EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Garbrand Wiersema, Alissa M. Kleinnijenhuis,
Esti Kemp, Thom Wetzer

Higher-order exposures

No 3091

Abstract

Traditional exposure measures focus on direct exposures to evaluate the losses an institution is exposed to upon the default of a counterparty. Since the Global Financial Crisis of 2007-2008, the importance of indirect exposures via common asset holdings is increasingly recognized. Yet direct and indirect exposures do not to capture the losses that result from shock propagation and amplification following the counterparty's default. In this paper, we introduce the concept of "higher-order exposures" to refer to these spill-over losses and propose a way to formalize and quantify these. Using granular data on the South African banking and investment fund sectors and a contagion model that captures the most commonly studied contagion channels and their interactions, we demonstrate that higher-order exposures make up a significant part of exposures – particularly during times of financial distress when exposures matter most. We also show that higher-order exposures cannot simply be extrapolated from direct or indirect exposures, since they depend strongly on the network structure and the robustness of individual institutions. Our findings suggest that exposures should be properly understood as consisting of direct, indirect and higher-order exposures in the design and calibration of those tools in the regulators' arsenal where exposures matter – including large exposure limits, capital requirement calibration, stress test design and resolution. Failure to do so may result in both lax *ex-ante* regulation and ill-informed *ex-post* handling of financial crises.

Keywords

Financial Contagion, Systemic Risk, System-Wide Stress Test, Financial Exposures, Non-Bank Financial Institutions (NBFIs)

JEL Classification: G01, G17, G21, G23, G28

1 Non-Technical Summary

The global financial crisis of 2007-2008 revealed how financial institutions are deeply interconnected, and how risks can spread across the entire financial system, sometimes in ways not immediately visible from traditional measures. When a financial institution defaults, the losses it causes can ripple through the system, affecting others beyond those with direct ties to the failing institution.

Traditionally, exposures between financial institutions have been measured by direct exposures—the immediate financial links such as loans or investments one institution has with another. More recently, there has been growing interest in indirect exposures—losses that arise because different institutions hold overlapping assets. For example, if two banks hold large amounts of the same securities, a forced sale by one can depress prices and cause mark-to-market losses to the other. While these direct and indirect exposures provide important insights, they still do not capture the full picture.

This paper introduces a new concept called higher-order exposures, which refer to losses caused by the spread and amplification of financial shocks beyond direct and indirect exposures. When one institution fails, the resulting shock can propagate through multiple intermediaries before impacting another institution. These spill-over effects can lead to losses that are not accounted for in traditional exposure measures. Ignoring higher-order exposures hence risks underestimating potential losses and systemic vulnerabilities.

Using detailed data from South Africa's financial banking and investment fund sectors and a contagion model that simulates various channels of shock propagation, we find that higher-order exposures represent a significant portion of total exposure—especially during periods of financial stress when risks matter most. Importantly, higher-order exposures cannot simply be extrapolated from direct or indirect exposures. Their magnitude and distribution depends heavily on the structure of the financial network and the resilience of individual institutions. Therefore, higher-order exposures require explicit modeling to be accurately estimated.

The findings have direct implications for financial regulation and policy. Many regulatory tools, such as capital requirements, large exposure limits, resolution mechanisms, and stress testing frameworks, currently rely on exposure measures that exclude higher-order effects. This oversight can lead to insufficient safeguards and ill-prepared responses in times of crisis. To better protect financial stability, regulators should incorporate higher-order exposures into their risk assessments and regulatory frameworks. This requires collecting more detailed, granular data across a wide range of financial institutions and using models that capture the complex, multi-layered network of financial interconnections. While this study focuses on South Africa, the concept and its policy relevance apply broadly to other financial systems, including the euro area.

2 Higher-Order Exposures

The Great Financial Crisis of 2007-2008 served as a stark illustration of systemic risk materializing. Now, escalating geopolitical tensions are raising concerns about the potential for another systemic crisis to emerge. Transmission channels through which distress spreads across the financial system may generate exposures that are not visible on institutions' balance sheets. In this paper, we introduce the concept of Higher-Order Exposures, which expands the traditional understanding of exposure to capture the ever-present systemic risk arising from endogenous shock propagation.

An institution's exposure to a counterparty is typically understood to be the maximum loss the institution stands to suffer when that counterparty defaults. The most commonly studied exposures to a counterparty are *direct exposures*. Direct exposures arise from loans, bonds, stocks and other investments in a counterparty, which depreciate to their recovery value upon the default of the counterparty.[1] More recently, *indirect exposures* have been captured in contagion models studying spill-over losses resulting from fire-sales[2]. Indirect exposures arise when the portfolios of tradable securities of two or more institutions overlap. When one of these institutions liquidates (part of) its portfolio, the market price of the securities falls, causing mark-to-market losses to the other institutions.[3]

In this paper, we show that institutions' exposures are neither limited to direct exposures as is traditionally assumed, nor to direct and indirect exposures as increasingly recognized; assuming either is likely to dangerously underestimate exposures. When institution j defaults, other institutions in the system may propagate the losses caused by the default. The propagation of losses by financial institutions has been widely studied and is referred to as contagion.[4] Consequently, when j defaults institution i may suffer losses on other assets than its investments in and portfolio overlap with j. Put differently, these losses are not captured by i's direct and indirect exposures to j. We refer to the losses that were caused by the default of j, and propagated by at least one intermediate institution k before being suffered by i, as i's "higher-order exposure" to j. The exposure that i has to j (i.e. the total losses it would suffer when j defaults) thus has a higher-order component:

$$\mathit{Exposure}_{ij} \overset{\text{def}}{=} \mathit{Direct\ Exposure}_{ij} + \mathit{Indirect\ Exposure}_{ij} + \mathit{Higher\text{-}Order\ Exposure}_{ij} \qquad (1)$$

By contrast, we refer to the direct and indirect component of the exposure as the "first-order exposure".

We propose that exposures should be properly understood as consisting of direct, indirect, and higher-order exposures. To highlight the dangers of ignoring higher-order exposures, we introduce the "higher-order share of exposure" (HSE), which expresses the share of an institu-

[1] See e.g. Kraft and Steffensen [2007], Jorion and Zhang [2009], Cont et al. [2010], Glasserman and Young [2016].

[2] Losses due to overlapping portfolios have been studied in e.g.Gai and Kapadia [2010], Cont and Wagalath [2013], Caballero and Simsek [2013], Elliott et al. [2014], Glasserman and Young [2015], Cont and Schaanning [2019]. (However, these papers do not consider implications for measuring exposures).

[3] See e.g. Coval and Stafford [2007], Krishnamurthy [2010], Shleifer and Vishny [2011], Caccioli et al. [2014], Greenwood et al. [2015], Cont and Schaanning [2017].

[4] See e.g. Allen and Gale [2000], Gai and Kapadia [2010], Elliott et al. [2014], Glasserman and Young [2015].

tion's exposure made up by higher-order exposures. The HSE of the exposure of i to j is given by

$$HSE_{ij} \stackrel{\text{def}}{=} \frac{\textit{Higher-Order Exposure}_{ij}}{\textit{Exposure}_{ij}}. \tag{2}$$

The HSE thus expresses the fraction of an exposure that is overlooked when only considering first-order (i.e. direct and indirect) exposures. When the HSE is high, conventional methods will substantially underestimate exposures. That adversely affects the efficacy of institutions' efforts to limit their exposures, and of prudential policy designed to prevent or mitigate financial crises.

Contribution. In this paper, we introduce the concept of higher-order exposures and propose a way to formalize and quantify them. To measure the exposure of an institution i to another institution j, we study the losses that i is exposed to when j defaults and j's tradable securities are liquidated. i's direct exposure to j is calculated as i's Loss Given Default (LGD) on its investments in j (i.e. the investments' value minus expected recoveries), and i's indirect exposure to j is calculated as the mark-to-market losses (given by the liquidity-weighted portfolio overlap between i and j[5]) that i suffers when j defaults and liquidates its securities. Any additional losses that i suffers, due to shock propagation following j's default, constitute higher-order exposures.

To simulate how losses propagate throughout the system following j's default, we model the most salient contagion channels and their interactions; overlapping portfolio contagion, counterparty risk contagion and shareholder contagion (although in principle higher-order exposures can be evaluated conditional on any set of channels). Shareholder contagion simply propagates an institution's losses to its shareholders as those are the ultimate owners of the institution's assets. Furthermore, overlapping portfolio contagion captures the propagation of losses through portfolio overlaps when institutions' tradable securities are liquidated upon default[6]. This liquidation depresses the market price of tradable securities which, through the marked-to-market accounting required by modern accounting practices, causes losses to other holders of the tradable securities. Finally, counterparty risk contagion reflects that an institution's probability of default rises when it suffers losses, which depresses the expected value of investments in the institution (and, in particular, reduces the value of the investments to their recovery value when the losses cause the institution to default)[7]. This propagates the losses suffered by the institution to its investors. Any contagion channel may further propagate losses caused by another channel, and failing to capture these interactions may underestimate the amplification of asset losses (see e.g. Kok and Montagna [2013], Wiersema et al. [2023]). By capturing the interactions between contagion channels, our model provides a comprehensive measure of higher-order exposures.

We demonstrate our approach using South-Africa as a case study, given the high degree of interconnectedness in the country's (relatively well-developed) financial system and because we have access to a unique and highly detailed dataset on the interconnections between various types of South-African financial institutions. We measure exposures across multiple sectors of the financial system as the data not only cover the banking sector, but, uniquely, also various

[5]See e.g. Cont and Schaanning [2017, 2019]

[6]See e.g. Cifuentes et al. [2005], Adrian and Shin [2010], Cont and Wagalath [2013], Greenwood et al. [2015], Clerc et al. [2016], Cont and Schaanning [2019], Duarte and Eisenbach [2021].

[7]See e.g. Bardoscia et al. [2017], Wiersema et al. [2023]

types of investment funds (i.e. money market funds, fund-of-funds and other funds). Funds' assets are sourced from Morningstar, which give any fund's investments per instrument type in individual counterparties (individual banks, funds, etc.). Banks' assets are given by the BA900 data published by the South African Reserve Bank (SARB [2016b]), which, together with the Morningstar data, allow us to infer likely interconnections among banks. These data inform us on the exposures that arise from deposits, bonds, shares and money market instruments owned by South African banks and funds. In particular, our model captures how a single tradable security generates exposure both through counterparty risk to the security's issuer as well as through overlapping portfolio contagion when the security is fire-sold, which is a substantial improvement over existing models of overlapping portfolio contagion[8]. We capture the data in a multi-layered network model of the South-African financial system, which is a natural representation of the financial system in the context of studying systemic risk (Allen and Babus [2009]).

Our results show that higher-order exposures make up the vast majority of the exposures in a significant part of the system. We focus on exposures to the South Africa's six largest banks, which form the core of the South-African financial system, and hold over 90% of banking sector assets (SARB [2017a]). We estimate that the average HSE of banks and funds to one of these six banks South-African banks is between 45-55%, so about half of the exposures are overlooked when ignoring high-order exposures. We observe substantial heterogeneity across sectors in terms of their HSE to these banks. In particular, fund-of-fund's HSEs are close to 100% as, due to their particular strategy of investing in other funds, fund-of-funds have virtually no direct or indirect exposures to the banks but substantial higher-order exposures. Hence, higher-order exposures cannot be extrapolated from direct and/or indirect exposures as they depend strongly on the network structure and the robustness of individual institutions.

We find that higher-order exposures are particularly pronounced during times of financial crises. Unfortunately, this is exactly when exposures matter most, as defaults are more likely occur and hence exposures are more likely to materialize into losses. Higher-order exposures are increased during crises by falling market liquidities, which exacerbates the overlapping portfolio contagion channel, and by institutions becoming less resilient to shocks, which increases counterparty risk contagion. We formulate a severe stress scenario, which depresses market liquidities and worsens institutions' resilience (SARB, 2016a), and find that the average HSE of banks and funds to the failure of one of the large banks may reach up to 85% (while not exceeding 55% in benign times). These results suggest that higher-order exposures should inform the design and calibration of those tools in the regulators' arsenal where exposures matter, including large exposure limits, capital requirements, stress testing and resolution. Failure to do so may result in both excessively lax *ex-ante* regulation and ill-informed *ex-post* handling of financial crises.

Link to the Literature. We contribute to the literature on systemic risk[9] and financial crises[10], and to the financial exposures literature in particular, which includes a large body

[8] see e.g. Cifuentes et al. [2005], Adrian and Shin [2010], Cont and Wagalath [2013], Greenwood et al. [2015], Clerc et al. [2016], Cont and Schaanning [2019], Duarte and Eisenbach [2021]

[9] See e.g. Eisenberg and Noe [2001], Haldane and May [2011], Adrian and Brunnermeier [2011], Acharya et al. [2017], Brownlees and Engle [2017].

[10] See e.g. Bernanke [1983], Reinhart and Rogoff [2008], Brunnermeier [2009], Baron et al. [2021].

of work, dating back many years, that measures direct exposures between counterparties[11]. Since the Great Financial Crisis, the importance of indirect exposures has been emphasized and measures thereof have been introduced. These contagion models have sought to quantify the systemic implications of indirect exposures[12]. We contribute to the literature on financial exposures by introducing higher-order exposures as a third component of exposure.

Our introduction of higher-order exposures is motivated by the large theoretical and empirical literature on the importance of contagion and amplification of financial shocks to the dynamics of financial crises. The idea that higher-order interconnections between one institution i and another j could pose a financial risk to institution i if institution j defaults (through the process of financial contagion) is well-understood and has often been modeled.[13] Yet, these "higher-order exposures" have neither been conceptually introduced nor formally measured before. We propose that one institution's exposure to another should be understood as the sum of its direct, indirect and higher-order exposures, rather than just its direct exposures alone.

Our methodology contributes to the literature on contagion in financial networks by formalizing the quantification of higher-order exposures, using techniques which are fairly standard in the contagion literature[14]. The multi-layered financial network literature has demonstrated that contagious losses may be amplified by the interactions between different contagion channels, and are underestimated when these interactions are ignored[15]. By capturing how a single tradable security generates exposure both through counterparty risk to the security's issuer as well as through overlapping portfolio contagion when the security is fire-sold, we innovate on the literature that studies contagion in securities markets[16].

Allen and Babus [2009] point out that financial institutions like banks and funds are interconnected through both their assets and liabilities, and contagion may propagate across both. Empirical studies like Greenwood et al. [2015] and Duarte and Eisenbach [2021] investigate the spread of contagion via the asset side of banks, and Fricke and Fricke [2021] the spread of contagion across funds' assets. Sydow et al. [2024] cover the European banking system as well as funds investing in equities and/or bonds. They capture both asset- and liability-side contagion and find that the inclusion of the fund sector amplifies losses in the banking sector. Building on this literature, we focus on the asset- and liability-side interactions between the highly concentrated South-African banking sector and a diverse network of South African funds, including money market funds, fund-of-funds and a broad range of other funds. We find that exposures to the banking sector differ significantly across this wide variety of funds.

Our empirical case study of the South-African financial system further shows that even when

[11]See e.g. Altman and Saunders [1997], Crouhy et al. [2000], Jorion and Zhang [2009], Duffie and Singleton [2012], Begenau et al. [2015], Bluhm et al. [2016].

[12]See e.g. Cifuentes et al. [2005], Adrian and Shin [2010], Cont and Wagalath [2013], Caballero and Simsek [2013], Greenwood et al. [2015], Clerc et al. [2016], Cont and Schaanning [2017], Calimani et al. [2017], Aymanns et al. [2018], Cont and Schaanning [2019], Aldasoro et al. [2020].

[13]See e.g. Allen and Gale [2000], Gai and Kapadia [2010], Kok and Montagna [2013], Elliott et al. [2014], Glasserman and Young [2015], Wiersema et al. [2023], Farmer et al. [2020].

[14]See e.g. Gai and Kapadia [2010], Cont et al. [2010], Elliott et al. [2014], Acemoglu et al. [2015], Bardoscia et al. [2017].

[15]See e.g. Brunnermeier and Pedersen [2009], Krishnamurthy [2010], Brunnermeier and Oehmke [2013], Kok and Montagna [2013], Wiersema et al. [2023].

[16]See e.g. Caccioli et al. [2014], Greenwood et al. [2015], Duarte and Eisenbach [2021], Cont and Schaanning [2017].

both direct and indirect exposures are considered, exposure is still significantly underestimated as higher-order exposures are often substantial, even when direct and indirect exposures are minimal or completely absent. This finding reinforces the importance of taking all components of exposure into account when calculating institutions' exposures for prudential regulatory purposes, thereby contributing to the literature on (macro)prudential regulation.[17] To the best of our knowledge, we are the first to point out that large exposure limits (which have been widely adopted following the Great Financial Crisis across various jurisdictions, including in South-Africa, in line with the Basel accords[18]) and other elements of the prudential toolkit that hinge on exposure measures may be significantly misguided if based solely on direct exposures – which is the current standard practice.

Outline. The remainder of this paper is structured as follows: We first illustrate how higher-order exposures materialize in a styled model of a financial system in section 3. In section 4, we discuss our data and methodology for measuring higher-order exposures in the South African financial system. Section 5 presents our results. We analyze system-wide, sectoral, individual and stressed exposures in the South-African system, with particular focus on the higher-order component of these exposures. Section 6 discusses the implications of our findings.

3 Exposures in a Stylized Financial System.

To demonstrate how higher-order exposures materialize, we evaluate the exposures in a stylized model of a financial system. Consider three financial institutions, i, j and k, which are part of a larger economic system. Institution i has extended a loan l_{ik} to institution k. Furthermore, i has external assets a_i and external debt d_i (i.e. investments in and debt to institutions other than i, j and k). Institution i's equity is denoted as e_i. Institution j holds a number s_j shares in stock S and has external assets a_j, external debt d_j, and equity e_j. Lastly, institution k holds a number s_k shares in stock S and has debt $d_{ki} = l_{ik}$ to i. k also has external assets a_k, external debt d_k, and equity e_k. The balance sheets are shown in panel 1 and our notation is summarized in table 4 in the Appendix. We assume for simplicity that the institutions' external assets a_q and external debt d_q, $q \in \{i, j, k\}$, do not generate exposures.

[17] See e.g. Persaud [2009], Freixas et al. [2015], Aikman et al. [2019], Jeanne and Korinek [2020].
[18] See BIS [2018b].

Assets	**Liabilities**
Loan l_{ik}	External Debt d_i
External Assets a_i	**Equity** e_i

(a) Balance sheet institution i

Assets	**Liabilities**
Shares s_j	External Debt d_j
External Assets a_j	**Equity** e_j

(b) Balance sheet institution j

Assets	**Liabilities**
Shares s_k	Debt d_{ki}
External Assets a_k	External Debt d_k
	Equity e_k

(c) Balance sheet institution k

Table 1: **Balance sheets of institutions i, j and k.**

Let us consider the direct exposures in this system first. The loan from institution i to k gives rise to a direct exposure E_{ik}; the value of the loan l_{ik} is written-down to its recovery value when k defaults. We assume throughout this paper that the *LGD* is 100% (so recoveries are assumed to be zero, but note that this assumption is easily modified). The time taken to resolve a default is much longer than the typical timescales over which contagion materializes so short-term recovery may be realistically assumed to be zero (Elsinger et al., 2006, Cont et al., 2010). The loan l_{ik} is written-off when k defaults, so i's corresponding direct exposure to k is given by

$$E_{ik} = l_{ik}. \tag{3}$$

This direct exposure is visualized by the red arrow in Figure 1.

Let us now discuss the indirect exposures in this system. The portfolio overlap between institutions j and k generates an indirect exposure; when j defaults, its assets are liquidated as part of the default resolution, which includes the sale $\Delta s = s_j$ of j's shares in stock S. Such a fire sale typically depresses the market price P^s of stock S, which is referred to as the sale's *price impact* (Dufour and Engle [2000], Lillo et al. [2003], Potters and Bouchaud [2003]). While movements in market prices are a function of supply as well as demand, we do not explicitly model the demand side. Rather, Following Cont and Schaanning [2017], we assume a linear price-impact function of the form

$$\Delta P^s = \frac{\Delta s}{D^s}, \tag{4}$$

where Δs is the number of shares in stock S sold, and the market depth D^s expresses the number of shares sold per unit change in price P^s. The market depth D^s is a measure of the liquidity of stock S and approximates the sensitivity of the price P^s to changes in the supply of shares in stock S. D^s depends, among other things, on the average daily trading volume and daily volatility of the security.[19]

[19] See e.g. Potters and Bouchaud [2003], Eisler et al. [2012], Cont et al. [2014], Cont and Schaanning [2017].

Modern accounting practices require mark-to-market accounting of tradable securities. The accounting value of a tradable security is thus driven by its market price. Hence, when a financial shock forces an institution to liquidate its tradable securities, other institutions that hold the same securities suffer mark-to-market losses. These institutions may then be forced to liquidate tradable securities in turn. The propagation of market-to-market losses through (fire)sales is typically referred to as *overlapping portfolio contagion* (Caccioli et al., 2014, Farmer et al., 2020). Institution k's indirect exposure E_{kj} to the default of j is given by k's mark-to-market loss on its shares s_k, resulting from the price-impact of the liquidation of j's shares:

$$E_{kj} = \Delta P^s s_k = \frac{s_j s_k}{D^s}. \tag{5}$$

As a consequence of the linear price-impact function we use, j's indirect exposure E_{jk} to the default of k is also, symmetrically, given by

$$E_{jk} = \frac{s_j s_k}{D^s}. \tag{6}$$

The right-hand side of equation (6) is referred to as the *liquidity-weighted portfolio overlap* of j and k in stock S (Cont and Schaanning, 2017, Poledna et al., 2021). These indirect exposures are visualized by the blue arrows in Figure 1.

Finally, we are ready to discuss the higher-exposures in this system. Assume that k's indirect exposure to j exceeds its equity buffer, i.e.

$$E_{kj} > e_k. \tag{7}$$

Hence, when j defaults and its shares s_j in stock S are liquidated, k's mark-to-market loss from the price-impact causes k to default too, resulting in a write-off of i's loan to k. Hence, i has a higher-order exposure to the default of j equal to the size of i's loan to k;

$$E_{ij} = l_{ik} \tag{8}$$

This higher-order exposure is visualized by the purple arrow in Figure 1. Because i has no direct or indirect exposures to j, its higher-order share of exposure to j is given by

$$HSE_{ij} = \frac{l_{ik}}{0 + 0 + l_{ik}} = 100\%. \tag{9}$$

In sum, without considering higher-order exposures, i completely overlooks its exposure to j.

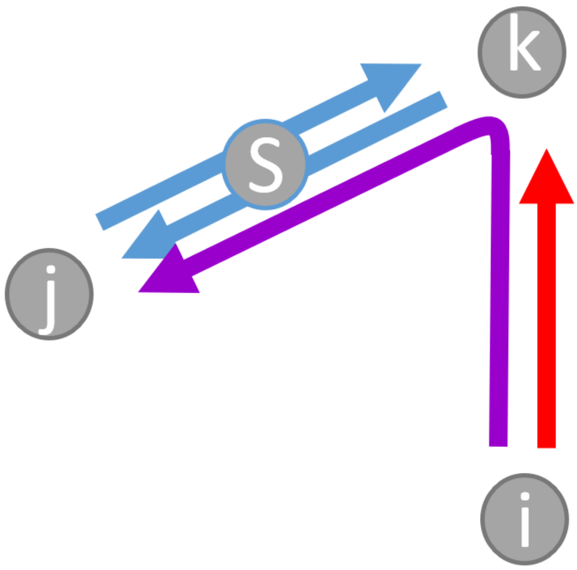

Figure 1: **Stylized example of a higher-order exposure.** Bank i lends to bank k. This *direct exposure* is depicted by the red arrow. Banks j and k have a large portfolio overlap through their position in stock S. When either defaults, its position in s is liquidated, which results in mark-to-market losses to the other bank, causing it to default too. This *indirect exposure* is depicted by the blue arrows. Bank i has neither a direct nor an indirect exposure to bank j. Yet, when bank j defaults, k defaults as a consequence, causing k's debt to i to be written-off. Hence, bank i has a *higher-order exposure* to bank j, which is depicted by the purple arrow. The magnitude of i's higher-order exposure to j is equal to the loss i stands to suffer when j defaults, i.e. the value of i's loan to k.

This example illustrates that a regulator needs granular data in order to understand the exposures of regulated institutions, as well as models that use these data to simulate how financial distress propagation results in higher-order exposures. Furthermore, the example also highlights the importance of capturing all relevant contagion channels and their interactions in the same model. Studying either the counterparty default contagion channel or overlapping portfolio contagion channel individually would result in overlooking the higher-order exposure of i to j completely. In the example, i's higher-order exposure to j was generated by the knock-on default of k when j defaults. As discussed below, higher-order exposures occur in the absence of knock-on defaults too, as institutions may propagate losses through counterparty risk and shareholder contagion without defaulting (and, potentially, through other contagion channels not considered in this paper as well).

4 Exposures in the South African Financial System[20]

We apply the concept of higher-order exposures to the South African financial system. South Africa is a small open economy with a relatively well-developed financial market compared to other African or emerging-market economies. The South African debt market is liquid and well-developed in terms of the number of participants and their daily activity, and its equity market dominates the region in terms of capitalisation (Andrianaivo and Yartey, 2010). The

[20]All data presented in this section are as of Q42016.

local currency is South African Rand, which will henceforth be referred to as ZAR.

Banking sector assets exceed GDP in aggregate terms, but are smaller than the assets held by the non-bank financial intermediation sector, which includes entities such as insurers, pension funds and collective investment schemes, which are henceforth referred to as "funds". Since the Global Financial Crisis, the share of assets held by banks has decreased, as the growth of assets held by the non-bank financial sector – in particular funds - has outpaced that of banks (Kemp, 2017). Non-bank financial intermediaries are an important source of funding for banks and direct linkages among banks and non-bank financial intermediaries other than pension funds and insurers are relatively high, amounting to 15% of bank assets (FSB, 2018).

4.1 Institutions

In this study, we focus on banks and funds domiciled in South Africa. Pension funds and insurers are not included due to data limitations, but we do not expect this to affect our results significantly, as pension funds and insurers typically do not generate contagion in our model. Non-financial corporates, henceforth referred to as the corporate sector, and the South African government are not modeled. However, we include the tradable securities that corporates and the government issue through our data of banks' and funds' investments in these securities.

4.1.1 Banks

The South African banking sector comprises 34 registered banks, local branches of foreign banks and mutual banks. The sector is concentrated, with the five largest banks by assets holding more than 90% of the banking sectors' assets (SARB, 2017a), as illustrated in Figure 2a. Overall, the banking sector is largely funded by deposits, but banks also issue debt instruments, such as bonds and money market instruments (MMIs), and equity shares.

We calculate higher-order exposures of banks and funds to each of the six largest banks, as they form the core of the South African financial system and generate the largest exposures. The six largest banks (by total assets) are the Standard Bank of South Africa Ltd (Standard Bank), FirstRand Bank Ltd (Firstrand), Absa Bank Ltd (Absa), Nedbank Ltd (Nedbank), Investec Bank Ltd (Investec) and Capitec Bank Ltd (Capitec). While the assets held by Capitec are significantly smaller than the assets held by the other five large banks, it is included in our analysis given that it is the second largest retail bank based on the number of customers.

4.1.2 Funds

Funds pool investors' money and purchase a portfolio of securities, thereby offering investors the opportunity to obtain exposure to a diverse portfolio of underlying securities, without having to purchase and trade securities directly. From the investor's perspective, funds provide investors with an opportunity to earn higher returns than those offered by deposits, in return for taking on greater risk. There are over 1200 open-ended funds registered South African funds with assets under management of about 2 trillion ZAR. We divide funds into three categories according to the instruments they invest in: money market funds (MMFs), fund-of-funds (FoFs) and other funds (OFs). Participation bond schemes and hedge funds are not within the scope of this study given data limitations and their relatively small size (see Kemp [2017]).

By investing in funds, investors buy fund shares, each of which represents ownership of a portion of the underlying portfolio. These shares are typically redeemable on a daily basis. The value of a fund share is given by its Net Assset Value (NAV), which is equal to the the fund's total asset value, divided by the fund's total shares outstanding. Fund shares can be either Variable NAV-valued (VNAV) or Constant NAV-valued (CNAV). When a fund makes a profit or loss, a VNAV fund adjusts the shares' NAV to reflect this while keeping the number of shares that shareholders own constant, whereas a CNAV fund adjusts the number of shares that each shareholder owns while keeping the NAV constant. South-African MMFs are CNAV-valued, while other South-African funds are VNAV-valued. While the mechanism through which VNAV and CNAV funds pass on their profits and losses to their shareholders is different, the impact on the value of an investor's share portfolio is identical. Therefore, we simply assume for that all fund shares in our model are VNAV-valued.

In the case of a default of, e.g., one of the five largest banks (which borrow substantially from funds), funds could act as shock absorbers by spreading losses across a diverse set of investors. In extreme circumstances, however, open-ended funds involved with credit intermediation together with leverage, liquidity or maturity transformation could be susceptible to "runs" – i.e. large-scale redemption requests, when investors anticipate or observe a substantial drop in their fund shares' value. Funds tend to observe capital outflows following an NAV loss (see e.g. Coval and Stafford [2007], Goldstein et al. [2017]). Redemptions may force funds to firesell assets to raise liquidity, which depresses prices and further reduces the fund's NAV. Vigilant investors anticipate redemptions following an NAV loss and exit first at favorable prices. This first-mover advantage may lead to a cycle of increasing redemptions and falling prices (Cont and Wagalath [2013]). When such a run is initiated, the fund may run out of liquid assets and become unable to meet redemptions, having to suspend redemptions and risk being wound up.

Money Market Funds

MMFs are formally designated according to legal requirements (Board, 2014): Board Notice 90 of 2014 (Board, 2014) restricts the money-market instruments that a fund manager may invest in, both in terms of maturity of the investments and in terms of the maximum counterparty exposure (inclusion limits).

For the restriction on maturity transformation, the weighted average legal maturity of the fund may not exceed 120 days, while the weighted average duration of the money-market instruments may not exceed 90 days. No single instrument that MMFs invest in may have a maturity exceeding 13 months. The regulations also limit the exposure in terms of the maximum percentage of the aggregate market value of the portfolio. These limitations include a maximum exposure of 30% of total fund assets to (MMF instruments issued by) local or foreign banks (registered in South Africa) of which the holding company is listed on the exchange if the market capitalisation of the listed group holding company exceeds R20 billion, and to 20% if the market capitalisation of the listed group holding company is between R2 billion and R20 billion. MMFs can also invest in money market instruments issued by any local or foreign entity that is listed on an exchange. This exposure is limited to 10% per issuer.

The MMF industry in South Africa is relatively small, amounting to 2% of total financial assets. The sector is concentrated - of the 49 MMFs in South Africa, 82% of assets are held by the 10 largest MMFs as illustrated in Figure 2b. Money-market instruments issued by banks and deposits placed with banks make up 90% of the overall portfolio of MMFs, with the remainder of the holdings made up of instruments issued by non-banks (SARB, 2017b). Amid these regulatory restrictions on instruments that MMFs are allowed to invest in, MMFs in South Africa have large exposures to banks making the portfolios of MMFs exceptionally similar (Kemp, 2017).

Fund of Funds

Although there is no formal distinction among non-MMF funds, a specific subset of these funds can be singled-out because of their particular investment strategy (Kemp, 2017). These "Fund-of-Funds" (FoFs) invest predominantly in other funds' shares and give investors exposure to multiple fund schemes and, consequently, diversification across management styles. Due to their investments in other funds, FoFs are highly exposed to instabilities in the fund sector. This potentially generates significant higher-order exposures, which is why this particular subset of funds is singled-out. We classify funds that invest more than 80% of their portfolio in the shares of other funds as FoFs. Our data include over 400 FoFs, and the distrubtion of their total asset sizes is show in Figure 2c.

Other Funds

The remaining (i.e. non-MMF, non-FoF) "Other Funds" (OFs) include equity funds (which predominantly invest in equity shares), non-MMF fixed income funds (which typically invest in debt instruments with longer maturities), multi-asset funds (whose investments include both equity shares and debt instruments), and real-estate investment trusts. As these OFs invest in a mixture of instruments issued both domestically and off-shore, their portfolios are highly variable. Our data include over 800 OFs and the distribution of their total asset sizes is show in Figure 2d.

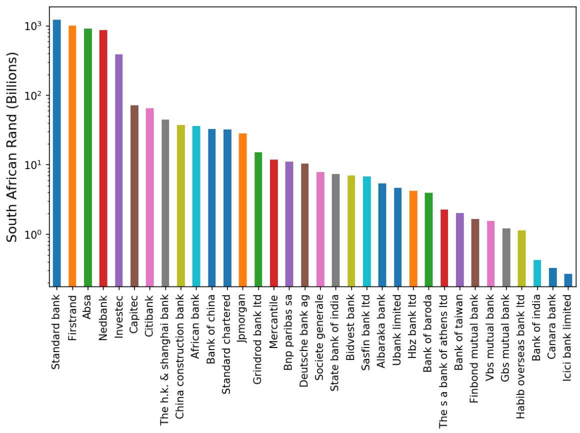

(a) Banks ranked by total assets

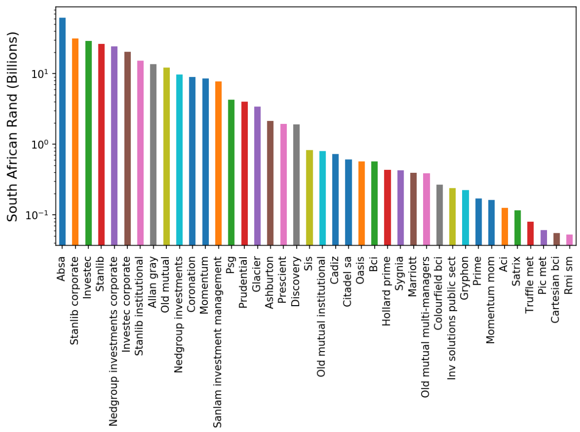

(b) MMFs ranked by total assets

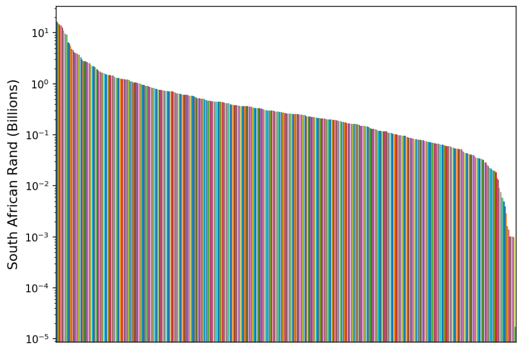

(c) FoFs ranked by total assets

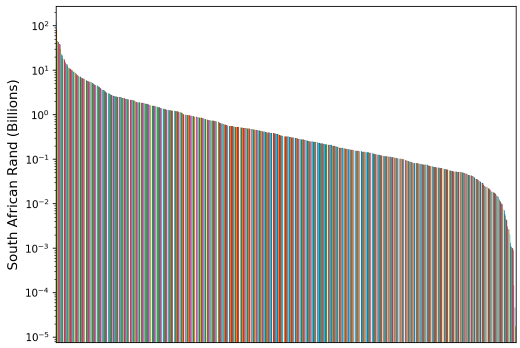

(d) OFs ranked by total assets

Figure 2: **Distribution of South African financial institutions by asset size.** The institutions are listed on the x-axis in decreasing order of total assets size and their total assets in billions of South African Rand are on the y-axis (log-scale). Note that the FoFs' and OFs' names are not listed because they are too numerous. The banking sector consist of a core of six large banks – the Standard Bank of South Africa Ltd (Standard Bank), FirstRand Bank Ltd (Firstrand), Absa Bank Ltd (Absa), Nedbank Ltd (Nedbank), Investec Bank Ltd (Investec) and Capitec Bank Ltd (Capitec) – and a periphery of 28 smaller banks. The FoF and OF sectors also show a strong concentration in terms of asset size, whereas the concentration in the MMF sector is less pronounced.

4.2 Network Data

The data used are sourced from two publicly available data sets, on for the banks and one for the funds. Aggregate balance sheet data (aggregate assets, liabilities and equity) on individual banks are sourced from the BA900 data published by the South African Reserve Bank (SARB, 2016b). Balance sheet entries are aggregated by asset type and counterparty type (e.g. "loans and deposits to domestic banks"). Data on funds' assets were sourced from Morningstar Inc and are highly granular. These data report funds' investments per instrument type in individual counterparties (e.g. "bonds in Absa"). The data do not explicitly report funds' shareholders. However, we know from the banks' balance sheet data that banks do not invest funds, and from the funds' asset data we know funds' holdings of shares in other funds. Hence, we have data at the level of individual counterparties on all of a fund's shareholders that are included in the model.

4.2.1 Balance Sheet Composition

We focus on five types of domestic assets (foreign assets are not modeled): loans and deposits (l), bonds (b), money market instruments, or *MMIs* (m), equity shares (e) and fund shares (f). Figure 2 shows where these assets appear on the stylised balance sheet of a bank, MMF and FoF/OF. The data also include investments in gold, repo, foreign institutions and the real economy. We do not model these assets, however, as they are not expected to cause substantial contagion through the channels included in this paper.[21]

Assets		**Liabilities**
Loans + Deposits (l)		Loans + Deposits (l)
Tradable securities	Bonds (b)	Bonds (b)
	MMIs (m)	MMIs (m)
	Equity Shares (e)	Other liabilities
Other assets		**Equity** (e)

(a) Stylised balance sheet of a bank.

Assets		**Liabilities**
Deposits (l)		Fund shares (f)
Tradable securities	MMIs (m)	
Fund shares (f)		
Other assets		

(b) Stylised balance sheet of a MMF.

Assets		**Liabilities**
Deposits (l)		Fund shares (f)
Tradable securities	Bonds (b)	
	MMIs (m)	
	Equity shares (e)	
Fund shares (f)		
Other assets		

(c) Stylised balance sheet of FoFs and OFs.

Table 2: **Stylised balance sheets of the types of modeled financial institutions in the South-African financial system**. We consider: (a) banks; (b) MMFs; and (c) FoFs and OFs. FoFs have the same balance sheet structure as OFs, but FoFs invest more than 80% of their assets in fund shares.

Loans & Deposits (l)

Only banks receive loans and deposits, because funds do not have debt. The bank data do

[21] The market for gold is very liquid and therefore the price of gold will not be affected substantially when one of the South African banks defaults and sells its gold. Furthermore, repo is collateralized and should therefore, at least theoretically, not be subject to counterparty default (risk) contagion (and the repo market is also quite small). Lastly, the foreign sector and real economy are not modelled.

not distinguish between deposits and loans (of any maturity), so they are all treated as one and the same and we only distinguish between debt to different counterparties.

Bonds (b)

Bonds are issued by banks, the corporate sector and the South African Government. We do not distinguish between bonds of different maturities, but only between bonds issued by different institutions. Contrary to loans, bonds are tradable.

MMIs (m)

Money market instruments are defined in line with Board Notice 90 of the Financial Sector Conduct Authority (Board, 2014), and include commercial paper, negotiable certificates of deposits, bankers acceptances and promissory notes. The data do not distinguish between these various types of MMIs so they are treated as one and the same and we only distinguish between MMIs issued by different counterparties. Like bonds, MMIs are tradable.

Equity Shares (e)

Equity shares are issued by banks and the corporate sector. The bank data distinguish between listed equity, unlisted equity, and redeemable preference shares, but the fund data do not. For simplicity, all three types of shares are all modeled as listed equity and we only distinguish between shares issued by different institutions. Therefore, all modeled equity shares are tradable.

Fund Shares (f)

Fund shares are issued by funds and are (almost always) redeemable on a daily basis. Therefore, fund shares are not traded. As explained, we assume that all fund shares are VNAV-valued for simplicity, so the shares' NAV is updated to reflect any losses that the issuing fund may suffer.

4.2.2 Initialization Values

We do not have data on the market prices or NAVs of the financial securities $\sigma \in \{b, m, e, f\}$ that institutions hold, nor the number of securities they hold, but only on the value of an institution's positions in a security (i.e. the market value of a position in a tradable security and the NAV times the number of shares of a position in shares issued by a fund). We normalize the (initial) NAV of each fund share and the (initial) market price of each tradable security $\tau \in \{b, m, e\}$ to 1 ZAR. This normalization has no effect on our results and is only for simplicity.

4.3 Network Construction

The data can be represented as a weighted, directed network, with the nodes representing the institutions and the edges their assets (a node's out-edges are given by its assets and its in-edges by its liabilities and/or issued shares). We refer to this as the "asset network". The network is multiplex, consisting of five layers, which each layer representing one of the asset types $\alpha \in \{l, b, m, e, f\}$. Each layer includes the same set of nodes, made up by the (individual) banks and funds, the node $\mathcal{G}$ representing the government, and a single "corporate" node $\mathcal{C}$ representing the domestic non-financial corporate sector, but we include the tradable securities that they issue.[22] As noted before, we do not model the government and corporate sector.[23] We use $\mathcal{B}$ to denote the set of banks, $\mathcal{F}$ the set of funds, $\mathcal{I} = \mathcal{B} \cup \mathcal{F}$ denotes the financial institutions, $\mathcal{H} = \mathcal{B} \cup \mathcal{G} \cup \mathcal{C}$ the set of securities-issuing institutions, and $\mathcal{A} = \mathcal{B} \cup \mathcal{F} \cup \mathcal{G} \cup \mathcal{C}$ the set of all institutions.

Let w_{ik}^{α} denote the weight of an edge pointing from node $i \in \mathcal{I}$ to node $k \in \mathcal{A}$ in the layer of the asset network corresponding to investments of type $\alpha \in \{l, b, m, e, f\}$. For loans and deposits (l), we set the weight w_{ik}^{l} equal to the sum of the principals of i's loans to and deposits at k. For securities $\sigma \in \{b, m, e, f\}$, we set the weight w_{ik}^{σ} equal to the *initial* value of i's position in securities of type σ issued by k. (Note that the values of securities may change when we simulate the losses following a default, by the weights w_{ik}^{α} do not.)

We only model the banks' and funds' investments of type $\alpha \in \{l, b, m, e, f\}$ in domestic institutions $k \in \mathcal{A}$. Therefore, each of a bank's or fund's modelled assets can be represented by a weighted, directed edge in the corresponding layer of the asset network. The directed edges corresponding to funds' assets are given directly by our detailed data on the funds' investments in individual banks and funds. The directed edges corresponding to banks' investments in individual counterparties are randomly generated, based on the aggregate data per asset type provided by the banks' balance sheets. We explain the algorithm to reconstruct the interbank network next.

4.3.1 Interbank Network Reconstruction

The technique used for the reconstruction of the banks' investments is is similar to Kok and Montagna [2013] and aims to reproduce the high heterogeneity of interconnections observed in financial networks. From the banks' balance sheet data, we know that banks do not invest in funds, and we know each bank's investments, per asset type, in the government $\mathcal{G}$ and corporate sector $\mathcal{C}$. Therefore, only the interbank investments require reconstruction. We assume that the (initial) market value of any security that a bank has issued is equal to the book value of that liability or equity share on the banks' balance sheets, and perform the following steps for each of the asset types $\beta \in \{l, b, m, e\}$ in which banks invest:

1. We subtract from each bank's aggregate liabilities (or equity) of type β the funds' investments of type β in that bank.

2. We pick a random pair of banks y and z, where bank y is the investor and bank z is the investee. Bank y is picked from the banks with nonzero aggregate assets of type β and z is picked from the banks with nonzero aggregate liabilities (or equity) of type β.

[22] We do not have disaggregate data on investments in the corporate sector, but do not expect this to affect our results substantially as this only affects the granularity of the overlapping portfolio contagion channel.

[23] The government node and corporate node are only included in the asset network to capture the indirect exposures generated by financial instruments these nodes issue. As we focus on exposures between financial institutions, the government node and corporate node are not included in the exposure networks introduced below.

3. We pick a random number $x \in U(0,1)$ and generate an investment of type β of bank y in bank z equal in size to the product of x and the minimum of y's aggregate assets of type β and z's aggregate liabilities (or equity) of type β.

4. The investment is added to the network layer of investments of type β (i.e. added to w_{yz}^{β}) and subtracted from y's aggregate assets of type β and z's aggregate liabilities (or equity) of type β.

5. Steps 2-4 are repeated until all banks' assets of type β are allocated.[24]

After step 5, the asset network is complete and all of its edges w_{ir}^{α} are defined.

4.3.2 Direct Exposure Network

The total value of all investments of an institution $i \in \mathcal{I}$ in a counterparty $k \in \mathcal{I}$ is written-off when the counterparty defaults (under the previously explained assumption of zero short-term recoveries), such that i's direct exposure to k is equal to

$$\hat{w}_{ik}^{\delta} = \sum_{\alpha \in \{l,b,m,e,f\}} w_{ik}^{\alpha}. \tag{10}$$

Here, $\hat{w}_{ik}^{\delta}$ gives the weight of the edge from i to k in the direct exposure network denoted by the superscript δ. Note that $i, k \in \mathcal{I}$, so the banks and funds are included in the direct exposures network, but the government and corporate nodes are not. The resulting direct exposure network is visualized in Figure 3.

[24] The last 500 ZAR are invested in a single chunk so the algorithm terminates: we set $x = 1$ when the minimum of y's aggregate assets of type β and z's aggregate liabilities (or equity) of type β is less than or equal to 500 ZAR.

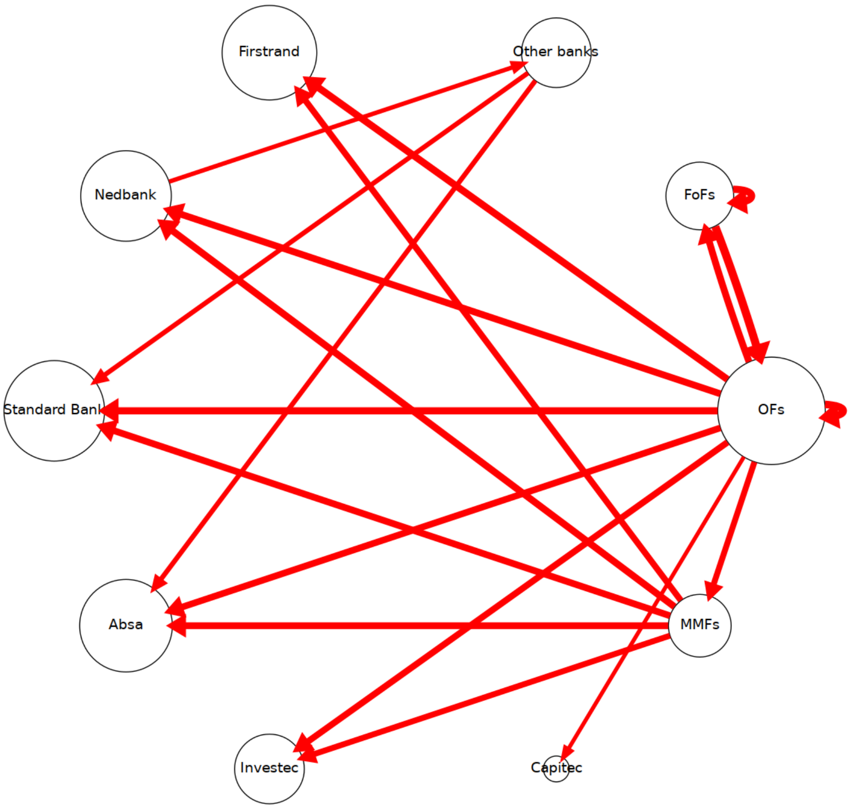

Figure 3: **Network of the largest direct exposures in the South African financial system**. We plot the largest direct exposures between the six largest banks, the rest of the banking sector, and the three fund sectors. Edge widths visualize the size of the exposure, varying between 5 bln. to 196 bln. South African Rand, and node sizes visualize the total asset size of the node, ranging from 717 bln. to 1.6 trn. South African Rand. The scaling is logarithmic for both edge widths and node sizes. Other than the six largest banks, all institutions of the same type are aggregated into a single node. Edges point in the direction of the exposures. For example, the edge from the MMFs to Absa denotes the sum of all MMFs' direct exposures to Absa. The OF and FoF sectors' self-loops denote the sums of all direct exposures between OFs or FoFs. The figure shows that the largest exposures are those of the MMF and OF sectors to the six large banks, and the OF and FoF sectors' exposures to itself.

4.3.3 Indirect Exposure Network

Indirect exposures are generated by tradable securities $\tau \in \{b, m, e\}$. As all tradable securities in our model are issued by nodes in the network, the indirect exposure between two financial institutions i and k is calculated by summing their liquidity-weighted portfolio overlap across all tradable securities $\tau \in \{b, m, e\}$ issued by all institutions $q \in \mathcal{H}$ (note that funds do not issue tradable securities):

$$\hat{w}^{\phi}_{ik} = \sum_{\tau \in \{b,m,e\}} \sum_{q \in \mathcal{H}} \frac{w^{\tau}_{iq} w^{\tau}_{kq}}{D^{\tau}_{q}}, \tag{11}$$

where $\hat{w}^{\phi}_{ik}$ gives the weight of the edge from i to k in the indirect exposure network denoted by the superscript ϕ. The edge weight gives the mark-to-market losses that i suffers when k defaults, and is equal to the mark-to-market losses that k suffers when i defaults because we assume a linear price impact. The market depth parameter D^{τ}_{q} of the tradable security τ issued by institution q is measured in units of ZAR, as explained below. (Note that the market depth described in the stylized example in section 3 had different units for simplicity.) The fraction w_{kq}^{τ}/D_q^{τ} reflects the price impact, as a fraction of the security τ's initial price, that results from k's liquidation of its position in τ. The product of this price impact and the initial value w_{iq}^{τ} of i's position gives the mark-to-market loss that i suffers.

When regulators evaluate indirect exposures, they may calibrate the exposures to real market estimates of the market depths D_q^{τ} to estimate indirect exposures more accurately. Do to data limitations, we do not have access to such estimates, however, so we use

$$D_q^{\tau} \stackrel{\text{def}}{=} \frac{C_q^{\tau}}{\mu} \tag{12}$$

as a proxy of the market depth throughout this paper, in line with Wiersema et al. [2023]. C_q^{τ} denotes the market capitalization of the tradable security τ issued by institution q[25]. Furthermore, μ is a nondimensional constant of order one, which denotes the ratio of the security's market cap to market depth, or simply *cap-to-depth ratio*. We use a security's market capitalization as our baseline estimate of the security's market depth (i.e. $\mu = 1$). Comparison with estimates for other financial markets suggests that the market capitalization most likely provides an conservative lower bound to the true market depth[26]. We explore various values of the cap-to-depth ratio μ in the results section to understand the sensitivity of our findings to the market depth.

Like the direct exposure network, the indirect exposure network only includes banks and funds. Figure 4 visualizes the indirect exposure network. Comparing equations (10) and (11) reveals that the same tradable securities that generate indirect exposures are also included in the direct exposure network. The literature that studies contagion via tradable assets typically models the indirect linkages only, while the direct linkages are ignored (Farmer et al., 2020). We improve on this approach by recognizing that the same tradable asset can generate both direct and indirect exposures.

[25]The market capitalization of South African government bonds as of Q4 2016 is sourced from the Q1 2017 SARB Quarterly Bulletin; https://www.resbank.co.za/content/dam/sarb/publications/quarterly-bulletins/quarterly-bulletin-publications/2017/7718/07Statistical-tables—Public-Finance.pdf. The total value $\sum_{k\in\mathcal{I}} w_{kr}^{\tau}$ of government bonds held held by institutions $k \in \mathcal{I}$ included in our data is equal to 21% of the market capitalization of South African government bonds'. As we cannot observe the market capitalization of other securities due to data limitations, we assume that the institutions $k \in \mathcal{I}$ we include also own 21% of the market capitalization of these securities. Hence, we estimate the market capitalization of these securities using $C_q^{\tau} = \sum_{k\in\mathcal{I}} w_{kr}^{\tau}/21\%$.

[26]See e.g. Potters and Bouchaud, 2003, Lillo et al., 2003, Cont and Schaanning, 2019, Fukker et al., 2022, Sydow et al., 2024. Under the assumption of linearity, the market depth gives the maximum amount of liquidity that can be extracted from the market for a security before the security's price falls to zero. Since prices cannot fall below zero even when all holders of the security liquidate their positions simultaneously (i.e. the entire market capitalization is sold at once), the market capitalization gives a conservative lower bound to the market depth under the linear approximation. In reality, the price impact is most likely a concave function of the volume sold (see aforementioned citations), which could theoretically allow the price impact of small volumes sold to exceed the linear approximation with market depth equal to the market capitalization, but such a severe price impact would probably only materialize during highly adverse market conditions, if at all.

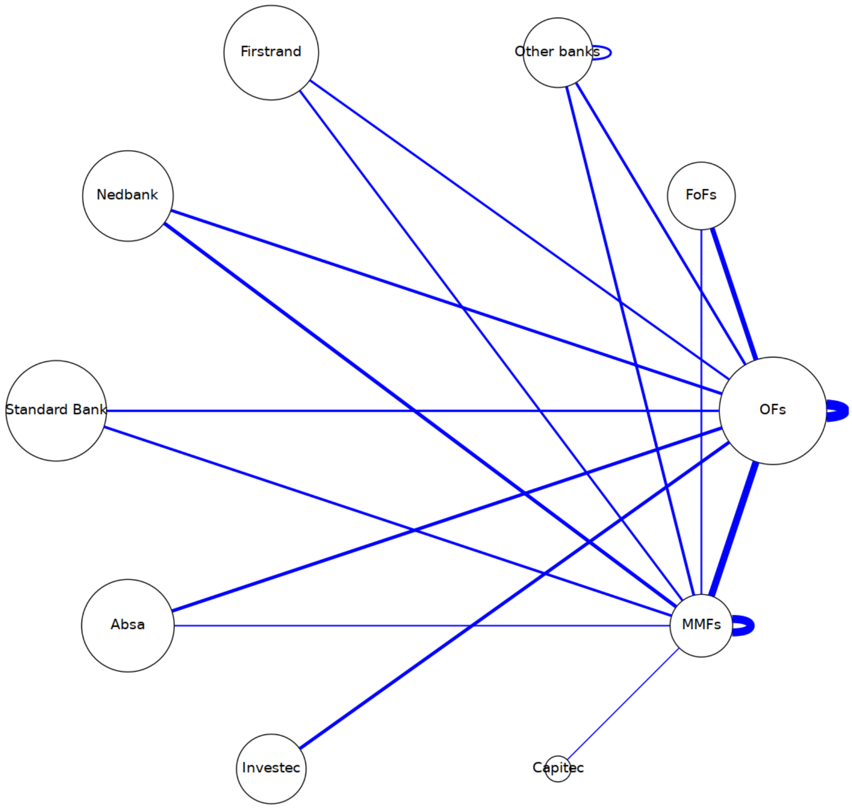

Figure 4: **Network of the largest indirect exposures in the South African financial system**. We plot the largest indirect exposures (i.e. institutions' liquidity-weighted portfolio overlaps) between the six largest banks, the rest of the banking sector, and the tree fund sectors. Edge widths visualize the size of the exposure, varying between 2 bln. and 187 bln. South African Rand, and node sizes visualize the total asset size of the node, ranging from 717 bln. to 1.6 trn. South African Rand. The scaling is logarithmic for both edge widths and node sizes. Other than the six largest banks, all institutions of the same type are aggregated into a single node. Indirect exposures are symmetric, so they are drawn as an undirected network. The MMF and OF sectors' self-loops denote the sums of all indirect exposures between MMFs or OFs. These self-loops make up the system's largest indirect exposures. Note that the MMF sector is much smaller than the OF sector in total asset size. Hence, the large indirect exposures between MMFs highlight their exceptionally similar portfolios.

4.3.4 First-Order Exposure Network

Having inferred the direct and indirect exposure networks from the asset network, we can now calculate the first-order exposure network as the sum of the edges in the direct and indirect exposure networks. The edges in this network are given by

$$\hat{w}_{ik} = \hat{w}_{ik}^{\delta} + \hat{w}_{ik}^{\phi}, \tag{13}$$

where $\hat{w}_{ik}$ gives the weight of the edge from i to k in the first-order exposure network. The first-order exposure network is visualized in Figure 5, where red edges indicate exposures that are predominantly made up of direct exposures, blue edges exposure that are predominantly indirect, and shades of purple, a mix between red and blue, indicate exposures with both a substantial direct and indirect component (the more red the shade is, the larger is the direct component and the more blue, the larger the indirect component). Note that the exposures to the six largest banks are all predominantly direct exposures, because these banks have relatively small liquidity-weighted portfolio overlaps with other institutions in the South African financial system.

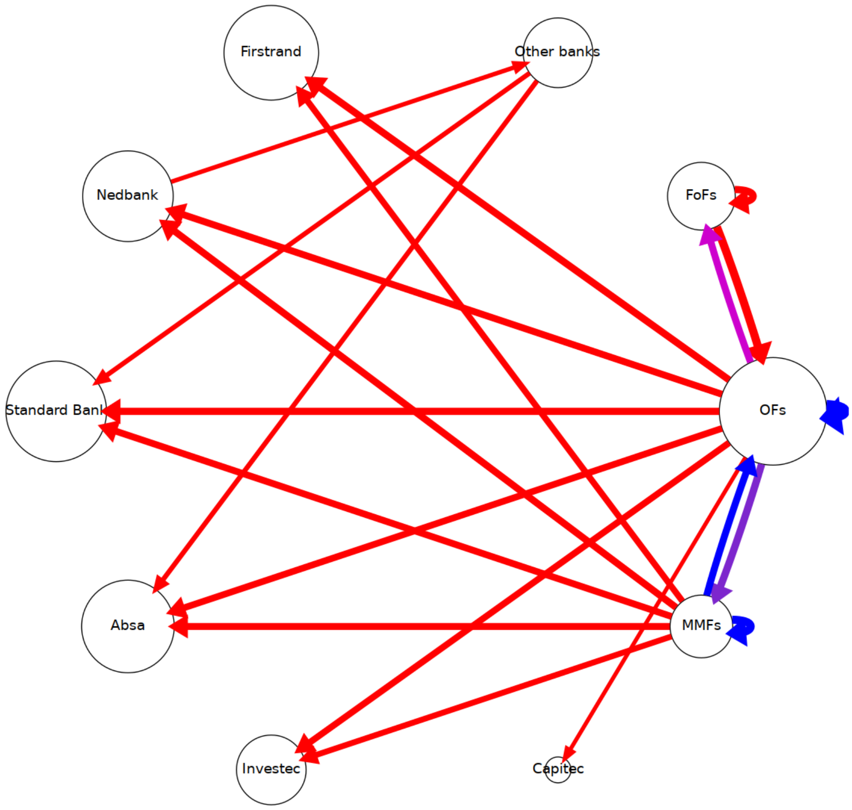

Figure 5: **Network of the largest first-order exposures in the South African financial system.** We plot the largest first-order exposures (i.e. direct + indirect exposures) between the six largest banks, the rest of the banking sector, and the tree fund sectors. Edge widths visualize the size of the exposure, varying between 5.7 bln. and 383 bln. South African Rand, and node sizes visualize the total asset size of the node, ranging from 717 bln. to 1.6 trn. South African Rand. The scaling is logarithmic for both edge widths and node sizes. Other than the six largest banks, all institutions of the same type are aggregated into a single node. Edges point in the direction of the exposures. For example, the edge from the MMFs to Absa denotes the sum of all MMFs' first-order exposures to Absa. The color of an arrow indicates the composition of that first-order exposure: A red edge indicates an exposure that is predominantly direct and a blue edge indicates an exposure that is predominantly indirect. Shades of purple, a mix between red and blue, indicate exposures with both a substantial direct and indirect component (the more red the shade is, the larger is the direct component and the more blue, the larger the indirect component). The figure shows that all large exposures to banks are direct and that OFs and MMFs have very large indirect exposures between them.

In Figure 6, we summarize our network construction by illustrating the correspondence between the layers of the asset network in the left column and the exposure networks in the right column. As each of the three tradable securities layers $\tau \in \{b, m, e\}$ (as well as the corresponding exposures) can be visualized identically, we only plot one tradable securities layer for simplicity. The figure shows on the one hand that (exposures corresponding to) multiple asset types can appear in the same exposure network. On the other hand, and more importantly, the figure visualizes that the same asset can generate exposures in both the direct and indirect network. Such assets appear as both direct and indirect exposures in the first order exposure network.

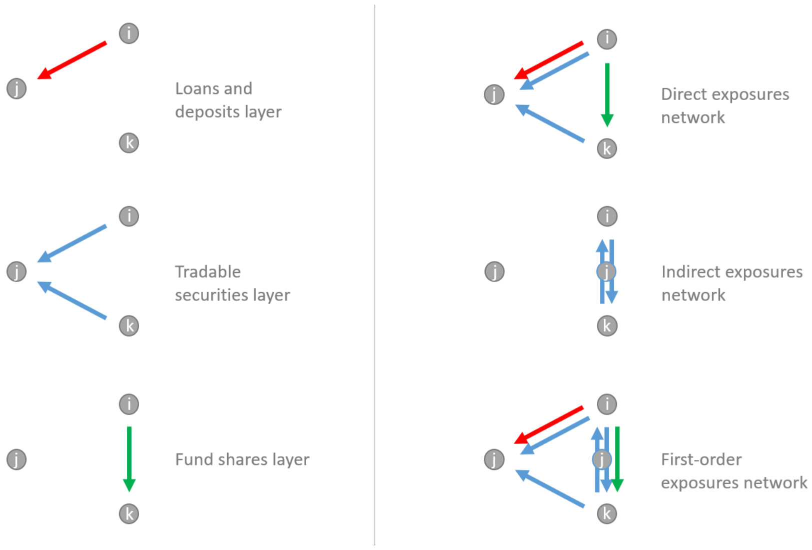

Figure 6: **Stylized correspondence between the asset network and exposure networks.** The figure summarizes our explanation of the network construction by illustrating the correspondence between the exposure networks and layers of the asset network. The left column shows the layers of the asset network and the right column the corresponding exposure networks. Because each of the three tradable securities layers $\tau \in \{b, m, e\}$ (as well as the corresponding exposures) can be visualized identically, we only plot one of the tradable securities layers for simplicity. Arrows in the left column point in the direction of the (principal) cashflow: FoF i deposits at bank j and buys fund shares in MMF k, so i suffers a write-off when either j or k defaults. The deposits are visualized by the red arrow and the fund shares by the green arrow. Both i and k buy MMIs in j, so either suffers a mark-to-market loss when the other liquidates their MMIs, and both suffer a write-off on the MMIs when j defaults. The MMIs are visualized by the blue arrows. Arrows in the right column point in the direction of the exposure and their color reflects the asset that generates the exposure. For example, an arrow from i to k represents an exposure of i to k and the color of the arrow is equal to that of the asset that generates the exposure. Exposures in the direct network are made-up of the write-offs that an institution would suffer on investments in a counterparty when that counterparty defaults. Exposures in the indirect network are made-up of institutions' liquidity-weighted portfolio overlaps, i.e. the mark-to-market losses suffered by an institution that has a portfolio overlap with another institution, when that other institution liquidates its portfolio. The first-order network is the sum of the direct and indirect networks. Note that the same MMIs generate both direct and an indirect exposures, causing exposures to multiple counterparties for both i and k in the first-order exposures network.

4.4 Measuring Higher-Order Exposures

Having set-up the asset and exposure networks, we are ready to compute the higher-order exposures. In the stylised example in Section 2, we calculated the exposure to an institution by recording all losses that would follow from the default of that institution. We did so by simulating how losses would propagate upon the default of the institution, where the propagation was dictated by the contagion mechanisms we formulated. Here, we develop a more comprehensive model of contagion in order to quantify the higher-order exposures more realistically and take into account the various types of assets included in our model. Yet, the approach is conceptually identical to the example and in fact the same approach may straightforwardly be generalized to calculate exposures in any financial system:

1. We formulate a model of how contagion channels propagate losses across the assets $\alpha \in \{l, b, m, e, f\}$ in the network.

2. We reconstruct the interbank investments from the banks' aggregate balance sheet data and initialize the asset network, based on the fund data and reconstructed interbank investments.

3. We use the contagion model to simulate how losses propagate throughout the network upon the (idiosyncratic) default of one of the six large banks. All losses incurred over the course of the simulation are recorded as exposures to this bank.

4. We repeat steps 2-3 1000 times for each of the six largest banks to average institutions' exposures to these banks over the random realizations of the reconstructed interbank network.

Each simulation thus starts with assuming the (idiosyncratic) default of one of the six large banks, to calculate the exposures to that bank. We use j to denote the bank that defaults at the start of the simulation. Furthermore, we assume discrete time dynamics and use n to denote the n^{th} round following j's default. (Note that any parameter that lacks the subscript denoting the n^{th} round after j's default is a constant.) The losses $L_{ij,1}$, $i \in \mathcal{I}$, in the first round following j's default constitute the first-order exposures $E_{ij,1}$ and are given by the first-order exposure network (eq. 13):

$$E_{ij,1} = L_{ij,1} = \hat{w}_{ij} \tag{14}$$

The losses $L_{ij,n}$ in round $n > 1$ constitute the higher-order exposures and are given by the contagion channels' propagation of the losses in the $(n-1)^{th}$ round. We specify the round in which losses were incurred as the order of the exposure, so i's n^{th}-order exposure to j is given by

$$E_{ij,n} = L_{ij,n}, \tag{15}$$

where the losses are found by summing across all asset types:

$$L_{ij,n} = \sum_{\alpha \in \{l,b,m,e,f\}} L^{\alpha}_{ij,n}. \tag{16}$$

We refer to i's losses incurred up to and including the n^{th} round as i's exposure up to n^{th} order to j and define the Higher-Order Share of Exposure (HSE) up to n^{th} order as the fraction of an exposure up to n^{th} order made up by higher-order exposures, i.e.

$$HSE_{ij,n} \stackrel{\text{def}}{=} \frac{\sum_{t=2}^{n} E_{ij,t}}{\sum_{t=1}^{n} E_{ij,t}}. \tag{17}$$

Table 3 summarizes our terminology.

<table>
<tr><td colspan="6">Exposure</td></tr>
<tr><td colspan="2">First order exposure</td><td colspan="4">Higher-order exposure</td></tr>
<tr><td>Direct exposure</td><td>Indirect exposure</td><td>Second order exposure</td><td>$\cdots$</td><td>n^{th} order exposure</td><td>$\cdots$</td></tr>
<tr><td colspan="2">$\cdots$</td><td colspan="3">Higher-Order Share of Exposure (HSE) up to n^{th} order</td><td>$\cdots$</td></tr>
<tr><td colspan="5">Exposure up to n^{th} order</td><td>$\cdots$</td></tr>
</table>

Table 3: **Terminology.**

4.4.1 Contagion Channels

The contagion channels introduced below determine the losses $L^{\alpha}_{ij,n+1}$ per investment type α in the $(n+1)^{th}$ round from the losses $L_{ij,n}$ in the n^{th} round. Once we have computed these, we can determine the higher-order exposures using equation (16). Assets are generally affected by multiple, interacting contagion channels. We first explain each contagion channel individually, before we discuss how the combination of interacting contagion channels determines the losses per investment type $L^{\alpha}_{ij,n+1}$.

Established methods for evaluating exposures assume that institutions' portfolios are fixed in the run-up to a default; direct exposures are calculated under the assumption that institutions do no liquidate their investments in a counterparty in the run-up to that counterparty's default and indirect exposures are calculated assuming fixed portfolio overlaps. We assume, in-line with these methods, that an institution does not liquidate any of its assets before it defaults, and that only upon its default are all of its tradable securities liquidated as part of the resolution process. (Assuming liquidation of tradable securities upon default is in line with contagion literature such as Burrows et al. [2012] and Caccioli et al. [2015], and practice.) This assumption restricts our model to the following three contagion channels: Counterparty (default) risk contagion, overlapping portfolio contagion, and what we refer to as "shareholder contagion" - shareholder contagion simply distributes an institution's losses across its shareholders, as explained below.

In reality, the strategic decisions and/or behavioral actions institutions' managers take during times of stress can lead to additional channels of contagion. In particular, institutions may start liquidating assets which are perceived as risky, leading to liquidity spirals (Hoerova et al., 2009, Acharya and Skeie, 2011, Gai et al., 2011). Institutions' responses to stress are inherently uncertain, however, so by assuming no liquidation of assets prior to default we avoid this source of uncertainty. Specifically, as institutions are assumed not to withdraw their deposits from banks, banks cannot default through illiquidity in our model.

Shareholder Contagion

Shares in a bank or fund represent part ownership of that institution's portfolio. Consequently, any losses incurred by the institution are distributed across its shareholders; when a fund incurs a loss, the NAV of the fund shares it has issued falls, so all holders of these shares suffer losses. Similarly, when a bank suffers a loss, the book value of the equity shares it has issued falls so all holders of the shares suffer losses accordingly. (As discussed below, equity shares are marked-to-market, but in efficient markets any book value loss is accompanied by a corresponding drop in the shares' market price.) We refer to this decrease in the NAV of a fund share or book value of an equity share as shareholder contagion.

Overlapping Portfolio Contagion

As explained previously, the liquidity-weighted portfolio overlap between two institutions gives the mark-to-market loss that either suffers when the other defaults, due to the price impact of liquidating the defaulted institution's portfolio. We refer to these mark-to-market losses that institutions suffer on their portfolios of tradable securities as overlapping portfolio contagion. Hence, the overlapping portfolio contagion that an institution suffers in round $n+1$ is driven by the price impacts of liquidating the portfolios of the institutions that default in round n.

Counterparty Risk Contagion

When an institution defaults, the value of its liabilities or fund shares is written-off. Consequently, the expected value of the liabilities or fund shares issued by any institution with non-zero default risk is below their nominal value; the higher the risk of default, the lower the risk-adjusted value. This is reflected in capital regulation (e.g. BIS [2019]), modern accounting standards (e.g. IFRS [2021]), and recent contagion literature (e.g. Bardoscia et al. [2017], Wiersema et al. [2023]). When an institution incurs a loss, its probability of default rises and the expected value of its liabilities or fund shares falls. Accordingly, its creditors or fund share holders suffer losses on the risk-adjusted value of these assets. Note that we do not claim that institutions necessarily recognize all risk-adjustment losses that they suffer, nor that their accounting standards require them to do so. We simply observe that the institutions suffer losses because the expected value of their assets is diminished, regardless of whether the institutions' accounting (accurately) reflects this.

We refer to the propagation of losses through risk adjustments as counterparty risk contagion. Equity shares are not subject to counterparty risk contagion, as equity value is (by definition) zero by the time the issuing bank defaults. All other investments in banks and funds are subject to risk adjustments due to counterparty risk contagion, as all investments other than equity shares have residual value that is written-off when the counterparty defaults. (Investments in the government and corporate sector are not subject to counterparty risk contagion, because these institutions cannot incur losses or default in our model).

4.4.2 Contagion Equations

We are now ready to discuss how to calculate the terms $L^{\alpha}_{ij,n}$ in equation (16) for each investment type $\alpha \in \{l, b, m, e, f\}$. We first introduce the value of an investment in round n and then discuss how the interacting contagion channels drive the losses in the investments' values.

Loans and Deposits

In our model, loans and deposits are only affected counterparty risk contagion and their (expected) value is given by their risk-adjusted principals. We use $\epsilon^{\rho}_{kj,n}$ to denote the ratio of the risk-adjusted value to the nominal value of an investment $\rho \in \{l, b, m, f\}$ in institution k in round n. The value $v^{l}_{ikj,n}$ of institution i's loans to and deposits at institution k in round n can therefore be written as the product of the risk-adjustment factor $\epsilon^{l}_{kj,n}$ and the sum of loans' to and deposits' principals w^{l}_{ik};

$$v^{l}_{ikj,n} = \epsilon^{l}_{kj,n} w^{l}_{ik}. \tag{18}$$

The risk adjustment factor $\epsilon^{\rho}_{kj,n}$ reflects the expected loss on the investment and may be based on sophisticated models that estimate losses using k's capital, profitability, collaterals, and other institution-specific and market-level features (as is common in expected loss modeling for capital and accounting purposes). Although our framework allows for calculating higher-order exposures using advanced models of the risk-adjustment factor, for our present purposes it suffices to derive a simple estimate of the risk-adjustment factor $\epsilon^{\rho}_{kj,n}$ based only k's *buffer*.

Institution k's buffer $B_{kj,n}$ gives the amount of losses that k can absorb before it defaults.[27] A bank is assumed to default when its equity is reduced to zero, because of insolvency. Funds cannot become insolvent as they do not have debt, but are vulnerable to runs that could make them default through illiquidity. We assume that a run on a fund is initiated (causing the fund to default) when the fund loses 45% of its total asset value. This threshold is in line with Cont and Wagalath [2013] and we explore how this assumption affects exposures in the results section. Hence, if k is a bank, its initial buffer is equal to its equity $e_{kj,1}$ (as given by the data), and when k is a fund, its initial buffer is equal to 45% of its total asset A_k:

$$B_{kj,1} = \begin{cases} e_{kj,1}, & \text{if } \ k \in \mathcal{B} \\ 0.45 A_k, & \text{if } \ k \in \mathcal{F} \end{cases}. \tag{19}$$

Banks' and funds' buffers are updated according to

$$B_{kj,n+1} = \max \left\{ B_{kj,n} - L_{kj,n}, 0 \right\}, \tag{20}$$

so the buffer cannot become negative. For any bank k, we have that

$$e_{kj,n} = B_{kj,n}, \tag{21}$$

so we can use $e_{kj,n}$ and $B_{kj,n}$ interchangeably for banks. Furthermore, because we do not model the government $\mathcal{G}$ and corporate sector $\mathcal{C}$ explicitly, they cannot incur losses and do not default. For notational convenience, we set $B_{gj,n} = e_{gj,n} = 1, \quad g \in \mathcal{G} \cup \mathcal{C}$.

We now calculate risk-adjustment factors based on institutions' buffers. Assume for simplicity that the initial risk-adjustment factor $\epsilon^{\rho}_{kj,1} = 1$ for any investment $\rho \in \{l, b, m, f\}$, i.e. that the initial risk-adjusted value of any investment is equal to the investment's nominal value given by the data. Hence, we assume that all counterparties initially have zero default risk. This is only for simplicity as it does not affect our results. As explained previously, we assume that short-term recoveries are zero, which implies that the risk-adjusted value of an investment is zero once the counterparty's buffer is exhausted (causing default). This, together with the assumption that the risk-adjusted value of an investment is proportional to the counterparty's

[27]Note that the model continues to record losses on an institution's assets after it defaults. This is for two reasons. First, when (in reality) a defaulted institution is resolved, the residual value of its assets determines the recoveries, so the depreciation of a defaulted institution's assets has economic relevance. Second, we vary institutions' buffers to study how this affects stability. When we do not record an institution's losses after its default, the size of the institution's buffer limits the amount of losses that an institution can incur. Hence, decreasing institutions buffers would artificially limit the amount of losses that institutions can incur and therefore the destabilizing impact of decreasing buffers would not be properly reflected in the losses. Figure 20 in the appendix shows that the assumption that institutions continue to incur losses after their default does not significantly affect our main results.

buffer (as in Bardoscia et al. [2017] and Wiersema et al. [2023]), implies that the risk-adjustment factor in round n for any investment $\rho \in \{l, b, m, f\}$ in k is equal to

$$\epsilon^{\rho}_{kj,n} \stackrel{\text{def}}{=} \frac{B_{kj,n}}{B_k}, \tag{22}$$

and we denote reductions in the risk-adjustment factor as

$$\Delta\epsilon^{\rho}_{kj,n} = \epsilon^{\rho}_{kj,n} - \epsilon^{\rho}_{kj,n+1}. \tag{23}$$

Note that the linear approximation in equation (22) probably forms an upper bound to the true risk adjustment function, which in reality is most likely a sublinear function (Bardoscia et al. [2017]). Hence, our risk-adjustment factor provides a conservative estimate of the counterparty credit risk channel.

Using equations 18 and 23, we find that the risk-adjustment loss that i suffers on its portfolio of loans and deposits in round $n+1$ after j's default due to counterparty risk contagion is given by

$$L^{l}_{ij,n+1} = \sum_{k\in\mathcal{B}} v^{l}_{ikj,n} - v^{l}_{ikj,n+1} = \sum_{k\in\mathcal{B}} \Delta\epsilon^{\rho}_{kj,n} w^{l}_{ik}, \tag{24}$$

where the sum runs over banks $k \in \mathcal{B}$ only because funds do not take loans or deposits.

Fund Shares

The value of a fund shares is affected by both risk adjustment and shareholder contagion. When fund k suffers a loss $L_{kj,n}$, the NAV $N_{kj,n}$ of the shares it has issued falls through shareholder contagion. As the NAV is given by k's total asset value divided by the number of shares k issued, the NAV loss is given by

$$\Delta N_{kj,n} = \frac{L_{kj,n}}{A_{kj,1}} N_{kj,1}, \tag{25}$$

where fund k's initial total assets $A_{kj,1}$ is given by the data, the initial NAV $N_{kj,1}$ is normalized to 1 ZAR, and $A_{kj,1}/N_{kj,1}$ gives the number of shares issued. The NAV evolves according to

$$N_{kj,n+1} = N_{kj,n} - \Delta N_{kj,n}. \tag{26}$$

Still, the NAV reduction following a loss does not cover the fund's increased default risk, so the (expected) value of the shares in our model is given by their risk-adjusted NAV. The value $v^{f}_{ikj,n}$ of institution i's shares in fund k in round n is therefore given by

$$v^{f}_{ikj,n} = \epsilon^{f}_{kj,n} \frac{N_{kj,n}}{N_{kj,1}} w^{f}_{ik}, \tag{27}$$

where $\epsilon^{f}_{kj,n} N_{kj,n}/N_{kj,1}$ gives the risk-adjusted NAV in round n as a fraction of the initial NAV. We use equations (23) and (26) to find that the loss that i suffers on its fund shares portfolio in round $n+1$ after j's default due to counterparty risk and shareholder contagion losses is given by

$$L^f_{ij,n+1} = \sum_{k\in\mathcal{F}} v^f_{ikj,n} - v^f_{ikj,n+1} = \sum_{k\in\mathcal{F}} \left(\Delta\epsilon^f_{kj,n} \frac{N_{kj,n}}{N_{kj,1}} + \epsilon^f_{kj,n+1} \frac{\Delta N_{kj,n}}{N_{kj,1}} \right) w^f_{ik}. \tag{28}$$

The first term within the brackets on the right hand side of equation (28) gives the counterparty risk contagion that i suffers due to risk adjustment $\Delta\epsilon^f_{kj,n}$, multiplied by the "interaction term" $N_{kj,n}/N_{kj,1}$. The second term gives the shareholder contagion that i suffers due to the fractional NAV loss $\Delta N_{kj,n}/N_{kj,1}$, multiplied by the interaction term $\epsilon^f_{kj,n+1}$. The interaction terms reflect that the fraction of an asset's value already lost to one contagion channel can no longer be lost to another contagion channel. As such, the interaction terms ensure that i's cumulative losses over the rounds can never exceed the initial asset values. We apply the counterparty risk contagion channel first and the shareholder channel second, which is why the interaction term $\epsilon^f_{kj,n+1}$ already reflects k's loss in round n but $N_{kj,n}/N_{kj,1}$ does not. (The choice of which contagion channel to apply first is arbitrary as it does not affect our results.)

Tradable Securities

In informationally efficient markets, the price of a security reflects the security's supply and demand, as well as the performance of the security's issuer. Overlapping portfolio contagion depresses the value of the security through a sudden increase in the security's supply, whereas the counterparty risk and shareholder contagion channels depress the security's value based on the issuer's performance, as reflected in the losses the issuer suffers. We first explain how the overlapping portfolio contagion and shareholder contagion channels reduce the value of an equity share, after which we discuss how the value of a bond or MMI is depressed by overlapping portfolio contagion and counterparty risk contagion.

To determine the overlapping portfolio contagion affecting an equity share, let $\Delta r^\tau_{kj,n}$ denote the price impact of a security τ issued by k. Furthermore, we use

$$\mathcal{D}_{j,n} = \{k | B_{kj,n} = 0 \cap B_{kj,n-1} > 0\} \tag{29}$$

to denote the set of institutions who exhaust their buffer during round n, causing default. (Note that as an institution can only default once, it can only be included in $\mathcal{D}_{j,n}$ for one value of n.) We apply the linear price impact used in the liquidity weighted portfolio overlap (eq. 11), such that the liquidation of the portfolios of all in-default institutions $\mathcal{D}_{j,n}$ causes a price impact

$$\Delta r^\tau_{kj,n} = \sum_{q\in\mathcal{D}_{j,n}} \frac{w^\tau_{qk}}{D^\tau_k}, \tag{30}$$

where the market depth (eq. 12) is assumed to be constant and we only consider market depths large enough so the market price $P^\tau_{kj,n}$ remains positive. $r^\tau_{kj,n}$ denotes the "liquidity factor" of a tradable security $\tau \in \{b, m, e\}$ issued by institution $k \in \mathcal{H}$ and reflects the price impacts that the security accumulates over the n rounds. The liquidity factor is initially normalized to one ($r^\tau_{kj,1} = 1$) and is depressed by the price impacts accumulated over the rounds according to:

$$r^\tau_{kj,n+1} = r^\tau_{kj,n} - \Delta r^\tau_{kj,n}. \tag{31}$$

Absent counterparty risk or shareholder contagion, the liquidity factor $r^{\tau}_{kj,n}$ gives the market price $P^{\tau}_{kj,n}$ in round n as fraction of initial market price $P^{\tau}_{kj,1}$ (which is normalized to 1 ZAR).

Up until now, we have discussed the price impact in absence of other contagion channels[28]. We now consider the combined impact of overlapping portfolio and shareholder contagion on an equity share. Using that the share's book value reflects the losses that the issuer has accumulated over the rounds, and the liquidity factor the accumulated price impacts, we assume that the market price of an equity share e issued by k in round n is given by

$$P^{e}_{kj,n} = \frac{e_{kj,n}}{e_{kj,1}} r^{e}_{rj,n} P^{\tau}_{kj,1}, \tag{32}$$

The first term on the right hand side gives the share's book value in round n relative to its initial book value, and the liquidity factor $r^{e}_{rj,n}$ reflects that overlapping portfolio contagion may drive the market value of the shares below their book value. Consequently, the value $v^{e}_{ikj,n}$ of institution i's equity shares in k in round n is given by

$$v^{\tau}_{ikj,n} = \frac{P^{e}_{kj,n}}{P^{e}_{kj,1}} w^{e}_{ik} = \epsilon^{e}_{kj,n} r^{e}_{kj,n} w^{e}_{ik}, \tag{33}$$

where $P^{e}_{kj,n}/P^{e}_{kj,1}$ gives the shares' market price in round n as a fraction of their initial market price, and we have used that $\epsilon^{e}_{kj,n} = e_{kj,n}/e_{kj,1}$ (see equations 21 and 22).

Let us now discuss how bonds and MMIs are affected by overlapping portfolio contagion and counterparty risk contagion. Similar to the market price of equity shares, we assume that the market price of bonds or MMIs $\tau \in \{b, m\}$ issued by k is given by

$$P^{\tau}_{kj,n} = \epsilon^{\tau}_{kj,n} r^{\tau}_{rj,n} P^{\tau}_{kj,1}, \tag{34}$$

which reflects that the market price of bonds and MMIs is depressed by both counterparty risk contagion (through the risk-adjustment factor $\epsilon^{\tau}_{kj,n}$) and overlapping portfolio contagion (through the liquidity factor $r^{\tau}_{rj,n}$). Hence, the value $v^{\tau}_{ikj,n}$ of institution i's bonds or MMIs τ in securities-issuing institution k in round n is given by

$$v^{\tau}_{ikj,n} = \frac{P^{\tau}_{kj,n}}{P^{\tau}_{kj,1}} w^{\tau}_{ik} = \epsilon^{\tau}_{kj,n} r^{\tau}_{kj,n} w^{\tau}_{ik} \tag{35}$$

Comparing equations (35) and (33) shows that the values of all tradable securities are modeled equivalently, so the corresponding losses are too. Note that this is only the case because of our particular choice of risk-adjustment factor (eq. 22) and is not true in general. The loss that i suffers on its tradable securities of type $\tau \in \{b, m, e\}$ in round $n+1$ after j's default is given by the decease in the securities' value due to overlapping portfolio contagion and counterparty

[28]For simplicity, we have defined the price impact in absence of other contagion channels. Consequently, when counterparty risk or shareholder contagion depress the market price, the effect of overlapping portfolio contagion on the market price is smaller than the price impact. This is captured by the interaction term $\epsilon^{\tau}_{kj,n+1}$ in equation 36

risk or shareholder contagion;

$$L^{\tau}_{ij,n+1} = \sum_{k\in\mathcal{H}} v^{\tau}_{ikj,n} - v^{\tau}_{ikj,n+1} = \sum_{k\in\mathcal{H}} \left(\Delta\epsilon^{\tau}_{kj,n} r^{\tau}_{kj,n} + \epsilon^{\tau}_{kj,n+1}\Delta r^{\tau}_{kj,n}\right) w^{f}_{ik}, \tag{36}$$

where we have used equations (23), (31), and (35).

The first term within the brackets on the right hand side of equation (36) gives the counterparty risk or shareholder contagion that i suffers due to j's buffer loss $\Delta\epsilon^{\tau}_{kj,n}$, multiplied by an interaction term; the liquidity factor $r_{kj,n}$. Similarly, the second term gives the overlapping portfolio contagion that i suffers due to price impact $\Delta r_{kj,n}$, multiplied by another interaction term: the risk-adjustment factor $\epsilon^{\tau}_{kj,n+1}$. As in equation 28, the interaction terms reflect that asset values already lost to one contagion channel can no longer be lost to another contagion channel, and ensure that i's cumulative losses over the rounds cannot exceed the initial asset values. We apply the counterparty risk shareholder contagion channel first and the overlapping portfolio contagion channel second, which is why the interaction term $\epsilon^{\tau}_{kj,n+1}$ already reflects k's loss in round n but $r_{kj,n}$ does not. This completes the calculation of the losses $L^{\alpha}_{ij,n+1}$ in equation (16) and, accordingly, the n^{th}-order exposures in equation (15).

5 Results

Our results provide the following key insights into higher-order exposures and their measurement in the case study of the South African financial system:

1. *Higher-order exposures in the South African financial system are substantial.* Whereas previous literature concluded that ignoring indirect exposures may lead to a significant underestimation of exposure (Cont and Schaanning, 2019), we show that even when capturing both direct and indirect losses (referred to as "first-order exposures") exposures are still underestimated. Due to higher-order exposures, institutions may be materially exposed even to banks to which they have no first-order exposures whatsoever.

2. *Higher-order exposures must be modelled explicitly at the level of individual institutions.* We show that higher-order exposures cannot be extrapolated from simpler proxies, due to the complex nature of the financial system that generates the exposures. Accurate calibration also requires contract-level data on the institutions' assets and liabilities, as reconstructing exposures from aggregate data may be highly inaccurate.

3. *Higher-order exposures are largest when they matter most; during financial crises.* Exposures are most relevant during times of financial distress, as that is when defaults are mostly likely to occur, causing the exposures to materialize into losses. We show that during financial crises, higher-order exposures become particularly pronounced and, in most cases, dominate first-order exposures (which are not exacerbated by the crisis).

These results are derived in the following sections. We discuss system-wide exposures, sectoral exposures, individual institutions' exposures, and stressed exposures, in that order. Each exposure depicted below is the mean exposure calculated over 1000 realizations of the reconstructed interbank network. The number of samples is sufficiently large that the standard errors of the results are negligible. As explained in previous sections, our market depth and risk-adjustment estimates are probably conservative, so the results in this section most likely provide an upper bound to the true exposures during normal times. Exposures during times of distress in the financial system are discussed in section 5.4.

5.1 System-Wide Exposures

Figure 7 shows the exposures, as % of the system's total assets, and HSEs up to n^{th} order of the South African financial system to the default of bank j, where j is one of the six large banks. The system's exposure up to n^{th} order is calculated as the sum of the exposures of all banks' and funds' exposures up to n^{th} order. The figure shows that higher-order exposures are substantial, in particular in comparison to the first-order exposures. This is highlighted by the HSEs plotted in Figure 7b, which show that exposures to Absa and Capitec may be underestimated by over 50% when ignoring higher-order exposures. Of the higher-order exposures, the second-order exposure is the most severe, as can be seen from the substantial jump in exposure from $n = 1$ to $n = 2$, and the higher-order exposures level out as $n \to 5$. (Note that the HSE up to first order is zero by definition.)

The six banks appear in the legend in descending order of total asset size. As this ordering is not reflected in the exposures, exposures to a counterparty cannot be inferred from the counterparty's total asset size alone. Furthermore, note that the first-order exposure to Absa is slightly smaller than Nedbank's, while the exposure up to fifth order to Absa is almost as Standard bank's. This shows that these higher-order exposures are qualitatively different and cannot simply be "extrapolated" from first-order exposures.

Both figures are limited to $n \leq 5$. The reason for this is threefold. First, as noted in Section 4.4, any model has finite accuracy, and because inaccuracies compound, the accuracy of the higher-order exposures worsens for larger n. As shown in Figure 17a in the appendix, the distribution of exposures, resulting from the random realizations of the reconstructed interbank network, fans out to quite substantial levels by $n = 5$. As we do not know the true interbank network, the true exposures may lie anywhere within this distribution. Second, exposures and HSEs up to n^{th} order level out for large n. The vast majority of higher-order exposures materialize as second and third order exposure. This also suggests that the vast majority of higher-order exposures accumulate when our confidence in the accuracy of the model is greatest. Third, when losses are particularly large and percolate through the financial system for an increasingly large number of rounds, it becomes more likely that the central bank will intervene to stabilize the system. Hence, large exposures of very high order are unlikely to materialize.

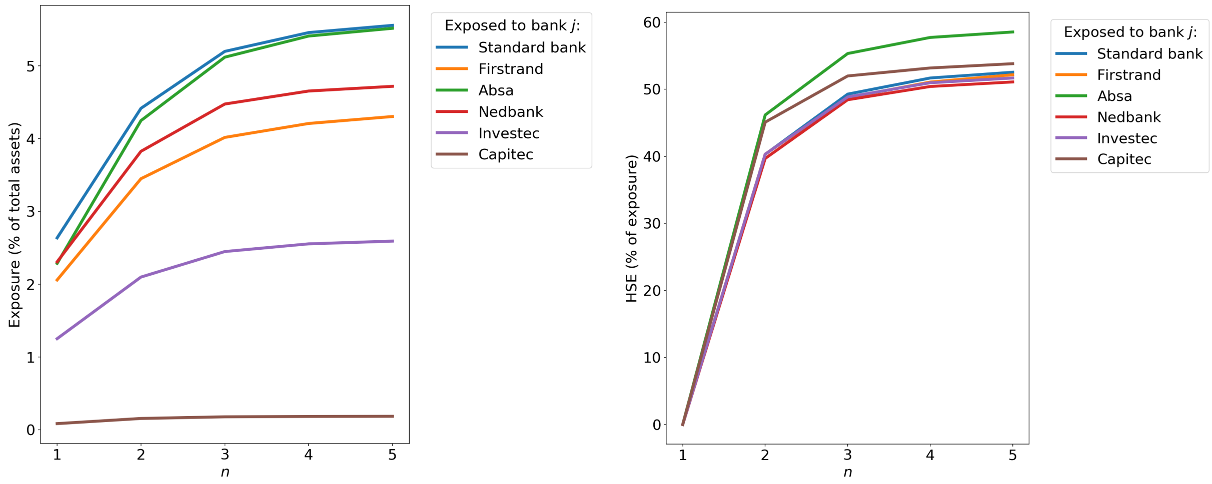

(a) Exposure up to n^{th} order

(b) HSE up to n^{th} order

Figure 7: **Exposure and higher-order share of exposure (HSE) up to n^{th} order of the South African financial system to the six largest banks.** Plot (a) shows the exposure (as % of the system's total assets) up to n^{th} order of the South African financial system to the default of bank j, where j is one of the six large banks. (Note that the six banks appear in the legend in descending order of total asset size.) The system's exposure is calculated as the sum of all banks' and funds' exposures. The exposure increases substantially from $n = 1$ to $n = 2$, which corresponds to the second-order exposures. Further increases in exposure level out as n approaches 5. Plot (b) shows HSEs up to n^{th} order, which are substantial for all $n > 1$. (Note that the HSE up to order $n = 1$ is zero by definition.) In particular, the HSEs indicate that exposures are underestimated by more than 50% when ignoring higher-order exposures.

5.2 Sectoral Exposures

Figures 8-11 break down the exposures in Figure 7 by sector. Figure 8 shows the exposures and HSEs of the banking sector to the six largest banks, Figure 9 those of the MMF sector, Figure 10 those of the OF sector and Figure 11 those of the FoF sector.

Figure 8 shows that exposures of the banking sector to the six large banks are small, but that the HSE is substantial (above 25 percent) in all cases. This demonstrates that banks in South Africa have significant higher-order exposures to one another. Furthermore, note that Figure 8 looks qualitatively different from Figure 7, as the vertical ordering of exposures to the six banks is different across figures. This is also reflected in Figures 9-11. Hence, the sectoral exposures vary qualitatively across sectors and cannot by proxied by the system-wide exposures.

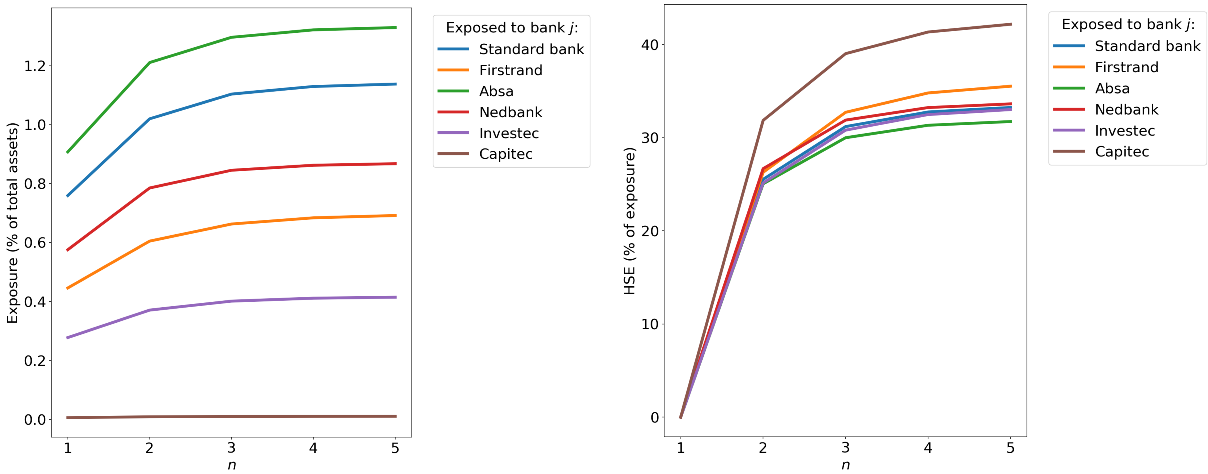

(a) Exposure up to n^{th} order

(b) HSE up to n^{th} order

Figure 8: **Exposure and HSE up to n^{th} order of the banking sector to the six largest banks.** Plot (a) shows the exposure (as % of the sector's total assets) up to n^{th} order of the South African banking sector to the default of bank j, where j is one of the six large banks. The sector's exposure is the sum of the banks' exposures. Plot (b) shows the corresponding HSE up to n^{th} order. Exposures of the banking sector to the largest six banks are small, yet the HSE of the exposures is substantial.

The MMF sector makes large investments in South Africa's six largest banks, which is reflected in the substantial first-order exposures shown in Figure 9. Although the sector's first-order exposures to all six banks remain below the 30% large exposure limit, the higher-order exposures push the exposure to Absa, Nedbank and Standard bank beyond this limit. Note that regulation only requires direct exposures to fall below 30% large exposure limit, but higher-order exposures should arguably be recognized too. Capitec shows a particularly large HSE, but the MMF sector's higher-order exposures to Capitec are negligible in absolute terms.

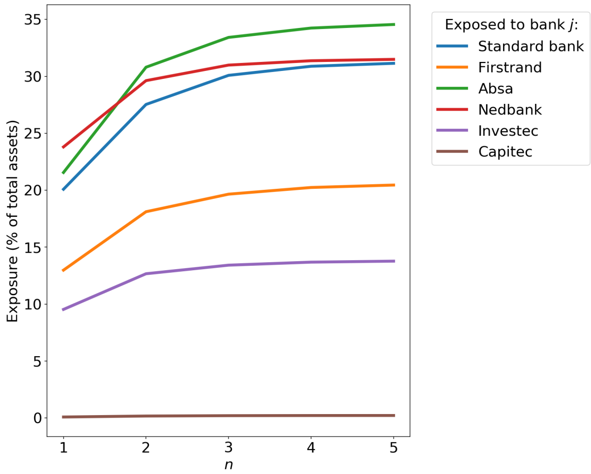

(a) Exposure up to n^{th} order

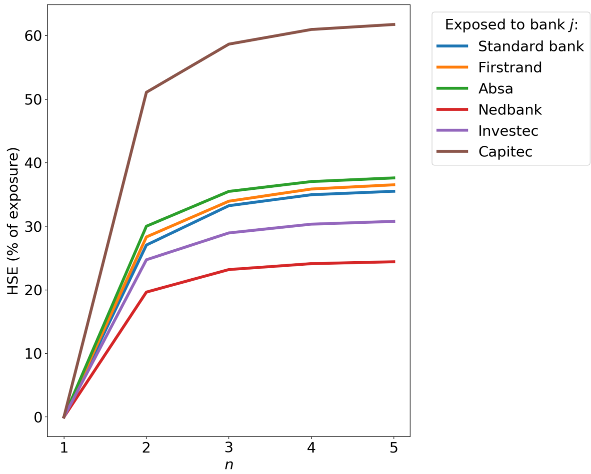

(b) HSE up to n^{th} order

Figure 9: **Exposure and HSE up to n^{th} order of the MMF sector to the six largest banks.** Plot (a) shows the exposure (as % of the sector's total assets) up to n^{th} order of the South African MMF sector to the default of bank j, where j is one of the six large banks. The sector's exposure is the sum of the MMFs' exposures. Plot (b) shows the corresponding HSE up to n^{th} order. The MMF sector has substantial first-order exposures to the largest six banks, as banks are the main recipients of MMFs' investments. Moreover, the higher-order exposures push the MMF sector's exposures to Absa, Nedbank and Standard bank beyond the regulatory limit of 30%.

Figure 10 shows that exposures of the OF sector to the six large banks are smaller than those of MMF sector. However, HSEs to all banks but Capitec are substantially larger. Ignoring the higher-order exposures to Absa would be particularly problematic and underestimate exposure by more than 60%.

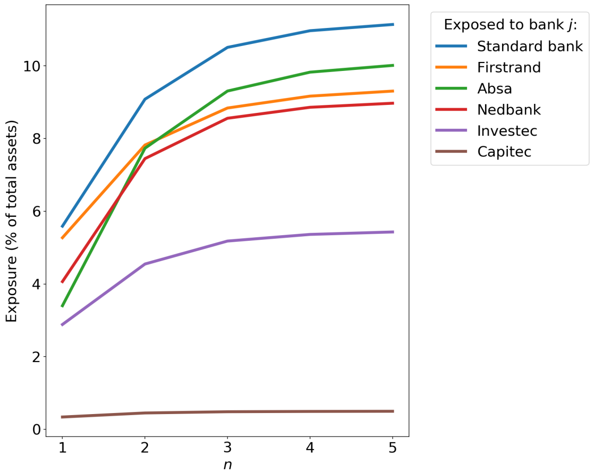

(a) Exposure up to n^{th} order

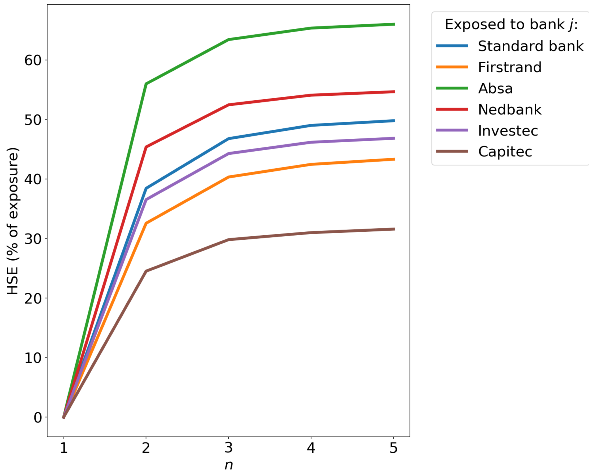

(b) HSE up to n^{th} order

Figure 10: **Exposure and HSE up to n^{th} order of the OF sector to the six largest banks.** Plot (a) shows the exposure (as % of the sector's total assets) up to n^{th} order of the OF sector to the default of bank j, where j is one of the six large banks and the sector's exposure is the sum of the OFs' exposures. (b) shows the corresponding HSE up to n^{th} order. Exposures of the OF sector to the six largest banks are smaller than those of the MMFs but HSEs are vast, reaching up to 60% in the case of Absa. Hence, OFs are expected to underestimate their exposure to Absa by more than 60% when only taking first-order exposures into account.

Funds of funds predominantly invest in other funds. Accordingly, Figure 11 shows that the FoFs' first-order exposures to the banks are virtually non-existent and the HSEs are very close to 100%. As a result, conventional exposure metrics capturing direct and even indirect exposures would find no significant exposures, even though exposures can reach as high as 18% of the FoF sector's total assets in the case of Standard bank. Furthermore, while the FoFs' first-order exposures are much smaller than those of the OFs', the FoFs' higher-order exposures are much larger. This demonstrates that the FoFs' strategy of investing in other funds makes them particularly exposed to systemic risk.

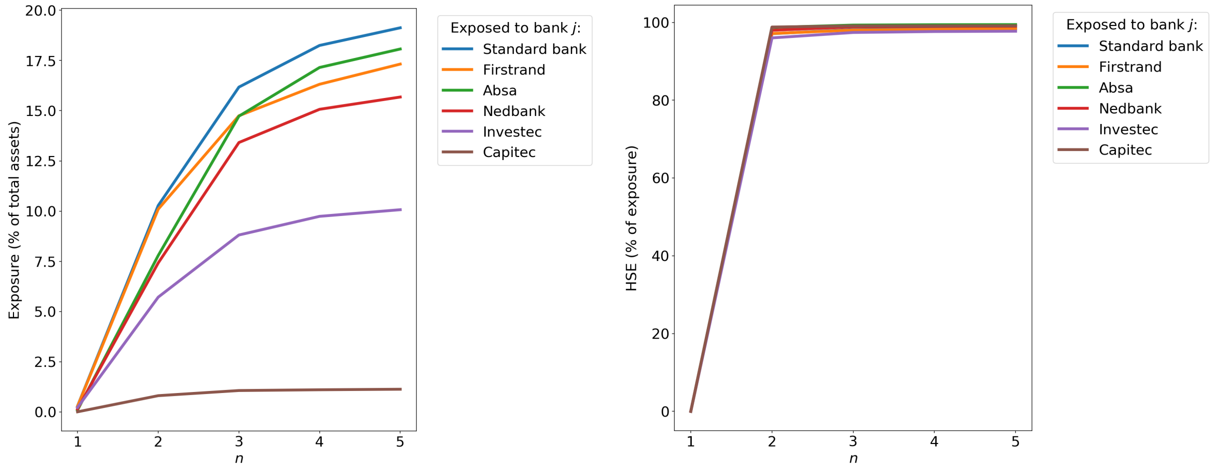

(a) Exposure up to n^{th} order (b) HSE up to n^{th} order

Figure 11: **Exposure and HSE up to n^{th} order of the FoF sector to the six largest banks.** Plot (a) shows the exposure (as % of the sector's total assets) up to n^{th} order of the FoF sector to the default of bank j, where j is one of the six large banks and the sector's exposure is the sum of the FoFs' exposures. (b) shows the corresponding HSE up to n^{th} order. FoFs' first-order exposures to the six largest banks are small, as FoFs typically invest in other funds. The FoF sector's higher-order exposures are substantial; exposures to the four largest banks are around 15% of the sector's total assets. Furthermore, because first-order exposures to the largest six banks are close to zero, the FoF sector's HSEs are close to 100%. Hence, FoFs completely overlook their exposure to these banks when not taking higher-order exposures into account.

5.3 Individual Exposures

We have already seen that higher-order exposures are qualitatively different from the first-order exposures at the system-wide and sectoral level. Here, we compare first-order and higher-order exposures of individual institutions. We focus on the MMFs' and OFs' exposures, because the banks' exposures are very modest in our data and the FoFs have no first-order exposures to compare the higher-order exposures to. We plot the individual institutions' HSEs to compare their first-order and higher-order exposures. We also plot the first-, second-, and fifth-order exposures to compare them individually. (We plot the second-order exposures because they are the largest of the higher-order exposures, and the fifth-order exposures because they are the highest-order exposures we measure.)

Figure 12 shows the HSEs and exposures, as % of the fund's total assets, of the ten largest MMFs and Figure 13 those of the ten largest OFs. The figures use a color gradient to indicate the magnitude of the exposure of a fund on the vertical axis to one of the six large banks on the horizontal axis. Note that the color gradients vary across plots, in order to accommodate the heterogeneity of exposure sizes.

We are primarily interested in variations in HSEs and exposures across the funds, i.e. along the columns. The MMFs' HSEs in Figure 12 show little variation; the HSEs to all banks but Capitec are quite modest. This homogeneity of exposures across MMFs is due to the MMFs' highly similar portfolios, as discussed in section 4.1. The only exception is MMF6, which has virtually no investment in Investec and a near 100% HSE to Investec accordingly. Furthermore, as we already saw in Figure 9, all HSEs to Capitec's are very high simply due to universally small first-order exposures.

Of the first-, second-, and fifth-order exposures, the first-order exposures show the most variation across MMFs, driven by whatever small differences exist between the MMFs' portolios. In the second-order exposures and, in particular, fifth-order exposures, most of the variation in exposures across the MMFs has vanished. Instead the exposures show a clear pattern of variation across the banks, i.e. along the rows. The pattern closely resembles the fifth-order exposures of the MMFs at the sectoral level (the fifth-order exposure can be inferred from the difference between the exposure up to fourth- and up to fifth order in Figure 9a). This suggests that for higher-order exposures, the systemic impact of the defaulting bank dominates the small variations across the MMFs' portfolios. However, as shown below, this is not true for all funds, but is specific to the MMFs and due to their exceptionally similar portfolios.

Contrary to the MMFs, Figure 13 shows that OFs' exposures vary substantially across the OFs. The figure also shows that the distribution of first-order exposures across the OFs is substantially different from the second-order exposures' distribution. This highlights that higher-order exposures cannot be extrapolated from first-order exposures. Furthermore, the variation of HSEs across the OFs also reflects that exposure up to n^{th} order cannot simply be estimated as a multiple of the first-order exposure. Notably, OF8 has some of the lowest first-order exposures, but highest second-order exposures of all ten OFs, highlighting the risk of attempting to extrapolate higher-order exposures from first-order exposures.

While most of the variation in exposures across the OFs' has vanished in the fifth-order exposures, the second-order exposures vary substantially OFs. Furthermore, because the second-order exposures dominate the higher-order component of the exposures, the HSEs also show strong heterogeneity across te OFs. Hence, while we already established that higher-order exposures cannot be extrapolated from first-order exposures, we conclude that individual institutions' higher-order exposures generally cannot be inferred from sectoral exposures either. MMFs are the exception, as their portfolios are so similar that individual institutions' higher-order exposures resemble the sector-level higher-order exposures. This conclusion is supported by Figure 18 and Figure 19 in the appendix, which show the exposures of the individual banks and FoFs.

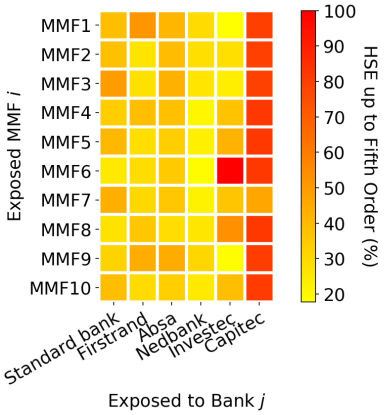

(a) HSE up to fifth-order

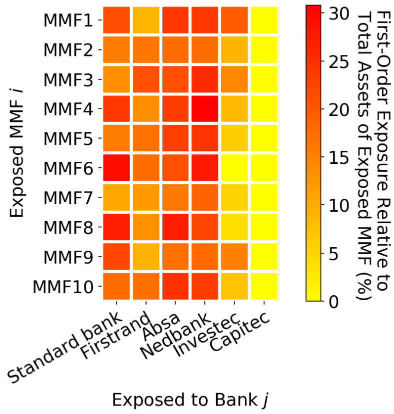

(b) First-order exposure

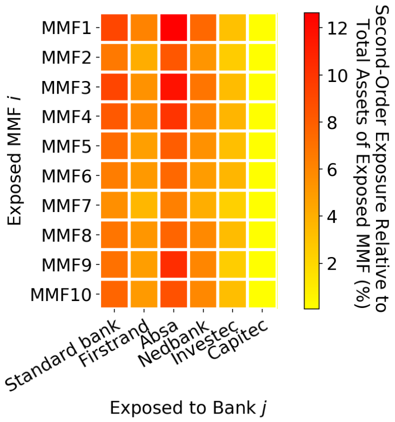

(c) Second-order exposure

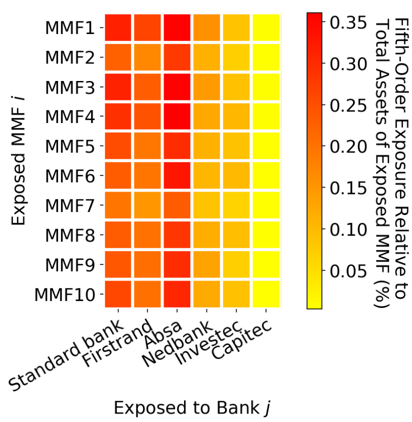

(d) Fifth-order exposure

Figure 12: **Individual MMFs' exposures to the six largest banks**. The figure shows the exposures, as % of total assets, and HSEs of the ten largest MMFs. The plots use a color gradient to indicate the exposure or HSE of an MMF on the vertical axis to a bank on the horizontal axis. Furthermore, the MMFs are numbered in descending order of total asset size and the banks are ordered from left to right by descending total asset size. The plots show little variation in HSEs and exposures across the MMFs, which is due to the MMFs' exceptionally similar portfolios. The first-order exposures (which are predominantly driven by direct exposures) in (b) show that all MMFs' investments in the four largest banks are close to the 30% limit, with Investec receiving the remainder of the investments and Capitec virtually nothing. Plots (c) and, in particular, (d) show that the higher the order of the exposure, the more the variation across the MMFs is damped out and gets dominated by the variation across the banks to which the MMFs are exposed.

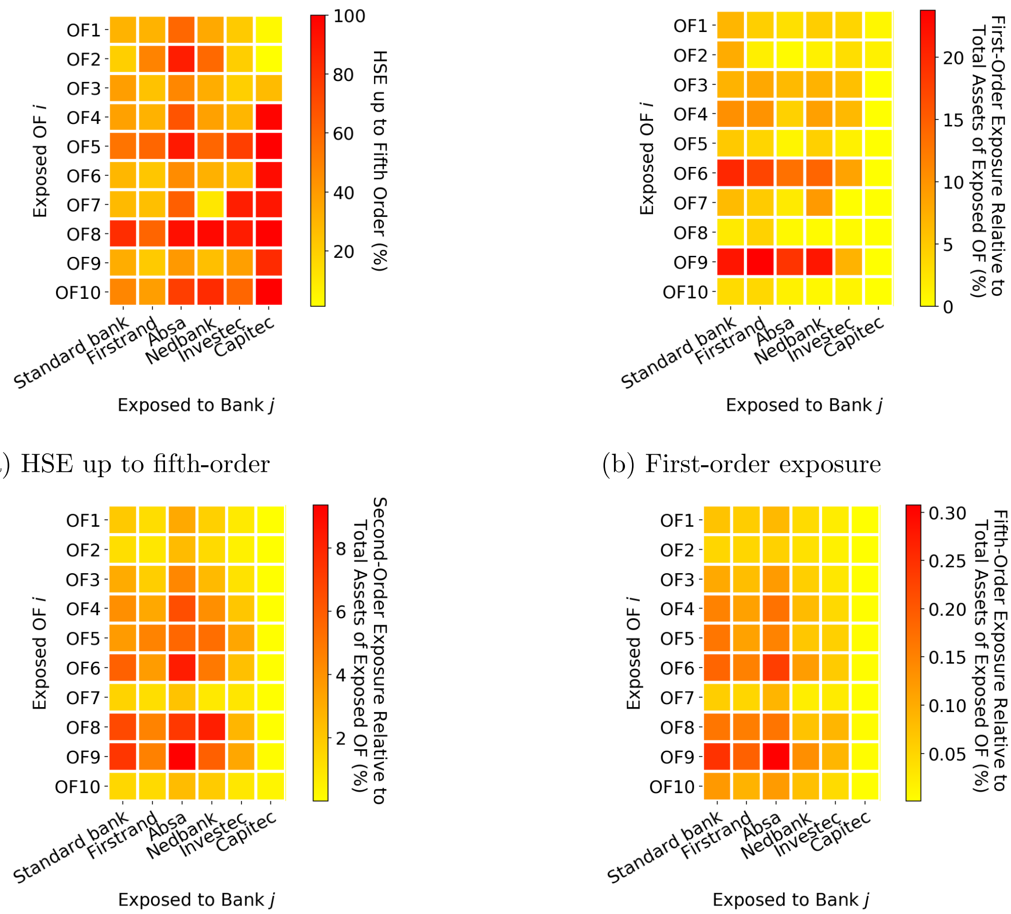

(a) HSE up to fifth-order

(b) First-order exposure

(c) Second-order exposure

(d) Fifth-order exposure

Figure 13: **Individual OFs' exposures to the six largest banks**. The figure shows the exposures, as % of total assets, and HSEs of the ten largest OFs. The plots use a color gradient to indicate the exposure or HSE of an OF on the vertical axis to a bank on the horizontal axis. Furthermore, the OFs are numbered in descending order of total asset size and the banks are ordered from left to right by descending total asset size. Plot (a) shows that HSEs vary strongly across OFs, and comparison of the first-order exposures in (b) to the second-order and fifth-order exposures in (c) and (d) clearly shows that the second-order and fifth-order exposures are qualitatively different from the first-order exposures. Hence, the higher-order exposures cannot be extrapolated from first-order exposures. Furthermore, similar to the MMFs, (d) shows that in the fifth order exposures, most of the variation across the OFs is damped out and gets dominated by the variation across the banks to which the OFs are exposed. However, this is not the case for the second-order exposures, which show substantial variation across the OFs.

5.4 Stressed Exposures

In this section, we vary institutions' buffers (which is analogous to varying the risk-adjustment factor) and securities' market depths, to understand how sensitive our results are to these parameters. As noted previously, our baseline market depth and risk-adjustment factor are probably conservative during benign times. The more extreme estimates of these parameters that we explore here are likely to arise only during severe crises, if at all. Nevertheless, these "stressed exposures" provide a useful illustration of how higher-order exposures are affected by rising credit risk adjustments and falling market depths during crises.

Figures 14, 15 and 16 show that higher-order exposures become particularly pronounced during times of financial distress. Put differently, higher-order exposures are at their greatest exactly when they matter most, i.e. in times of crisis when defaults are most likely to occur and, consequently, exposures are most likely to materialize into losses.

Due to adverse the macroeconomic conditions in a crisis scenario, institutions may incur unexpected losses, reducing their buffers. This exacerbates the counterparty risk channel and, in turn, pushes up higher-order exposures. Figure 14 shows the exposure up to second-order and up to fifth-order when all banks' and funds' (initial) buffers are reduced by 20%, 40% or 60%. For comparison's sake, the exposures are also shown for the case the institutions' buffers are not changed ("0%"), or increased by 20%[29]. The gray bars show the direct exposures, so all exposures exceeding the gray bars are overlooked when ignoring indirect and higher-order exposures exposures. (Note that while higher-order exposures are affected by buffers, direct and indirect exposures are not.) The figure shows that when buffers are reduced, higher-order exposures increase substantially and start to dominate exposure. In particular, (a) shows that exposures up to second order increase strongly when buffers fall by 60%. Furthermore, (b) shows that exposures up to fifth order to Standard Bank, Absa and Nedbank grow substantially when buffers are reduced by just 20%.

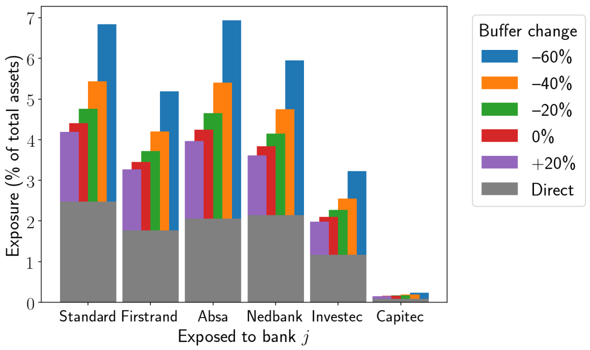

(a) Exposure up to second order

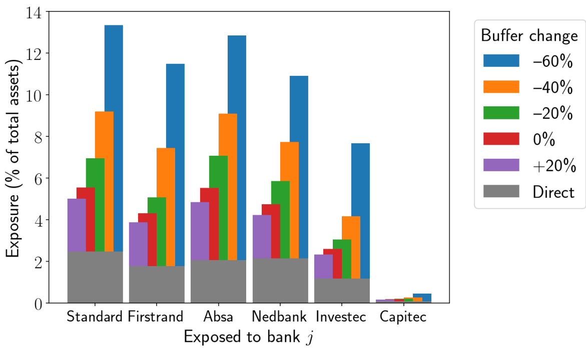

(b) Exposure up to fifth order

Figure 14: **Exposure of the South-African financial system to the six largest banks for various values of institutions' initial buffers.** (a) shows the exposure up to second order and (b) up to fifth order of the South-African financial system to the default of bank j, where j is one of the six large banks. The system's exposure is the sum of the banks' and funds' exposures, and is express as % of total bank and fund assets. The colors indicate the percentage change applied to all banks' and funds' initial buffers. The gray bars show the direct exposures, which are not affected by the buffers. The figures show that higher-order exposures become particularly pronounced when institutions' buffers are reduced, with exposures to Standard bank, Absa and Nedband increasing most strongly. More specifically, (b) shows that the higher-order component of exposures up to fifth order to Standard bank, Absa and Nedband increases by about 50% when initial buffers are reduced by just 20%.

Reduced buffers are not the only reason why higher-order exposures increase in times of financial distress. During crises, market liquidity for tradable securities $t \in \{b, m, e\}$ typically falls. In illiquid markets the demand for securities falls, which is reflected in reduced market

[29]Banks may increase their buffers by raising capital. As funds' buffers give the losses that funds can absorb before investors start running, funds increase buffers by raising investor confidence, e.g. through increasing funds' liquidity so the fund can repay redemptions without resorting to firesales.

depth, so the price impact of liquidating a defaulted institution's portfolio increases. Figure 15 shows the exposure up to second-order and up to fifth-order for various values of the cap-to-depth ratio μ. When the cap-to-depth ratio $\mu = 0$, market depths are infinite so the overlapping portfolio contagion channel is effectively turned off, whereas when the cap-to-depth ratio $\mu = 4$, market depths are reduced by 75% and the overlapping portfolio contagion channel is strongly amplified. The gray bars again show the direct exposures, which are not affected by the market depths, while the indirect and higher-order exposures are. The figure shows that higher-order exposures are exacerbated when market depths fall. However, within the parameter ranges explored, lowered market depths have less of an impact than diminished buffers.

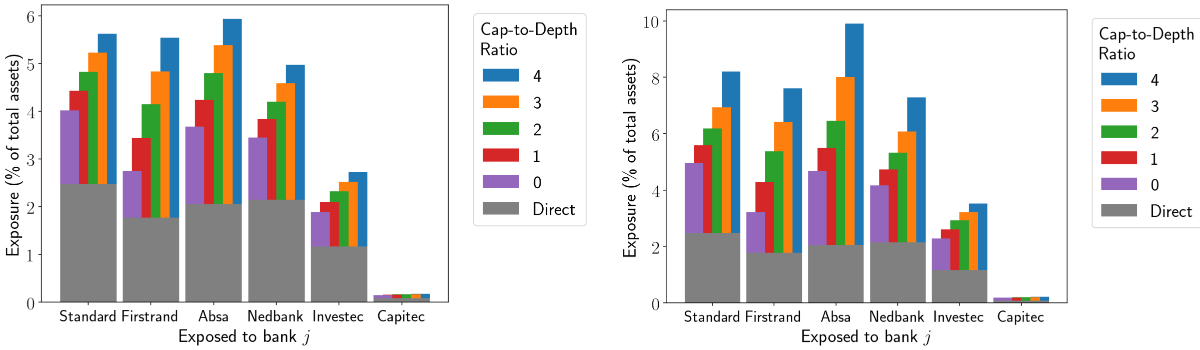

(a) Exposure up to second order

(b) Exposure up to fifth order

Figure 15: **Exposure of the South-African financial system to the six largest banks for various market depths.** (a) shows the exposure up to second order and (b) up to fifth order of the South-African financial system to the default of bank j, where j is one of the six large banks. The system's exposure is the sum of the banks' and funds' exposures, and is expressed as % of total bank and fund assets. The colors indicate the cap-to-depth ratio μ used for all securities. The gray bars show the direct exposures, which are not affected by market depths. The figures show that both the exposures up to second and fifth order to all banks but Capitec increase substantially when market depths fall.

Based on the SARB Financial Stability Review (SARB, 2016a), we formulate a severe macroeconomic stress scenario for illustrative purposes. The scenario consists of a 25% reduction in institutions' buffers and 50% reduction in the market depth for all tradable securities. Figure 16 shows that when all banks and funds are subjected to the stress scenario, higher-order exposures are substantially larger (compared to Figure 7). Furthermore, exposures no longer level out by $n = 5$, so the practice adopted in this paper of only considering exposure up to order $n = 5$ may still underestimate exposure during the worst of a crisis. This is confirmed by Figure 17 in the Appendix, which shows that exposures only level out by about $n \approx 8$ when institutions are subjected to the stress scenario.

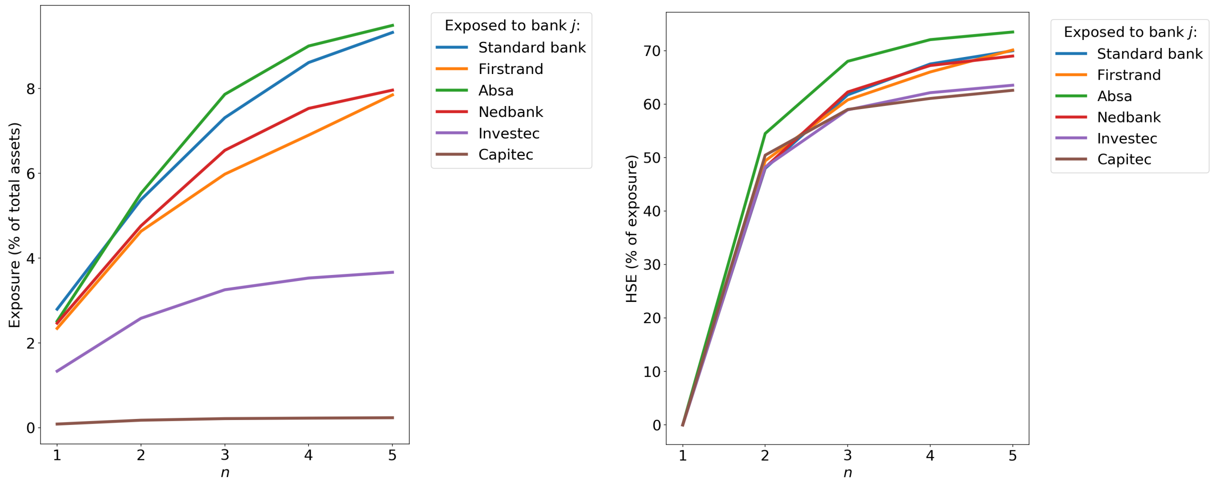

(a) Exposure up to n^{th} order

(b) HSE up to n^{th} order

Figure 16: **Stressed exposure and HSE up to n^{th} order of the South African financial system to the six largest banks.** All banks and funds are subjected to the macroeconomic stress scenario, which consists of a 25% reduction in the institutions' buffers and 50% reduction in the market depth of all tradable securities (i.e. for the cap-to-depth ratio $\mu = 2$). (a) shows the exposure (as % of the system's total assets) up to n^{th} order of the South African financial system to the default of bank j, where j is one of the six large banks and the system's exposure is the sum of the banks' and funds' exposures. (b) shows the corresponding HSE up to order n. Figure (a) shows that exposures no longer level out by $n = 5$. Moreover, figure (b) shows that the stressed HSEs exceed 60% across all six banks. Hence, higher-order exposures become particularly dominant during crises, which is exactly when exposures are most important (as that is when defaults are most likely to cause exposures to materialize into losses).

6 Discussion

In this paper we have introduced the concept of higher-order exposures and proposed a way to measure these. As shown by the findings of our case study of the South African financial system, the concept of higher-order exposures has substantial implications for prudential policymakers, regulators, and supervisors with financial stability mandates. We have shown that direct and indirect exposures, which are traditionally used to calculate exposures, only capture part of the exposures between financial institutions. Higher-order exposures can be significant, heterogeneous, and particularly high in times of crisis – when exposures matter most. Moreover, these exposures cannot easily be extrapolated from traditional measures of exposure and therefore require complementary analysis.

Analysis of higher-order exposures should inform the design and calibration of various tools in the regulatory arsenal. This essentially applies to all elements of the (macro)prudential toolkit where exposures matter. While it is beyond the scope of this paper to comprehensively discuss how each tool could be adjusted, we highlight the areas where incorporating higher-order exposures is especially salient:

1. *Large-exposure limits.* The risk of large losses associated with the demise of a single counterparty is not captured by the risk-based capital standards that apply to financial institutions. That is why many jurisdictions, following the Basel Committee on Banking

Supervision's recommendation in the Basel Accords (BIS, 2018b), have introduced large exposure regimes which limit a financial institution's exposure to any other institution or group of connected counterparties to a certain percentage (the details differ across jurisdictions) of its capital. For our purposes, the key point is that the measures of exposure that are used to calculate whether an institution violates the regime do not capture higher-order exposures. This not only implies that some exposures between institutions are not considered at all (i.e. those between institutions that are connected only through higher-order exposures), but also that institutions' exposures might unknowingly exceed the limit set by regulators. At a minimum, regulators should measure higher-order exposures in order to identify risks that would otherwise go unnoticed. More ambitiously, large exposure regimes could be overhauled to recognize the full extent of exposure.

2. *Capital requirements.* Capital requirements are generally calibrated to direct exposures only, thereby failing to account for the indirect and higher-order exposure. For Global Systemically Important Banks (G-SIBs) this omission is most striking. G-SIBs are required to hold additional capital (referred to as the "G-SIB surcharge") according to their level of systemic importance; G-SIBs that are more systemically important thus face a higher capital surcharge. The current Basel accords use an "indicator-based measurement approach" to measure a G-SIBs systemic importance, which quantifies systemic importance in terms of the impact that the bank's failure could have on the global financial system and the wider economy (BIS, 2018a). While this could be achieved through measuring the direct, indirect and higher-order exposure of the system to its failure, in practice five proxies for its systemic importance are used. One of these is "interconnectedness", which accounts for 20% of the bank's total score. Currently, interconnectedness is assessed through "intra-financial system assets and liabilities" and "securities outstanding", which strongly correlate to the size of the bank's balance sheet and are not sensitive to how portfolios overlap (Cont and Schaanning, 2019). Our results suggest that this approach may, in some relevant cases, significantly underestimate exposures and would thereby underestimate the systemic importance of the bank, resulting in a surcharge that is too low. We have shown that higher-order exposures are not necessarily correlated with balance sheet size, so extrapolating from the current indicators would not address the problem. Since this methodology is also used to assess the systemic importance of financial institutions more broadly, some systemically important institutions that generate large higher-order exposures may not be identified at all. Incorporating higher-order exposures in calibration exercises of the capital requirements could thus offer insights that are, at the very least, complementary to those obtained using existing methods – expanding the set of systemically important institutions and increasing the capital surcharge for some of them.

3. *Stress test models.* Today's regulatory stress tests are microprudential in nature. Microprudential stress tests are forward-looking exercises that assess the resilience of an institution (as e.g. measured by capital levels) to adverse economic conditions. While these tests measure asset losses resulting from direct exposures, they fail to measure the additional asset losses that could materialize through indirect and higher-order exposures. For stress tests to fulfil their basic function of assessing exposure and institutions' resilience

to risk, capturing higher-order exposures is essential. To be able to do that, stress test models should include multiple interacting contagion channels and be designed to study system-wide dynamics (Farmer et al., 2020), because those elements drive higher-order exposures. We have demonstrated that compensating for the lack of explicit system-wide models with direct loss multipliers is inadequate as it gives distorted outcomes. To better assess the resilience of financial institutions, stress tests should thus not only measure the capital impact of asset losses from direct exposures, but also from indirect and higher-order exposures.

4. *Resolution.* In the 2007-2008 financial crisis regulators stood before the terrible choice of either bailing-out a SIB or liquidating it in a potentially disorderly manner. Since then, new resolution regimes have been developed around the world that enable regulators to resolve banks in an orderly manner through a bail-in, thereby aiming to avoid undue disruption to the bank's activities and contagion effects to the rest of the economy. In a bail-in, the debt of the bank's creditors is written down and in part converted to equity to recapitalize it and revive its short-term viability. To decide whether a bank should be liquidated or bailed-in, authorities must determine that a failing bank cannot go through normal insolvency proceedings without harming public interest and causing financial instability (Kleinnijenhuis et al., 2021). To make this call, regulators must assess the stability implications of a bank failure. This requires a measurement of the losses that the system could suffer if the bank were liquidated. Put differently, it precisely requires the regulator to measure first-order and higher-order exposure of other institutions to the bank's failure. Unfortunately, a measure of higher-order exposure is completely missing in this toolkit.

The overarching takeaway is that, without explicitly capturing higher-order exposures, regulators and supervisors are flying blind. That leaves them ill-equipped to assess the resilience of the financial system they oversee, and ill-prepared to respond to crises once they inevitably materialise. Our results on the South African financial system illustrate these risks. To state the obvious, these results generalise to other jurisdictions – and so does their policy relevance. Conceptually, higher-order exposures fit neatly within the trend towards increasingly widely-adopted macroprudential regulation (Aymanns et al., 2018), combining the structural and network-sensitive (Enriques et al., 2019) and time-variant elements of such policies (Armour et al., 2016).

To operationalize the analysis of higher-order exposures, data-gathering mandates should cover more granular data across a wider range of financial institutions. Because higher-order exposures can only be quantified with contagion models that explicitly capture the multi-layered financial network, it is important that regulators and supervisors have the requisite data, i.e.: bilateral, contract-level data on individual institutions' assets and liabilities. The width of the distributions of exposures that we find (resulting from the random realisations of the reconstructed interbank network) highlight the vital importance of having insight into the financial system's network structure. The observation that substantial higher-order exposures may exist between seemingly unconnected parts of the financial system suggests that the scope of data-gathering mandates should be wide, spanning the entire financial system and potentially parts of the real economy [30]. To support this point, we note that without combining our data on the investment fund sector with that on the banking sector, we would not have been able to make the observation that the fund-of-fund sector is highly exposed to the banking sector in South-Africa, even though it has virtually no investments in it.

While this paper has provided a proof of principle for how the concept of higher-order exposures can be measured in practice using South-Africa as a case study, our method should by no means be seen as the gold standard. Regulators may want to consider various nuances, adjustments and extensions to the model, as well as to calibrate it more carefully to data. Depending on the financial system that is being studied, a different set of interconnections and associated contagion mechanisms might be relevant for higher-order exposures. Even if the same contagion mechanisms apply as in this study, they could be modeled differently. The counterparty risk contagion channel could be modeled with a different risk-adjustment rule and a different failure regime. Rather than assuming zero recovery in the short run on direct exposures to failed banks as we do (in line with Elsinger et al. [2006]), a regulator could capture that SIBs will likely be resolved (e.g. via a bail-in) while non-SIBs will be liquidated. Making this distinction has implications for the LGD that will apply in the counterparty risk contagion mechanism (Kleinnijenhuis et al., 2021). It also has implications for the overlapping portfolio contagion channel. Failed banks that are bailed-in do not have to liquidate their assets, potentially at discounted prices, while liquidated institutions do. Furthermore, overlapping portfolio contagion could be modeled using a more accurate price impact function. To be useful for regulatory purposes, our model could also be better calibrated to the prevailing and stressed market depths of tradable securities. It could also use more sophisticated stress test scenarios to determine stressed exposures.

[30] See e.g. Farmer et al., 2021, Ullersma and van Lelyveld, 2021, Sydow et al., 2024

References

D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015.

V. V. Acharya and D. Skeie. A model of liquidity hoarding and term premia in inter-bank markets. *Journal of Monetary Economics*, 58(5):436–447, 2011.

V. V. Acharya, L. H. Pedersen, T. Philippon, and M. Richardson. Measuring systemic risk. *The review of financial studies*, 30(1):2–47, 2017.

T. Adrian and M. K. Brunnermeier. Covar. Technical report, National Bureau of Economic Research, 2011.

T. Adrian and H. S. Shin. Liquidity and leverage. *Journal of financial intermediation*, 19(3): 418–437, 2010.

D. Aikman, J. Bridges, A. Kashyap, and C. Siegert. Would macroprudential regulation have prevented the last crisis? *Journal of Economic Perspectives*, 33(1):107–30, 2019.

I. Aldasoro, A.-C. Hüser, and C. Kok. Contagion accounting. 2020.

F. Allen and A. Babus. Networks in finance. *The network challenge: strategy, profit, and risk in an interlinked world*, 367, 2009.

F. Allen and D. Gale. Financial contagion. *Journal of political economy*, 108(1):1–33, 2000.

E. I. Altman and A. Saunders. Credit risk measurement: Developments over the last 20 years. *Journal of banking & finance*, 21(11-12):1721–1742, 1997.

M. Andrianaivo and C. A. Yartey. Understanding the growth of african financial markets. *African Development Review*, 22(3):394–418, 2010.

J. Armour, D. Awrey, P. L. Davies, L. Enriques, J. N. Gordon, C. P. Mayer, and J. Payne. *Principles of financial regulation*. Oxford University Press, 2016.

C. Aymanns, J. D. Farmer, A. M. Kleinnijenhuis, and T. Wetzer. Models of financial stability and their application in stress tests. *Handbook of Computational Economics*, 4:329–391, 2018.

M. Bardoscia, S. Battiston, F. Caccioli, and G. Caldarelli. Pathways towards instability in financial networks. *Nature Communications*, 8(1):1–7, 2017.

M. Baron, E. Verner, and W. Xiong. Banking crises without panics. *The Quarterly Journal of Economics*, 136(1):51–113, 2021.

J. Begenau, M. Piazzesi, and M. Schneider. Banks' risk exposures. Technical report, National Bureau of Economic Research, 2015.

B. S. Bernanke. Non-monetary effects of the financial crisis in the propagation of the great depression. Technical report, National Bureau of Economic Research, 1983.

BIS. Global systemically important banks: revised assessment methodology and the higher loss absorbency requirement. Technical report, Bank of International Settlements, 2018a.

BIS. The treatment of large exposures in the basel capital standards - executive summary. Technical report, Bank of International Settlements, 2018b.

BIS. Rbc20 calculation of minimum risk-based capital requirements. Technical report, Bank of International Settlements, 2019.

C. Bluhm, L. Overbeck, and C. Wagner. *Introduction to credit risk modeling*. Chapman and Hall/CRC, 2016.

F. S. Board. Board notice 90 of 2014. Technical report, Financial Sector Conduct Authority (FSCA), 2014.

C. Brownlees and R. F. Engle. Srisk: A conditional capital shortfall measure of systemic risk. *The Review of Financial Studies*, 30(1):48–79, 2017.

M. K. Brunnermeier. Deciphering the liquidity and credit crunch 2007-2008. *Journal of Economic perspectives*, 23(1):77–100, 2009.

M. K. Brunnermeier and M. Oehmke. Bubbles, financial crises, and systemic risk. *Handbook of the Economics of Finance*, 2:1221–1288, 2013.

M. K. Brunnermeier and L. H. Pedersen. Market liquidity and funding liquidity. *Review of Financial studies*, 22(6):2201–2238, 2009.

O. Burrows, D. Learmonth, J. McKeown, and R. Williams. Ramsi: a top-down stress-testing model developed at the bank of england. *Bank of England Quarterly Bulletin*, page Q3, 2012.

R. J. Caballero and A. Simsek. Fire sales in a model of complexity. *The Journal of Finance*, 68 (6):2549–2587, 2013.

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

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

S. Calimani, G. Hałaj, and D. Żochowski. Simulating fire-sales in a banking and shadow banking system. Technical report, ESRB Working Paper Series, 2017.

R. Cifuentes, G. Ferrucci, and H. S. Shin. Liquidity risk and contagion. *Journal of the European Economic association*, 3(2-3):556–566, 2005.

L. Clerc, A. Giovannini, S. Langfield, T. Peltonen, R. Portes, and M. Scheicher. Indirect contagion: the policy problem. Technical report, ESRB Occasional Paper Series, 2016.

R. Cont and E. Schaanning. Fire sales, indirect contagion and systemic stress testing. 2017.

R. Cont and E. Schaanning. Monitoring indirect contagion. *Journal of Banking & Finance*, 104: 85–102, 2019.

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

R. Cont, A. Moussa, et al. Network structure and systemic risk in banking systems. *Edson Bastos e, Network Structure and Systemic Risk in Banking Systems (December 1, 2010)*, 2010.

R. Cont, A. Kukanov, and S. Stoikov. The price impact of order book events. *Journal of financial econometrics*, 12(1):47–88, 2014.

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

M. Crouhy, D. Galai, and R. Mark. A comparative analysis of current credit risk models. *Journal of Banking & Finance*, 24(1-2):59–117, 2000.

F. Duarte and T. M. Eisenbach. Fire-sale spillovers and systemic risk. *The Journal of Finance*, 76(3):1251–1294, 2021.

D. Duffie and K. J. Singleton. *Credit risk: pricing, measurement, and management.* Princeton university press, 2012.

A. Dufour and R. F. Engle. Time and the price impact of a trade. *The Journal of Finance*, 55 (6):2467–2498, 2000.

L. Eisenberg and T. H. Noe. Systemic risk in financial systems. *Management Science*, 47(2): 236–249, 2001.

Z. Eisler, J.-P. Bouchaud, and J. Kockelkoren. The price impact of order book events: market orders, limit orders and cancellations. *Quantitative Finance*, 12(9):1395–1419, 2012.

M. Elliott, B. Golub, and M. O. Jackson. Financial networks and contagion. *The American economic review*, 104(10):3115–3153, 2014.

H. Elsinger, A. Lehar, and M. Summer. Risk assessment for banking systems. *Management science*, 52(9):1301–1314, 2006.

L. Enriques, A. Romano, and T. Wetzer. Network-sensitive financial regulation. *J. Corp. L.*, 45:351, 2019.

J. D. Farmer, A. M. Kleinnijenhuis, P. Nahai-Williamson, and T. Wetzer. Staff working paper no. 861 foundations of system-wide financial stress testing with heterogeneous institutions. 2020.

J. D. Farmer, A. M. Kleinnijenhuis, and T. Wetzer. Stress testing the financial macrocosm. *Handbook of Financial Stress Testing (forthcoming)*, 2021.

X. Freixas, L. Laeven, and J.-L. Peydró. *Systemic risk, crises, and macroprudential regulation.* Mit Press, 2015.

C. Fricke and D. Fricke. Vulnerable asset management? the case of mutual funds. *Journal of Financial Stability*, 52:100800, 2021.

FSB. Global shadow banking monitoring report 2017. Technical report, Financial Stability Board, 2018.

G. Fukker, M. Kaijser, L. Mingarelli, and M. Sydow. *Contagion from market price impact: a price-at-risk perspective.* Number 2692. ECB Working Paper, 2022.

P. Gai and S. Kapadia. Contagion in financial networks. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, page rspa20090410. The Royal Society, 2010.

P. Gai, A. Haldane, and S. Kapadia. Complexity, concentration and contagion. *Journal of Monetary Economics*, 58(5):453–470, 2011.

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

P. Glasserman and H. P. Young. Contagion in financial networks. *Journal of Economic Literature*, 54(3):779–831, 2016.

I. Goldstein, H. Jiang, and D. T. Ng. Investor flows and fragility in corporate bond funds. *Journal of Financial Economics*, 126(3):592–613, 2017.

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

A. G. Haldane and R. M. May. Systemic risk in banking ecosystems. *Nature*, 469(7330):351–355, 2011.

M. Hoerova, C. Holthausen, F. Heider, et al. Liquidity hoarding and interbank market spreads: the role of counterparty risk. In *2009 Meeting Papers*, number 929. Society for Economic Dynamics, 2009.

IFRS. Ifrs 9 financial instruments. Technical report, International Financial Reporting Standards Foundation, 2021.

O. Jeanne and A. Korinek. Macroprudential regulation versus mopping up after the crash. *The Review of Economic Studies*, 87(3):1470–1497, 2020.

P. Jorion and G. Zhang. Credit contagion from counterparty risk. *The Journal of Finance*, 64 (5):2053–2087, 2009.

E. Kemp. Measuring shadow banking activities and exploring its interconnectedness with banks in south africa. *SARB Occassional Papers*, 2017.

A. M. Kleinnijenhuis, C. Goodhart, and J. D. Farmer. Systemic implications of the bail-in design. Technical report, Forthcoming Working Paper, 2021.

C. Kok and M. Montagna. Multi-layered interbank model for assessing systemic risk. Technical report, European Central Bank, 2013.

H. Kraft and M. Steffensen. Bankruptcy, counterparty risk, and contagion. *Review of Finance*, 11(2):209–252, 2007.

A. Krishnamurthy. Amplification mechanisms in liquidity crises. *American Economic Journal: Macroeconomics*, 2(3):1–30, 2010.

F. Lillo, J. D. Farmer, and R. N. Mantegna. Master curve for price-impact function. *Nature*, 421(6919):129–130, 2003.

A. Persaud. Macro-prudential regulation. 2009.

S. Poledna, S. Martínez-Jaramillo, F. Caccioli, and S. Thurner. Quantification of systemic risk from overlapping portfolios in the financial system. *Journal of Financial Stability*, 52:100808, 2021.

M. Potters and J.-P. Bouchaud. More statistical properties of order books and price impact. *Physica A: Statistical Mechanics and its Applications*, 324(1-2):133–140, 2003.

C. M. Reinhart and K. S. Rogoff. This time is different: A panoramic view of eight centuries of financial crises. Technical report, National Bureau of Economic Research, 2008.

SARB. Financial stability review, second edition 2016. Technical report, South African Reserve Bank, 2016a.

SARB. Banks ba900 economic returns. Technical report, South African Reserve Bank, 2016b.

SARB. Annual report 2016. Technical report, South African Reserve Bank, 2017a.

SARB. Financial stability review first edition. Technical report, South African Reserve Bank, 2017b.

A. Shleifer and R. Vishny. Fire sales in finance and macroeconomics. *Journal of Economic Perspectives*, 25(1):29–48, 2011.

M. Sydow, A. Schilte, G. Covi, M. Deipenbrock, L. Del Vecchio, P. Fiedor, G. Fukker, M. Gehrend, R. Gourdel, A. Grassi, et al. Shock amplification in an interconnected financial system of banks and investment funds. *Journal of Financial Stability*, 71:101234, 2024.

C. Ullersma and I. van Lelyveld. Granular data offer new opportunities for stress testing. *Handbook of Financial Stress Testing (forthcoming)*, 2021.

G. Wiersema, A. M. Kleinnijenhuis, T. Wetzer, and J. D. Farmer. Scenario-free analysis of financial stability with interacting contagion channels. *Journal of Banking & Finance*, 146: 106684, 2023.

A Appendix

Notation	Description
α	Assets $\alpha \in \{l, b, m, e, f\}$
β	Interbank investments $\beta \in \{l, b, m, e\}$
ϵ	Risk-adjustment factor
ρ	Risk-adjusted investments $\rho \in \{l, b, m, f\}$
σ	Securities $\sigma \in \{b, m, e, f\}$
τ	Tradable securities $\tau \in \{b, m, e\}$
μ	Market depth divisor
$\mathcal{A}$	Set of all South African institutions $\mathcal{A} = \mathcal{B} \cup \mathcal{F} \cup \mathcal{C} \cup \mathcal{G}$
$\mathcal{B}$	Set of South African Banks
$\mathcal{C}$	South African (non-financial) corporate sector
$\mathcal{D}$	Set of in-default institutions
$\mathcal{F}$	Set of South African funds
$\mathcal{G}$	South African Government sector
$\mathcal{H}$	Set of South African securities-issuing institutions $\mathcal{H} = \mathcal{B} \cup \mathcal{C} \cup \mathcal{G}$
$\mathcal{I}$	Set of South African financial institutions $\mathcal{I} = \mathcal{B} \cup \mathcal{F}$
A	Total assets
B	Buffer
C	Market Capitalization
D	Market Depth
E	Exposure
L	Loss
Δr	Price impact
N	NAV of a fund share
P	Market price
r	Liquidity factor
S	Stock S
a	External assets
b	Bonds
c	Counterparty risk contagion
d	External debt
e	Equity (shares)
f	Fund shares
h	Shareholder contagion
i	Exposed institution
j	Initially defaulted bank
k	Propagating institution
l	Loans and Deposits
m	Money market instruments (MMIs)
n	Order of exposure / round of loss
p	Overlapping portfolio contagion
q	Miscellaneous institution
s	Shares in stock S
v	Expected value
w	Weight of an edge in the asset network
$\hat{w}^{\delta}$	Weight of an edge in the direct exposure network
$\hat{w}^{\phi}$	Weight of an edge in the indirect exposure network
$\hat{w}$	Weight of an edge in the first-order exposure network
x	Random number $x \in U(0, 1)$
y	Investor bank $y \in \mathcal{B}$
z	Investee bank $z \in \mathcal{B}$

Table 4: **Notation**

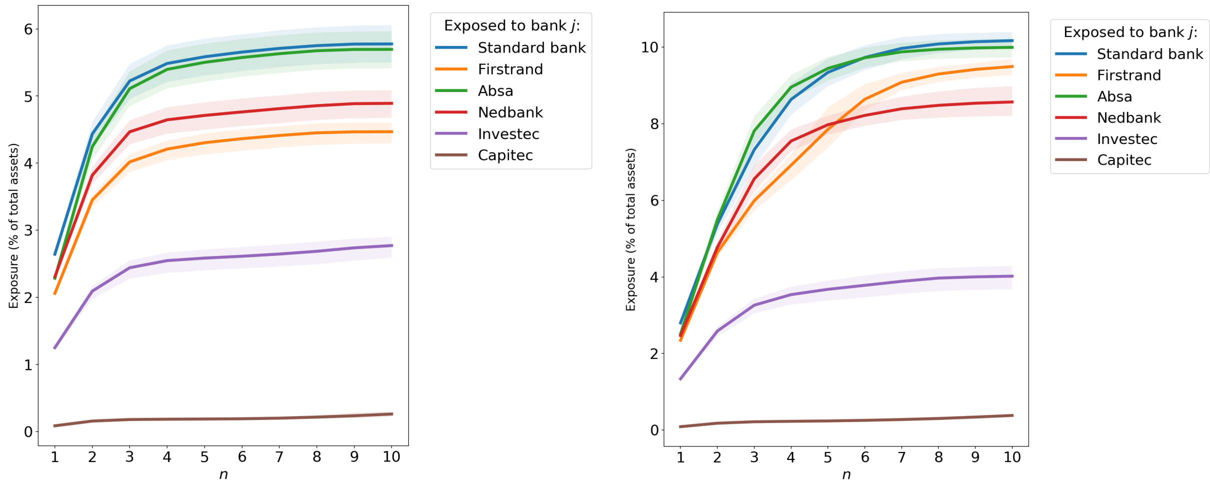

(a) Baseline exposure up to n^{th} order

(b) Stressed exposure up to n^{th} order

Figure 17: **25^{th} and 75^{th} percentiles of exposures of the South African financial system to the six largest banks.** We plot exposure (as % of the system's total assets) up to n^{th} order of the South African financial system to the default of bank j, where $n \leq 10$, j is one of the six large banks and the system's exposure is the sum of the banks' and funds' exposures. (a) shows the baseline exposures and (b) the exposures when institutions are subjected to the stress scenario, which consists of a 25% reduction in all institutions' buffers and a 50% reduction in the liquidity of all tradable assets. We find a distribution of exposures over the 1000 realized samples of the reconstructed interbank network. As we do not know the true interbank network, the true exposures may lie anywhere within this distribution. We plot the mean of the distribution as a solid line and the area between the 25^{th} and 75^{th} percentiles of the distributions as a shaded region in the same color as the mean. Plot (a) shows that baseline exposures level out around $n = 5$ and that the distribution fans out as the order of the exposure increases, which is to be expected because inaccuracies compound. However, (b) shows that stressed exposures to Standard bank, Absa and Nedbank only level out by $n \approx 8$, which suggests that the approach taken in this paper of calculating exposures up to fifth order may substantially underestimate exposures under stressed conditions.

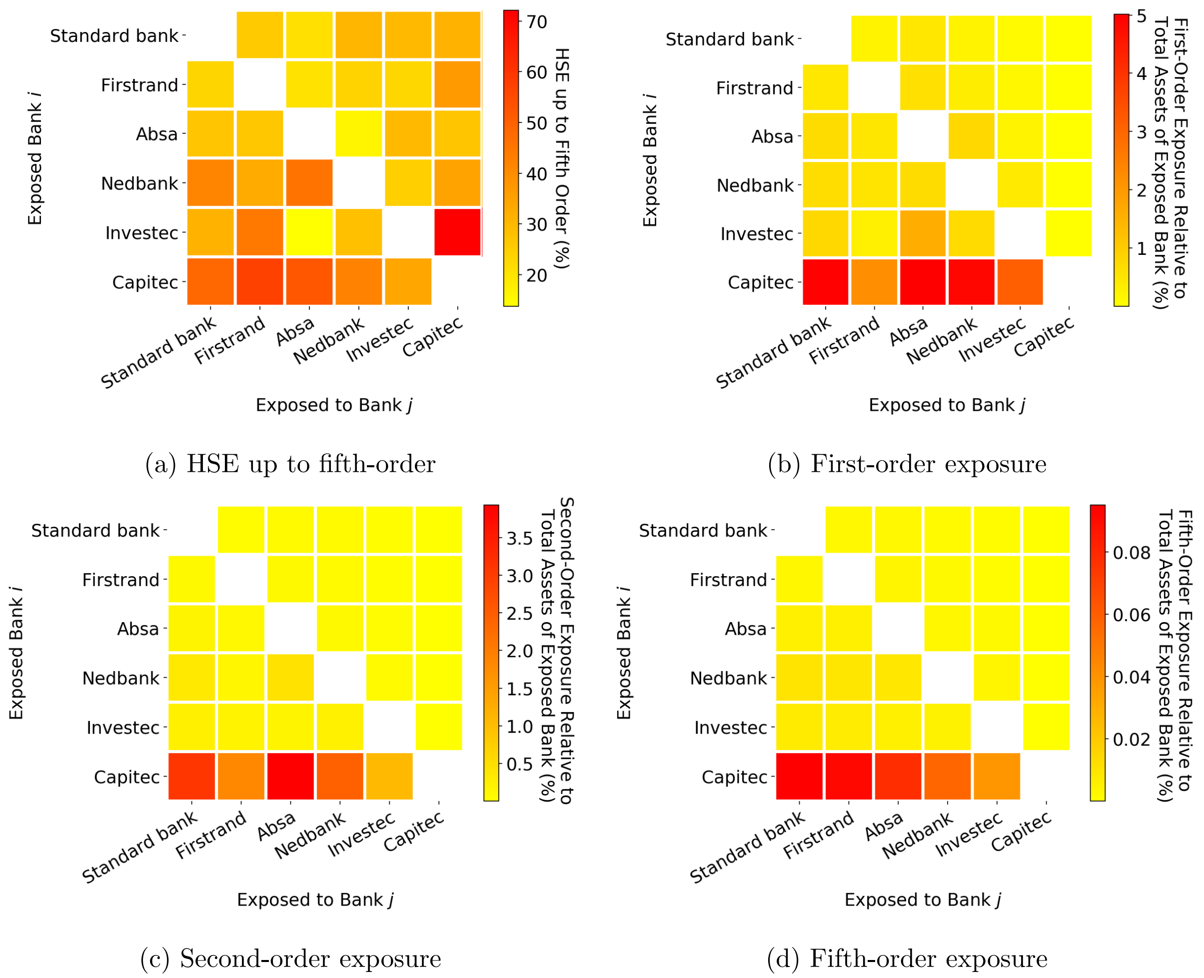

(a) HSE up to fifth-order

(b) First-order exposure

(c) Second-order exposure

(d) Fifth-order exposure

Figure 18: **Individual banks' exposures to the six largest banks**. The figure shows the exposures (as % of each bank's total assets) and HSEs between the six large banks. The banks are ordered by total asset size (descending from top to bottom) and the plots use a color gradient to indicate the HSE or exposure of one bank on the vertical axis to another bank on the horizontal axis. (Note that the banks on the horizontal axis are ordered from left to right by descending total asset size.) In general, figures (b)-(d) show that the banks have very modest exposures between them. Other than Capitec, the banks have both low first-order and low higher-order exposures. Yet, figure (a) highlights a few cases where the exposures' HSEs are substantially smaller or larger than the average. Hence, the higher-order exposures cannot be proxied by "scaled" first-order exposures in general.

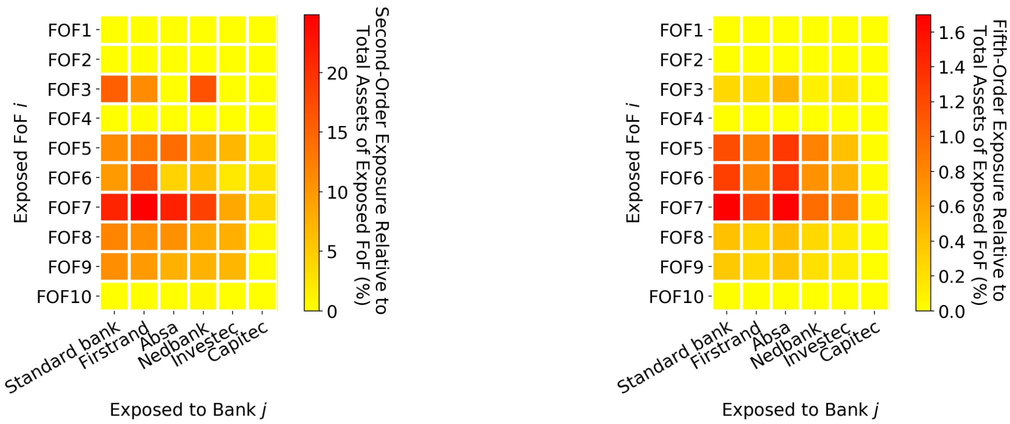

(a) Second-order exposure

(b) Fifth-order exposure

Figure 19: **Individual FoFs' exposures to the six largest banks**. The figure shows the second-order exposures in (a) and fifth-order exposures in (b) of the ten largest FoFs, as % of each FoF's total assets. The FoFs are ordered by total asset size (descending from top to bottom) and the plots use a color gradient to indicate the exposure of a FoF on the vertical axis to a bank on the horizontal axis. (Note that the banks are ordered from left to right by descending total asset size.) We do not show the FoFs' HSEs or first-order exposures, as the first-order exposures are all zero and, consequently, the HSEs are all equal to one hundred percent. The second and fifth-order exposures show strong variation across the funds. Hence, the FoFs' exposures cannot be modelled at the sectoral level but must be modelled explicitly for individual FoFs .

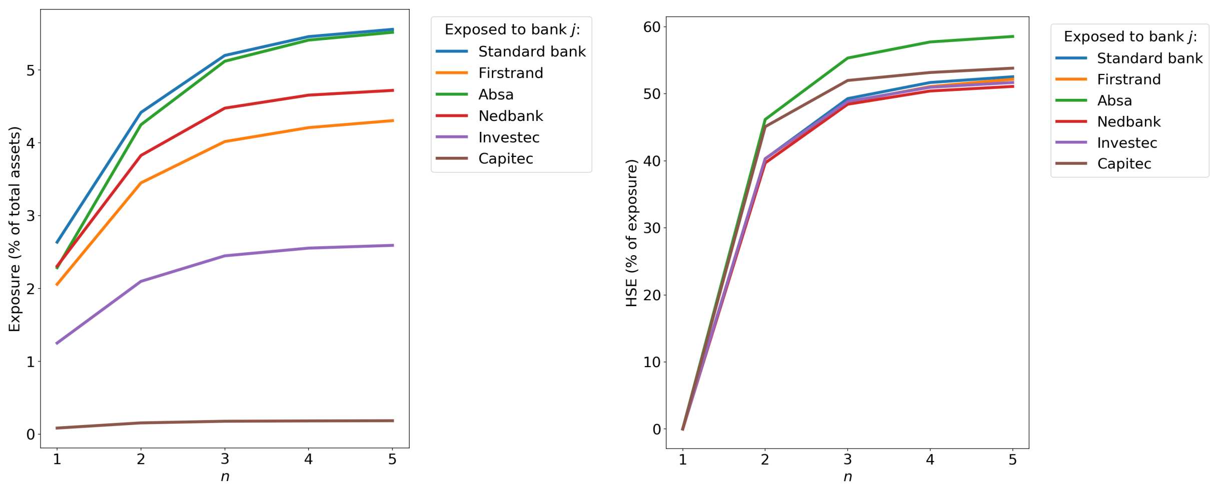

(a) Exposure up to n^{th} order

(b) HSE up to n^{th} order

Figure 20: **Exposure and higher-order share of exposure (HSE) up to n^{th} order of the South African financial system to the six largest banks when ignoring defaulted institutions' losses.** We reproduce the results in figure 7 under the assumption that institutions do not suffer losses after they default, and that the securities sold upon default do not cause any subsequent contagion after the sale. The figure shows that this assumption does not impact our results significantly.

Acknowledgements

We thank Tobias Adrian, Darrell Duffie, Luca Enriques, Doyne Farmer, David Ng, Anton Pichler, Maarten P. Scholl, and Adi Sunderam for their helpful discussions. We thank the participants of seminars at the Stanford Graduate School of Business, the Institute of New Economic Thinking at the Oxford Martin School, and the South African Reserve Bank. The authors take full responsibility for the contents of this paper and report limited use of Generative AI to improve the readability of the Non-Technical Summary. Working paper - in process of incorporating feedback.

This paper should not be reported as representing the views of the European Central Bank (ECB).

Garbrand Wiersema

Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Oxford, United Kingdom; European Central Bank, Frankfurt am Main, Germany; email: garbrandwiersema@outlook.com

Alissa M. Kleinnijenhuis

Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Oxford, United Kingdom; Cornell SC Johnson College of Business, Cornell University, Ithaca, United States; Imperial College Business School, Imperial College London, London, United Kingdom; email: alissa.kleinnijenhuis@gmail.com

Esti Kemp

International Monetary Fund, Washington D.C., United States; email: ekemp@imf.org

Thom Wetzer

Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Oxford, United Kingdom; Oxford Sustainable Law Programme, University of Oxford, Oxford, United Kingdom; Smith School of Enterprise and the Environment, University of Oxford, Oxford, United Kingdom; Faculty of Law, University of Oxford, Oxford, United Kingdom; email: thom.wetzer@law.ox.ac.uk

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

PDF ISBN 978-92-899- 7419-6 ISSN 1725-2806 doi:10.2866/5533965 QB-01-25-193-EN-N