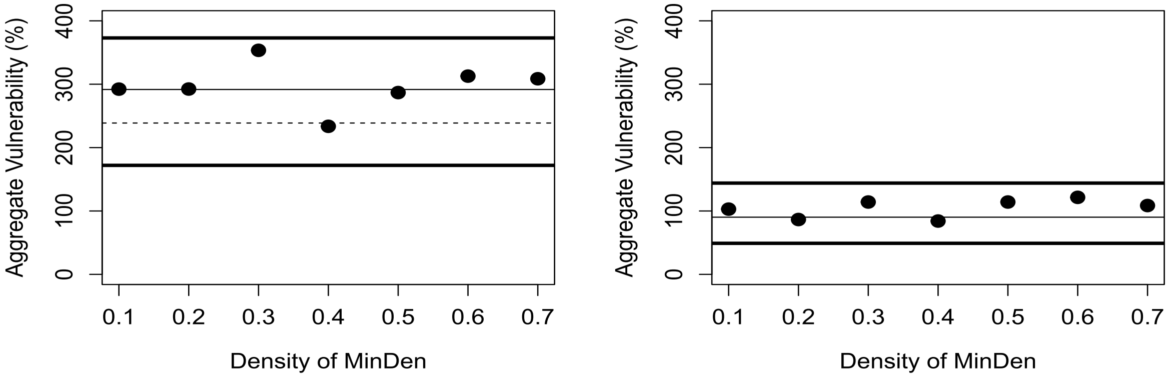

(c) MinDen under GIIPS 2011.

(d) MinDen under Bad Brexit 2011.

Fig. C.7. Aggregate vulnerabilities as a function of the number of known columns in the target matrix of the Entropy method (top), and aggregate vulnerabilities as a function of the network density for the MinDen method (bottom). This is for the 2011 data and a capitalisation-dependent price impact.

The corresponding optimisation problem that determines the minimum aggregate vulnerability can be defined in exactly the same way by minimising the aggregate vulnerability rather than maximising it.

Fig. C.6 shows examples of asset holding matrices that correspond to the minimum or maximum aggregate vulnerability for different shocks. We find that the matrices are generally very sparse. We observe large positions in corporate and retail assets, which correspond to 80% of the total value of asset holdings across all banks. It is striking to see how similar the two matrices are that correspond to the minimum and the maximum aggregate vulnerability for a given stress scenario. This shows the large influence of a small number of positions on aggregate vulnerability. It is possible that there are other matrices with different levels of sparsity that result in similarly small or large aggregate vulnerabilities.

C.3. Additional sensitivity analysis

C.3.1. Network reconstruction for different densities

In the following, we provide further details on the sensitivity of different network reconstruction methods with respect to the assumed density of the network.

In Section 3.5 we have already discussed how the aggregate vulnerability computed using the StatPhys, BayeER, and BayeEF network reconstruction methods depends on the choice of the target density of the network for the 2011 data and a capitalisation-dependent price impact, see Fig. 1. In particular, we find that the aggregate vulnerability can be estimated reasonably precisely even if the true density of the network is not available. Further analysis of these sensitivities for constant price impact and for the data from 2016 confirm these conclusions. We do not report the details here.

Next, we analyse how additional information can be included in the Entropy and MinDen methods.

For the Entropy method, we cannot just assume a target density but we will need to provide a suitable target matrix $\tilde{X}$ instead. We do this by replacing some columns in the target matrix with the true asset holding matrix. In particular, first, we assume that the first column of the target matrix consists of the true asset holdings and the remaining entries correspond to those in the Entropy method matrix i.e., $\tilde{X}_{n1} = X_{n1} \quad \forall n \in \mathcal{N}$ and $\tilde{X}_{nk} = X_{nk}^{\text{Entropy}}, \quad k \in [2, \dots, K]$. The column sums remain the same but the row sums are no longer consistent with the partial information. We, therefore, re-balance the matrix such that marginal sums are equal to the true matrix. We repeat this process for each column sequentially until the target matrix consists only of the true entries. We do this column by column.

Our results in Fig. C.7 show that the aggregate vulnerabilities are similar under the additional information for the Entropy method. For several points, incorporating additional information can result in a worse performance of the aggregate vulnerability. Although more information is known about the true matrix, the proportional scaling from the re-balancing method alters other entries. This leads to changes in other assets with high influence, for example, changes in position in UK assets within the Bad Brexit scenario. Only when the information about the shocked asset is included, we observe that the estimated aggregate vulnerability becomes closer to the true one.

For the MinDen method, it is possible to consider a generalisation, see Anand et al. (2015) for details, that can be calibrated to a target

Table C.7
The table presents average fire sale risk measures (averaged over the banks and additionally averaged over the reconstructed samples) for the 2011 EBA data for two different shock scenarios for the true matrix and for reconstructed matrices that were reconstructed from noisy observations of the row and column sums.

Capitalisation-dependent price impact ($l_k = \rho/c_k\ \forall k$)											
Matrix	True	MinDen		Entropy		StatPhy		BayeER		BayeEF	
σ	0	100	1000	100	1000	100	1000	100	1000	100	1000
GIIPS (%)											
$\mathcal{DV}11$	15.58	32.74	2.60	7.81	7.80	7.80	7.81	20.20	20.17	**17.69**	**17.53**
	-	-	-	-	-	*(0.54)*	*(0.53)*	*(3.09)*	*(3.05)*	*(3.01)*	*(3.01)*
$\mathcal{IV}11$	460.79	1075.10	293.85	416.77	416.17	**417.42**	**416.99**	611.87	610.83	544.72	542.05
	-	-	-	-	-	*(10.66)*	*(10.71)*	*(38.63)*	*(38.49)*	*(30.36)*	*(30.46)*
$\mathcal{AV}11$	291.70	289.01	316.26	288.21	287.85	288.98	288.42	**292.41**	**291.48**	294.22	293.44
	-	-	-	-	-	*(6.02)*	*(6.15)*	*(5.27)*	*(5.28)*	*(5.01)*	*(5.03)*
Bad Brexit (%)											
$\mathcal{DV}11$	1.47	**0.85**	4.03	3.01	3.01	3.02	**3.01**	8.37	8.36	6.47	8.36
	-	-	-	-	-	*(0.47)*	*(0.46)*	*(02.50)*	*(2.54)*	*(2.43)*	*(2.39)*
$\mathcal{IV}11$	120.19	175.99	225.61	**160.90**	**160.71**	161.18	160.91	239.96	239.41	207.53	206.42
	-	-	-	-	-	*(7.31)*	*(7.37)*	*(22.66)*	*(22.96)*	*(17.30)*	*(17.04)*
$\mathcal{AV}11$	90.23	144.05	119.90	**111.27**	**111.16**	111.59	111.30	112.14	111.91	113.67	113.47
	-	-	-	-	-	*(4.77)*	*(4.83)*	*(3.95)*	*(4.04)*	*(3.90)*	*(3.92)*
Constant price impact ($l_k = 5 \times 10^{-13}\ \forall k$)											
GIIPS (%)											
$\mathcal{DV}11$	Same results as for capitalisation-dependent price impact										
$\mathcal{IV}11$	506.76	312.90	321.79	523.55	**522.83**	**523.48**	522.85	275.16	274.39	325.14	325.60
	-	-	-	-	-	*(13.51)*	*(13.50)*	*(16.68)*	*(16.81)*	*(17.83)*	*(17.85)*
$\mathcal{AV}11$	357.49	226.56	204.36	**362.05**	**361.62**	362.43	361.73	275.50	274.60	293.51	293.20
	-	-	-	-	-	*(7.24)*	*(7.41)*	*(11.35)*	*(11.42)*	*(10.11)*	*(10.21)*
Bad Brexit (%)											
$\mathcal{DV}11$	Same results as for capitalisation-dependent price impact										
$\mathcal{IV}11$	155.05	251.16	119.47	202.13	201.90	202.13	201.76	103.91	103.81	**126.40**	**126.83**
	-	-	-	-	-	*(9.24)*	*(9.30)*	*(8.87)*	*(8.98)*	*(9.53)*	*(9.59)*
$\mathcal{AV}11$	109.02	165.48	67.01	139.78	139.65	139.94	139.59	**104.08**	**103.99**	114.60	114.72
	-	-	-	-	-	*(5.95)*	*(6.03)*	*(8.09)*	*(8.19)*	*(7.38)*	*(7.39)*
Bold	-	1	0	3	4	2	2	2	2	2	2

density. We find that the aggregate vulnerabilities, computed from the MinDen method that have been calibrated to different densities, can vary and different densities can lead to similar aggregate vulnerabilities.

We have also conducted the same sensitivity checks for a constant price impact and also for the 2016 data and come to the same conclusions, therefore we do not report them here.

C.3.2. Network reconstruction for noisy observations

Finally, we investigate how sensitive our results are, if the row and column sums of the asset holding matrix are not observed directly but with noise. To do so we add a noise term to the row and column sums of the true asset holding matrix, i.e., we consider the new row and column sums

$$\alpha_{n1}^{\text{noise}} = \alpha_{n1} + \epsilon_n^{(\alpha)}, \quad \forall n \in \mathcal{N},$$

$$c_k^{\text{noise}} = c_k + \epsilon_k^{(c)}, \quad \forall k \in \mathcal{S},$$

where $\epsilon_1^{(\alpha)}, \dots, \epsilon_N^{(\alpha)}$ and $\epsilon_1^{(c)}, \dots, \epsilon_K^{(c)}$ are i.i.d. normally distributed random variables with mean 0 and variance σ^2. We chose a realisation of the noise in which all new row and column sums are non-negative. We consider two different choices of the parameter $\sigma \in \{100, 1000\}$ (million EUR). Finally, we normalise row and column sums of the data with noise such that the new row sums $\tilde{\alpha}_{n1}^{\text{noise}}$ and column sums $\tilde{c}_k^{\text{noise}}$ with noise satisfy

$$\sum_{n=1}^{N} \tilde{\alpha}_{n1}^{\text{noise}} = \sum_{k=1}^{K} \tilde{c}_k^{\text{noise}} = \sum_{n=1}^{N} \sum_{k=1}^{K} X_{nk}.$$

For the network reconstruction with noise, we use the normalised row and column sums with noise as the available partial information.

Table C.7 reports the results. Overall, we observe only a small or no deviation from the results without noise for all reconstruction methods except the MinDen method. The results of the MinDen method are sensitive to noisy observation. The mean equity losses for reconstructed matrices under 1 bn standard deviation of the noise are further away than for 100mn but with a similar standard deviation for the sampling methods. This shows that our results are robust under noise and across different fire-sales measures. Overall we find that the addition of noise does not lead to any different conclusion in terms of the relative ranking of the different network reconstruction methods, see Table 4 for comparison.

References

Adrian, T., Shin, H.S., 2010. Liquidity and leverage. J. Financ. Intermed. 19, 418–437.

Amini, H., Filipović, D., Minca, A., 2016. Uniqueness of equilibrium in a payment system with liquidation costs. Oper. Res. Lett. 44, 1–5.

Anand, K., Craig, B., Von Peter, G., 2015. Filling in the blanks: network structure and interbank contagion. Quant. Finance 15, 625–636.

Anand, K., van Lelyveld, I., Banai, Á., Friedrich, S., Garratt, R., Hałaj, G., Fique, J., Hansen, I., Jaramillo, S.M., Lee, H., et al., 2018. The missing links: a global study on uncovering financial network structures from partial data. J. Financ. Stab. 35, 107–119.

Banerjee, T., Feinstein, Z., 2021. Price mediated contagion through capital ratio requirements with vwap liquidation prices. Eur. J. Oper. Res. 295, 1147–1160.

Bichuch, M., Feinstein, Z., 2022. Endogenous inverse demand functions. Oper. Res. 70, 2702–2714.

Braouezec, Y., Wagalath, L., 2019. Strategic fire-sales and price-mediated contagion in the banking system. Eur. J. Oper. Res. 274, 1180–1197.

Brazier, A., 2017. How to: Macropru. 5 principles for macroprudential policy. Speech given by Alex Brazier, Executive Director for Financial Stability Strategy and Risk. Member of the Financial Policy Committee, Bank of England, at the London School of Economics, Financial Regulation Seminar, Monday 13 February 2017. Available at https://www.bankofengland.co.uk/speech/2017/how-to-macropru-5-principles-for-macroprudential-policy.

Caballero, R.J., Simsek, A., 2013. Fire sales in a model of complexity. J. Finance 68, 2549–2587.

Capponi, A., Larsson, M., 2015. Price contagion through balance sheet linkages. Rev. Asset Pricing Stud. 5, 227–253.

Capponi, A., Weber, M., 2021. Systemic portfolio diversification. Available at SSRN 3345399.

Cifuentes, R., Ferrucci, G., Shin, H.S., 2005. Liquidity risk and contagion. J. Eur. Econ. Assoc. 3, 556–566.

Cimini, G., Squartini, T., Garlaschelli, D., Gabrielli, A., 2015. Systemic risk analysis on reconstructed economic and financial networks. Sci. Rep. 5, 15758.

Cont, R., Schaanning, E., 2017. Fire sales, indirect contagion and systemic stress testing. Norges Bank Working Paper. 2017.

Cont, R., Schaanning, E., 2019. Monitoring indirect contagion. J. Bank. Finance 104, 85–102.

Cont, R., Wagalath, L., 2013. Running for the exit: distressed selling and endogenous correlation in financial markets. Math. Finance 23, 718–741.

Detering, N., Meyer-Brandis, T., Panagiotou, K., Ritter, D., 2022. Suffocating fire sales. SIAM J. Financ. Math. 13, 70–108.

Di Gangi, D., Lillo, F., Pirino, D., 2018. Assessing systemic risk due to fire sales spillover through maximum entropy network reconstruction. J. Econ. Dyn. Control 94, 117–141.

Duarte, F., Eisenbach, T.M., 2021. Fire-sale spillovers and systemic risk. J. Finance 76, 1251–1294.

Erdős, P., Rényi, A., 1959. On random graphs, I. Publ. Math. (Debr.) 6, 290–297.

Feinstein, Z., 2017. Financial contagion and asset liquidation strategies. Oper. Res. Lett. 45, 109–114.

Feinstein, Z., 2020. Capital regulation under price impacts and dynamic financial contagion. Eur. J. Oper. Res. 281, 449–463.

Gandy, A., Veraart, L.A.M., 2017. A Bayesian methodology for systemic risk assessment in financial networks. Manag. Sci. 63, 4428–4446.

Gandy, A., Veraart, L.A.M., 2019. Adjustable network reconstruction with applications to CDS exposures. J. Multivar. Anal. 172, 193–209.

Greenwood, R., Landier, A., Thesmar, D., 2015. Vulnerable banks. J. Financ. Econ. 115, 471–485.

Huang, X., Vodenska, I., Havlin, S., Stanley, H.E., 2013. Cascading failures in bi-partite graphs: model for systemic risk propagation. Sci. Rep. 3, 1219.

Lebacher, M., Cook, S., Klein, N., Kauermann, G., 2019. In search of lost edges: a case study on reconstructing financial networks. arXiv preprint. arXiv:1909.01274.

Ramadiah, A., Caccioli, F., Fricke, D., 2020. Reconstructing and stress testing credit networks. J. Econ. Dyn. Control 111, 103817.

Sharpe, W.F., 1964. Capital asset prices: a theory of market equilibrium under conditions of risk. J. Finance 19, 425–442.

Shleifer, A., Vishny, R.W., 2010. Asset fire sales and credit easing. Am. Econ. Rev. 100, 46–50.

Squartini, T., Almog, A., Caldarelli, G., Van Lelyveld, I., Garlaschelli, D., Cimini, G., 2017. Enhanced capital-asset pricing model for the reconstruction of bipartite financial networks. Phys. Rev. E 96, 032315.

Upper, C., Worms, A., 2004. Estimating bilateral exposures in the German interbank market: is there a danger of contagion? Eur. Econ. Rev. 48, 827–849.