EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Garbrand Wiersema, Esti Kemp,
J. Doyne Farmer

Liquidity spirals

No 3169

Abstract

The financial crisis of 2007-2008 highlighted the risks that liquidity spirals pose to financial stability. We introduce a novel method for studying liquidity spirals and use this method to identify spirals before stock prices plummet and funding markets lock up. We show that liquidity spirals may be underestimated or completely overlooked when interactions between different types of contagion channels or institutions are ignored. We also find that financial stability is greatly affected by how institutions choose to respond to liquidity shocks, with some strategies yielding a "robust-yet-fragile" system. To demonstrate the method, we apply it to a highly granular data set on the South African banking sector and investment fund sector. We find that the risk of a liquidity spiral emerging increases when the pool of institutions' most liquid assets is reduced, while a liquidity injection by the central bank can dampen the spiral. We further show that a liquidity spiral may be due to the banking and fund sectors' collective dynamics, but can also be driven by an individual sector under some market conditions. The approach developed here can be used to formulate interventions that specifically target the sector(s) causing the liquidity spiral.

Keywords

Financial Contagion, Systemic Risk, System-Wide Stress Test, Liquidity Risk, Non-Banks Financial Institutions (NBFIs)

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

1 Non-Technical Summary

The 2007–2008 financial crisis underscored the dangers of sudden, self-reinforcing declines in market and funding liquidity, known as liquidity spirals. This study introduces a new method to detect such spirals early—before they lead to severe disruptions in financial markets and funding availability. A key takeaway is that liquidity spirals often arise from the interactions between different transmission channels of financial stress (known as contagion channels) and from interactions between banks and non-bank financial institutions (NBFIs). Ignoring these interactions can lead to underestimating or entirely overlooking emerging liquidity risks.

This research uses a stability measure based on a so-called the shock transmission matrix (Wiersema et al., 2023) to identify when liquidity spirals start to develop. It does not rely on any particular stress scenario but instead tracks when feedback effects within the system become strong enough to trigger a spiral. When this measure crosses a critical threshold, it signals that even a small liquidity shock can trigger a rapid deterioration in market and funding conditions. A key contribution of the shock transmission matrix is that captures multiple interacting contagion channels and both banks and NBFIs.

A novel insight from this work is that the way financial institutions decide to sell assets in response to a liquidity shock plays a crucial role in determining overall system stability. For instance, institutions that sell their most liquid assets first may be able to withstand small liquidity shocks quite well. However, when a larger shock occurs, they may quickly run out of liquid assets to sell, causing liquidity spirals to intensify. This creates what is described as a "robust-yet-fragile" system: strong against small shocks but highly vulnerable to larger ones.

To demonstrate the practical relevance of their approach, we apply the model to the South African financial system, because we have access to a highly granular dataset. The dataset includes security-level bilateral exposures of the vast network of over a thousand investment funds and balance-sheet level data on the strongly concentrated banking sector. While this study focuses on South Africa, the concept and its policy relevance apply broadly to other financial systems, including the euro area. The study shows that liquidity spirals can emerge from the combined actions of these sectors, but sometimes can be driven mainly by one sector alone. This insight has important implications for policy responses. For example, central bank liquidity support may need to be carefully targeted: simply providing liquidity to one sector may not be enough to stabilize the whole system if the other sector is driving the spiral.

In summary, the paper offers a comprehensive approach to detecting and analyzing liquidity spirals by considering how shocks spread through interconnected financial sectors. It underscores that financial institutions' choice of response to shocks—particularly regarding which assets they sell first in a crisis—can strongly influence whether a small shock remains contained or snowballs into a full-blown liquidity crisis. By providing a tool for regulators and policymakers to identify emerging liquidity spirals before they lead to sharp declines in market and funding liquidity, the framework can help identify sector-specific vulnerabilities and guide targeted policy responses to maintain financial stability. Such early warnings can inform policy and regulation and support decisions on whether—and how—to intervene in financial markets.

2 Introduction

The Global Financial Crisis of 2007/2008 demonstrated the dangers of systemic risk. Today, rising geopolitical tensions are making another such crisis increasingly likely. A prime example of the endogenous dynamics that drive systemic risk is a "liquidity spiral" (Brunnermeier and Pedersen, 2009).

The term liquidity spiral refers to the progressive worsening of market and funding liquidity due to positive feedback loops in the financial system. These positive feedback loops are made up of mechanisms that propagate financial shocks; so-called contagion channels (Allen and Gale, 2000). Various contagion channels have been studied in the literature [1] and the interactions between different contagion channels have been observed to severely amplify instabilities.[2] Furthermore, multiple types of institutions across various sectors may be involved in the contagion process (see e.g. Farmer et al., 2020, Wiersema et al., 2025). This highlights the importance of capturing the various types of interacting contagion channels and institutions in models of liquidity spirals.

We identify liquidity spirals as they emerge using a *shock transmission matrix* (Wiersema et al., 2023). The matrix captures the stability of various interacting contagion channels without relying on any specific, subjective stress scenarios and includes both banks and non-bank financial institutions (NBFIs) to study spillovers between the sectors. When the dominant eigenvalue of the matrix exceeds one, a liquidity spiral emerges that depresses market and funding liquidity. Using this threshold, we can study liquidity spirals and what causes them, without having to observe falling market and funding liquidities to identify a spiral.

The model immediately bears out a number of insights. We find that liquidity spirals may be severely underestimated or even completely overlooked when interactions between different types of contagion channels are ignored. In fact, under most conditions, no single contagion channel can individually cause a spiral and any emerging spiral is the result of multiple interacting contagion channels. We also find that the type of assets that institutions choose to liquidate in response to a liquidity shock greatly affects the

[1] See e.g. Kiyotaki and Moore, 1997, Shleifer and Vishny, 1997, Allen and Gale, 2000, Eisenberg and Noe, 2001, Gorton and Metrick, 2012

[2] See e.g. Caccioli et al. (2013), Poledna et al. (2015), Kok and Montagna (2013), Wiersema et al. (2023), Cont et al. (2020), Detering et al. (2021).

risk of a liquidity spiral to emerge. In particular, we identify liquidation strategies that yield a "robust-yet-fragile" system, which is resilient to small shocks, but may become highly unstable due to a single large shock to institutions' liquidity. This phenomenon, i.e. a system being resilient against one type of shock but vulnerable to another, was also discovered by Gai and Kapadia (2010) when studying a different characteristic of financial systems; its network topology. The identification of these robust-yet-fragile tendencies across multiple characteristics of financial systems underscores the importance of stability measures that assess a system's resilience to a wide range of shocks, like the eigenvalue-based stability measure used in this paper.

To demonstrate our method, we apply it the South African banking and investment fund sectors, because we have access to a highly granular data set (on the funds in particular). The South African financial system consists of a core of five large banks and a periphery of about 30 smaller banks. In addition to a highly concentrated banking sector, the financial system includes a complex, interconnected network of more than a thousand investment funds with highly variable portfolios. We evaluate the stability of the banking sector and fund sector individually, as well as their collective stability when these sectors interact. We find spirals which are due to the banking and fund sectors' combined dynamics, but also identify conditions under which a spiral emerges that is driven by an individual sector's instability.

In an attempt to counteract an emerging spiral, the South African Reserve bank could provide additional liquidity as a lender of last resort. However, our results show that interventions that target one sector to dampen a spiral caused by the other sector have little effect. Whether the central bank successfully negates the spiral therefore depends on how the liquidity is distributed across the system.

The approach developed here can be applied to financial systems around the world, provided that sufficient data are available. It enables the study of conditions under which a liquidity spiral may emerge, thereby serving as a tool to assess systemic risk within these systems. Additionally, this method can identify the sector(s) responsible for triggering a liquidity spiral and facilitate the design of targeted interventions for those sectors.

2.1 Contributions

Our results complement the previous literature on market and funding liquidity crises. Various mechanisms that may progressively worsen market and funding liquidity have been studied in, e.g., Brunnermeier and Pedersen (2009), Gorton and Metrick (2012), Thurner et al. (2012), Hałaj (2018) and Sydow et al. (2024b). However, such studies only include a subset of the contagion channels that our model captures and therefore may underestimate the severity of the liquidity spiral.[3] Our approach of identifying liquidity

[3] Caccioli et al. (2013), Kok and Montagna (2013), Poledna et al. (2015), Wiersema et al. (2023), Detering et al. (2021).

spirals based on the contagion dynamics' dominant eigenvalue rather than based on a specific, subjectively determined stress scenario further strengthens the comprehensiveness of the analysis (Borio et al., 2014, Wiersema et al., 2023).

Our main contribution is the insight that our method provides into the impact of institutions' choices on financial stability; which actions institutions choose to take in response to liquidity shocks strongly affects the potential for a liquidity spiral to emerge. Institutions' liquidation strategies have been empirically studied in Kim (1998), van den End and Tabbae (2012), and Ma et al. (2020), but to the best of our knowledge, these strategies' impact on financial stability has not been explicitly studied previously. Furthermore, our finding that certain liquidation strategies may yield a robust-yet-fragile financial system complements the identification of robust-yet-fragile network topologies by Gai and Kapadia (2010), and demonstrates that financial systems show this property across multiple dimensions.

We also contribute to the literature on the interconnectedness of the South African financial system (see e.g. Kemp, 2017, Wiersema et al., 2025). Using a similar data set as we do here, Wiersema et al. (2025) study counterparty exposures in the South African financial system and how they are affected by the solvency of South African banks and investment funds. The analysis presented in this paper broadens the understanding of the stability of the South African financial system by focusing on the liquidity-driven contagion dynamics of the banks and funds.

Finally, our analysis contributes to the broader literature on systemic risk in financial systems[4]. While the model is specifically applied here to study liquidity spirals, it also serves as a general measure of financial stability, as discussed in Wiersema et al. (2023). We demonstrate how this analysis can be calibrated with granular data on bilateral exposures between financial institutions, yielding valuable insights into the stability of real-world financial systems and the potential impact of various interventions.

2.2 Structure

The remainder of this paper is organized as follows: Section 3 presents our method for identifying liquidity spirals and the insights it offers. In section 4, we apply the framework to the South African financial system. We discuss our data on the South African banks and investment funds, and present the results of the calibration of our method to the South African financial system. Section 5 concludes by discussing the implications and limitations of our results, and provides avenues for further research.

[4]See e.g. Bernanke (1983), Eisenberg and Noe (2001), Brunnermeier (2008), Reinhart and Rogoff (2008), Haldane and May (2011), Adrian and Brunnermeier (2011), Acharya et al. (2017), Brownlees and Engle (2017), Baron et al. (2021).

3 Identifying Liquidity Spirals

We use the framework developed in Wiersema et al. (2023) to study the conditions under which liquidity spirals emerge in the South African financial system. This framework allows us to capture many interacting contagion channels and sectors without relying on any specific stress scenario, which are often subjectively defined (Borio et al., 2014). By capturing the contagion dynamics in a linear framework, we can use the dominant eigenvalue to identify a liquidity spiral without actually having to observe a liquidity crisis developing.

3.1 The Solvency-Liquidity Nexus

The contagion dynamics captured by are model are fundamentally shaped by the Solvency-Liquidity Nexus. At the heart of a severe financial crisis lies the default of one or more institutions (Brunnermeier, 2008, Roukny et al., 2013), which can occur due to either insolvency or illiquidity. Insolvency arises when an institution's liabilities surpass the value of its assets, causing its equity to become negative (Amini et al., 2016). Conversely, illiquidity-driven default occurs when an institution—despite being solvent—fails to meet its payment obligations because it lacks the cash or liquid assets to do so (Cont and Schaanning, 2017). While solvency and liquidity are closely related, they remain conceptually distinct: a liquidity shock can force even a solvent institution to default, just as insolvency can lead to default despite ample liquidity.

In times of financial stability, solvent institutions can typically borrow to bridge temporary liquidity shortfalls. However, during periods of crisis, borrowing becomes difficult or impossible as lending markets break down. Factors such as uncertainty about asset valuations, heightened collateral demands, liquidity hoarding, and capital flight can paralyze normal market functioning (Rochet and Vives, 2004, Gorton and Metrick, 2012). Consequently, liquidity shortfalls pose a more immediate threat during crises, even if the underlying problem is one of solvency.

We can analyze the stability of the financial system in terms of its resilience to shocks, which we can classify either as liquidity shocks or valuation shocks, depending on the type of default they threaten to cause. For the purposes of this paper, we define a liquidity shock as an unexpected outflux (or cancellation of an expected influx) of liquid assets and a valuation shock as a drop in the value of an institution's assets.[5] Although we are principally interested in how liquidity shocks depress market and funding liquidity, the relevance of valuation shocks will become clear in the next section, where we show that

[5] We consider *expected* inflows and outflows of liquid assets as part of regular day-to-day liquidity management, and therefore do not classify such flows as a liquidity shock. For simplicity, we assume that shocks are non-negative. In principle, the framework could also capture negative shocks (i.e. liquidity and asset value *gains*), but this would cause the framework to lose some of the convenient properties guaranteed by the Perron Frobenius theorem.

some contagion channels may convert valuation shocks into liquidity shocks (and vice versa).

3.2 Contagion Channels

We now demonstrate how we capture contagion channels in our model in terms of the propagation of liquidity and valuation shocks and the conversion of one type of shock into the other. We discuss the five contagion channels most likely to contribute to the emergence of liquidity spirals. Note that this set of channels differs from the contagion channels included in Wiersema et al. (2023), which highlights the flexibility of the framework.

Overlapping Portfolio Contagion

Overlapping portfolio contagion can materialize when two institutions hold common securities and either institution sells securities, which drives prices down and lowers the securities' value[6]: If institution i suffers a liquidity shock it may be forced to sell securities to raise liquidity. This depresses their price, causing mark-to-market losses to all institutions that have a position in these securities. Hence, any institution j that has a position in these securities experiences a valuation shock. Therefore, *overlapping portfolio contagion converts liquidity shocks to valuation shocks.* By increasing the demand for liquidity on trading markets, overlapping portfolio contagion depresses market liquidity.

Funding Contagion

Funding contagion occurs when an institution depends on short-term funding to provide liquidity and runs the risk that the investor might withdraw its funding[7]: If institution i depends on a short-term funding from institution j, if j suddenly withdraws the funding to meet a liquidity shock it receives, then this causes a liquidity shock to i. Hence, *funding contagion propagates liquidity shocks* and reduces funding liquidity by decreasing the supply of short-term funding.

Shareholder Contagion

Shareholder contagion occurs whenever an institution suffers losses, as those losses are passed on to its shareholders: If a valuation shock causes institution i's asset value to fall, the value of its issued shares (which represent ownership of i's assets) falls accordingly, causing losses to the shareholders[8]. Hence, *shareholder contagion propagates valuation shocks.*

[6] See e.g. Adrian and Shin (2010), Caccioli et al. (2013, 2014, 2015), Duarte and Eisenbach (2021), Cont and Schaanning (2017, 2019).

[7] See e.g. Diamond and Dybvig (1983), Acharya and Skeie (2011), Caccioli et al. (2013), Brandi et al. (2018).

[8] See e.g. Wiersema et al., 2025

Redemption Contagion

When an institution, e.g. an investment fund, issues shares that are redeemable on a short-term basis (typically daily), the institution is at risk of redemption contagion; when the institution suffers a loss and the value of its issued shares falls accordingly, shareholders may decide to redeem shares as part of risk-management or performance-based capital allocation schemes [9]. Specifically, if a valuation shock decreases institution i's asset value and its shareholders decide to redeem (some of) their shares, institution i is forced to pay back the value of those shares and thus suffers a liquidity shock. Hence, *redemption contagion converts valuation shocks to liquidity shocks.*

Deleveraging Contagion

Deleveraging contagion takes place when an institution uses borrowed funds to purchase assets.[10] Borrowing creates debt and the ratio of debt to equity is called the *leverage* λ. As part of good risk-management practices, it is common for financial institutions to target a particular leverage to control risk. If the value of assets drops, the debt burden remains constant but the equity value decreases, so leverage increases. This forces a leverage-targeting institution to pay off debt to maintain its leverage target, an action that drains the institution's liquidity.[11] Specifically, if a valuation shock decreases bank i's equity and its leverage rises accordingly, the institution must raise cash to pay off its debt to return to its target leverage. Hence, the institution essentially triggers a liquidity shock to itself, so *deleveraging contagion converts valuation shocks to liquidity shocks.* Note that institutions can also be forced to deleverage due to haircuts on collateralized debt (Brunnermeier and Pedersen, 2009); when the value of the collateral falls, the institution must pay back some of the debt (assuming that it cannot post additional collateral).

3.3 The Shock Transmission Matrix

Our model characterizes the collective stability of the five contagion channels described above without relying on any specific, subjective stress scenarios. We do so by capturing these interactions in a *shock transmission matrix* (Wiersema et al., 2023). Assuming discrete dynamics, we define the shock transmission matrix A_t at a particular time t, since contagion relationships evolve over time alongside changes in institutions' balance sheets.

Let $x^l_{t,i}$ denote the liquidity shock experienced by institution i at time t. The vector of all liquidity shocks is then the N-dimensional vector $\boldsymbol{x}^l_t$, where N is the number of

[9]See e.g. Cont and Wagalath, 2013, Goldstein et al., 2017, Aikman et al., 2019, Mirza et al., 2020, Wiersema et al., 2025, Fricke and Wilke, 2023, Sydow et al., 2024b.

[10]See e.g. Fostel and Geanakoplos (2008), Brunnermeier and Pedersen (2009), Adrian and Shin (2010), Geanakoplos (2010), Adrian and Shin (2014), Aymanns et al. (2016).

[11]We do not consider slower mechanisms to raise equity-capital, such as issuing new shares or retaining earnings.

financial institutions. Similarly, $x^v_{t,i}$ denotes the valuation shock to institution i at time t, and $\boldsymbol{x}^v_t$ covers the valuation shocks to all institutions. Combining these, the full *shock vector* $\boldsymbol{x}_t$ of dimension $2N$ is defined as:

$$\boldsymbol{x}_t = \begin{bmatrix} \boldsymbol{x}^l_t \\ \boldsymbol{x}^v_t \end{bmatrix}. \tag{1}$$

The evolution of the shock vector $\boldsymbol{x}_t$ is described by the $2N \times 2N$ shock transmission matrix A_t according to

$$\boldsymbol{x}_{t+1} = A_t \boldsymbol{x}_t. \tag{2}$$

Given that $\boldsymbol{x}_t$ is the concatenation of a liquidity and a valuation shock vector, we can partition A_t into four $N \times N$ quadrants:

$$A_t = \begin{bmatrix} A^{ll}_t & A^{vl}_t \\ A^{lv}_t & A^{vv}_t \end{bmatrix}, \tag{3}$$

so the contagion dynamics are given by

$$\boldsymbol{x}_{t+1} = \begin{bmatrix} \boldsymbol{x}^l_{t+1} \\ \boldsymbol{x}^v_{t+1} \end{bmatrix} = \begin{bmatrix} A^{ll}_t \boldsymbol{x}^l_t + A^{vl}_t \boldsymbol{x}^l_t \\ A^{lv}_t \boldsymbol{x}^l_t + A^{vv}_t \boldsymbol{x}^v_t \end{bmatrix}. \tag{4}$$

Here, the *diagonal blocks* (A^{ll}_t and A^{vv}_t) capture the propagation of liquidity and valuation shocks, respectively. The *off-diagonal blocks* A^{lv}_t and A^{vl}_t describe how liquidity shocks convert into valuation shocks and vice versa. Each of the five contagion channels maps onto a specific block in A_t, reflecting the type of shock it transmits or transforms, as summarized in table 1.

Contagion Channel	Transmission Type	Block
Overlapping Portfolio Contagion	Liquidity Shock → Valuation Shock	A^{lv}_t
Funding Contagion	Liquidity Shock → Liquidity Shock	A^{ll}_t
Shareholder Contagion	Valuation Shock → Valuation Shock	A^{vv}_t
Redemption Contagion	Valuation Shock → Liquidity Shock	A^{vl}_t
Deleveraging Contagion	Valuation Shock → Liquidity Shock	A^{vl}_t

Table 1: **Contagion Channels Shock Transmission**. The table shows which type of shocks (liquidity or valuation) the contagion channels propagate or convert into the other. The last column specifies the corresponding block of the shock transmission matrix that captures the contagion channel.

The shock transmission matrix offers a rigorous way to assess the system's stability and resilience. Because its entries are non-negative, the Perron-Frobenius theorem guarantees that A_t has a non-negative real eigenvalue whose magnitude is not exceeded by any other eigenvalue. This *dominant eigenvalue* describes the dominant dynamics of the financial system[12]. If the eigenvalue exceeds one, shocks aligned with the corresponding eigenvector are amplified over time, triggering escalating contagion and resulting in progressively worsening market and funding liquidities. We refer to this self-reinforcing process of falling liquidity driving further liquidity declines as a *liquidity spiral.* In contrast, if the dominant eigenvalue remains below one, contagion can still propagate and worsen liquidity beyond the initial shock, but the system ultimately stabilizes as the shock's impact decays.

3.4 Pecking Orders

Using the shock transmission matrix to identify liquidity spirals requires quantifying the matrix' entries. These are determined by the contagion that institutions transmit, which in turn depends on the actions that institutions choose to take in response to liquidity shocks. When a liquidity shock hits, an institution must liquidate assets to meet the shock. We assume that each institution has a *pecking order* that specifies the sequence in which the institution chooses to liquidate its assets (Kok and Montagna, 2013, Hałaj, 2018, Wiersema et al., 2023). For example, once an institution has fully sold its position in a given security, it may move on to selling another, less liquid, security. The assumption of a liquidity pecking order underpins the design of regulatory measures like the Liquidity Coverage Ratio and Net Stable Funding Ratio requirements (BIS, 2013, 2014) as well as many studies [13].

An institution with a *uniform pecking order* treats all liquid assets as equally viable, responding to shocks by simultaneously liquidating pro-rate portions of each asset. Institutions with *liquidity-differentiated pecking orders*, as the name suggests, distinguish between assets of various liquidities. Institutions across multiple sectors have been observed to liquidate assets in order of decreasing liquidity when responding to shocks[14]. This pecking order typically minimizes liquidation costs as long as shocks remain small, so we refer to it as as the *optimistic pecking order*. On the other hand, institutions with the *conservative pecking order* liquidate assets in increasing order of liquidity so as to conserve their most liquid assets for the worst circumstances. Institutions may employ the conservative pecking order in anticipation of a flight to quality (i.e. liquidity) during crises, as observed by e.g. De Haan and van den End (2013) and De Santis (2014), and to preemptively divest from illiquid securities. Otherwise, the institution may be forced to liquidate those securities during the worst of the crisis, when their price is well below their fundamental value (see e.g. Bernardo and Welch, 2004). Institutions may also decide to employ different pecking orders depending on the conditions. van den End and Tabbae (2012) observe Dutch banks to use optimistic pecking orders in normal times,

[12]See e.g. Caccioli et al. (2014), Bardoscia et al. (2017), Cont and Schaanning (2019), Wiersema et al. (2023).

[13]See e.g. Allen and Gale (2000), Kok and Montagna (2013), Hałaj (2018), Wiersema et al. (2023).

[14]See e.g. Kim, 1998, van den End and Tabbae, 2012, Ma et al., 2020, Choi et al., 2020.

while employing uniform pecking orders during crises. Jiang et al. (2021) observe similar dynamics in corporate bond mutual funds.

An institution with a liquidity-differentiated pecking order only liquidates the asset at the top of its pecking order in response to liquidity shocks until that asset is exhausted. After that, the institution is forced to move on from the depleted asset to liquidating the asset that is next in line. As long as the asset at the top of the institution's pecking order remains, that asset exclusively determines the contagion that the institution transmits in response to liquidity shocks. Hence, for shocks small enough not to exhaust the assets at the tops of institutions' pecking orders, the contagion transmitted in response to liquidity shocks is exclusively determined by institutions' most liquid assets when all institutions have the optimistic pecking order, and by institutions' least liquid assets when all institutions have the conservative pecking order.

Large shocks, on the other hand, exhaust the assets at the tops of institutions' pecking orders and force the institutions to liquidate assets lower on their pecking orders. This alters the system's dynamics: once the top asset is exhausted, the next asset in line moves to the top of the pecking order. With new assets at the top of pecking orders, the contagion dynamics associated with liquidating top assets change accordingly. These changes may be captured in a new shock transmission matrix, reflecting the new contagion dynamics of the system. Therefore, the evolving dynamics of the system in response to a large propagating shock may be described by a sequence of shock transmission matrices.

Each of these matrices characterizes the "instantaneous" stability of the system at a given point in time, much like a how a first-order Taylor approximation may be used to characterize the tangent at each point of a function. The matrix approximates the immediate response to the next shock, but its reliability diminishes as the system continues to evolve in response to shocks, in particular when shocks are large. At that point, a new shock transmission matrix must be calculated to regain accuracy. The more frequently the shock transmission matrix is recalculated, the more accurately the system's dynamics may be described by the sequence of matrices.

The valuation shock-induced contagion channels included in our model do not change strongly over time; deleveraging contagion remains constant as long as institutions maintain their leverage targets, shareholder contagion is stable as long as ownership structures remain roughly unchanged and the redemption contagion channel only needs to be updated once external holders have withdrawn a significant proportion of their fund shares. In contrast, the liquidity channels must be updated must be updated whenever pecking orders shift; funding contagion acts only through the assets at the top of institutions' pecking orders that can be redeemed or withdrawn, while overlapping portfolio contagion manifests itself solely through tradable securities at the top of institutions' pecking orders. Consequently, these two contagion channels are directly determined by institutions' pecking orders and how they evolve. Pecking orders thus play a pivotal role in shaping the potential for liquidity spirals to emerge, as we will demonstrate below.

3.5 Stylized Example

The framework that we have described allows us to study the impact of institutions' pecking orders on the potential for liquidity spirals to emerge. As an example, consider a simple banking system where all banks hold (only) two types of liquid assets; deposits at the central bank and at other banks in the system, and (irredeemable) equity shares in other banks in the system. The deposits can be withdrawn at no cost, which causes funding contagion when withdrawn from other banks (but not when withdrawn from the central bank), while the shares can only be sold at a discount due to the price impact of the sale, and cause overlapping portfolio contagion when sold. Hence, deposits sit at the top of optimistic pecking orders while shares sit at the top of conservative pecking orders. Finally, assume that banks maintain their current levels of leverage.

The shock transmission matrix (3) of this system looks as follows: The upper-left quadrant A^{ll} gives the propagation of liquidity shocks, capturing the propagation of deposits withdrawals across the network of interbank deposits. Hence, when all banks have the conservative pecking order, they only sell securities to raise liquidity and this quadrant is empty. The lower-left quadrant A^{lv}, which converts liquidity to valuation shocks, gives the firesale losses caused by selling securities. Consequently, this quadrant is empty whenever all banks have the optimistic pecking order and only withdraw deposits to raise liquidity. The lower-right quadrant A^{vv} gives the propagation of valuation shocks, capturing the propagation of shareholder losses across the network of interbank equity exposures. Finally, note that there is no redemption contagion in our stylized system made up of banks, as redemption contagion only affects funds. Therefore, the upper-right quadrant A^{vv}, which converts valuation shocks to liquidity shocks, only includes the banks' deleveraging.

The banks' pecking orders determine whether funding or market liquidity falls first during crises, as the contagion transmitted in response to a liquidity shock is determined exclusively by the assets at the top of an institution's pecking order until these assets are exhausted by the liquidity demand. Hence, when all banks have the optimistic pecking order, funding liquidity falls until some banks have withdrawn all their deposits and start selling shares, causing market liquidity to be reduced too. On the other hand, when all banks have the conservative pecking order, market liquidity declines until some banks have sold all their shares and start withdrawing deposits, which depresses funding liquidity. Finally, when institutions have the uniform pecking order, funding and market liquidity fall in tandem.

Pecking orders have a strong impact on financial stability, which we can illustrate using this simple setup. For example, we explained that the lower-left quadrant of the shock transmission matrix is empty when all banks have the optimistic pecking order, because there is no overlapping portfolio contagion when banks only withdraw deposits to raise liquidity. Therefore, when all banks have the optimistic pecking order, the shock transmission matrix is upper block triangular (while this is generally not the case for other pecking orders). When a matrix is block triangular, its dominant eigenvalue is determined completely by the blocks along the diagonal. Hence, **when all banks have the optimistic pecking order, the dominant eigenvalue is determined solely by the (diagonal) funding and shareholder contagion quadrants, and is not affected by the matrix' upper-right quadrant, i.e. the deleveraging and redemption quadrant.**

Banks' leverages are typically on the order of ten in well-developed financial systems, and consequently deleveraging strongly amplifies shocks and tends to raise the dominant eigenvalue (Bardoscia et al., 2017, Wiersema et al., 2023). Liquidity spirals are therefore less likely to occur when the deleveraging does not affect the dominant eigenvalue, i.e. when all banks have the optimistic pecking order so the shock transmission matrix is block-triangular. Hence, the system is generally more resilient to (small) shocks as long as all banks have the optimistic pecking order, rather than the conservative or uniform pecking order.

Furthermore, the example of all banks having the optimistic pecking order also serves to illustrate how a large shock may destabilize the system: When the shock is large enough to exhaust (some) banks' deposits, these banks are forced to liquidate shares to meet liquidity shocks. As explained, this changes the system's contagion dynamics, so we recalculate the shock transmission matrix. Because (some) banks will now sell securities and cause overlapping portfolio contagion, the new shock transmission matrix is no longer upper block-triangular. This causes deleveraging to have an impact on the dominant eigenvalue, which generally worsens the stability of the financial system. Under certain conditions therefore, a large shock that exhausts the assets at the top of institutions' pecking orders may cause a liquidity spiral to emerge in a previously stable financial system.

3.6 General Observations

The stylized example illustrated how a large shock may destabilize the system when banks have the optimistic pecking order. A financial system which is very resilient to small shocks but becomes unstable when hit by a large shock is referred to as "robust-yet-fragile". While we find that optimistic pecking orders may generate a robust-yet-fragile system, Gai and Kapadia (2010) find a similar phenomenon for certain topologies of financial networks. The identification of this robust-yet-fragile tendency across multiple characteristics of financial systems highlights the dangers of optimizing financial stability with respect to the small shocks that are incurred on a frequent basis; such a system may turn out to be highly fragile when a large shock eventually hits. It also underscores the risk of assessing the robustness of a financial system only to specific stress scenarios, as the system may well be quite resilient to one scenario, yet highly vulnerable to another.

Finally, we present two observations about the funding contagion channel individually causing a liquidity channel. These observations follow directly from the shock transmission matrix and the properties of block matrices, as shown in appendix A.2:

1. Out of all contagion channels considered in this analysis, the funding contagion channel is the only channel that may cause a liquidity spiral to emerge in the absence of any other contagion channels (i.e. when all other quadrants of the shock transmission matrix are empty).

2. The funding contagion channel can only individually cause a liquidity spiral when banks hoard liquidity in response to shocks, i.e. when they recover liquidity in excess of the shock incurred in order to build reserves against potential future shocks.

From observations 1 and 2 follows that, under the assumption of no liquidity hoarding, all liquidity spirals result from a combination of interacting contagion channels. Therefore, analyzing contagion channels in isolation and ignoring their interactions, as is often done, may fail to identify liquidity spirals in even the most unstable of systems.

In this section, we have developed a method for studying liquidity spirals that takes the various contagion channels and their interactions into account, and that does so without relying on any specific stress scenario. The conclusions derived from our stylized example have underscored the importance of stability measures that have these two features to comprehensively analyze the financial system. In the next section, we apply the framework to the South African financial system.

4 Measuring the Potential for Liquidity Spirals in the South African Financial System

We demonstrate our framework for identifying liquidity spirals by applying it to the South African banking and investment funds sectors. We first describe the relevant features of the South African financial system and of the institutions and assets that we cover, after which we discuss where these assets sit on institutions' pecking orders. We then define the contagion channels that follow from these pecking orders and, finally, we present the results of the application of our contagion model 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 (Kemp, 2017). 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). Due to the relative lack of well-developed peers in the region, South African institutions are very reliant on the domestic financial market, making it highly interconnected (Kemp, 2017).

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 (the latter are henceforth referred to as “investment 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 investment funds - has outpaced that of banks (Kemp, 2017). Non-bank financial intermediaries are an important source of funding for banks; banks’ funding provided directly by non-bank financial intermediaries other than pension funds and insurers amounts to 15% of bank assets (FSB, 2018).

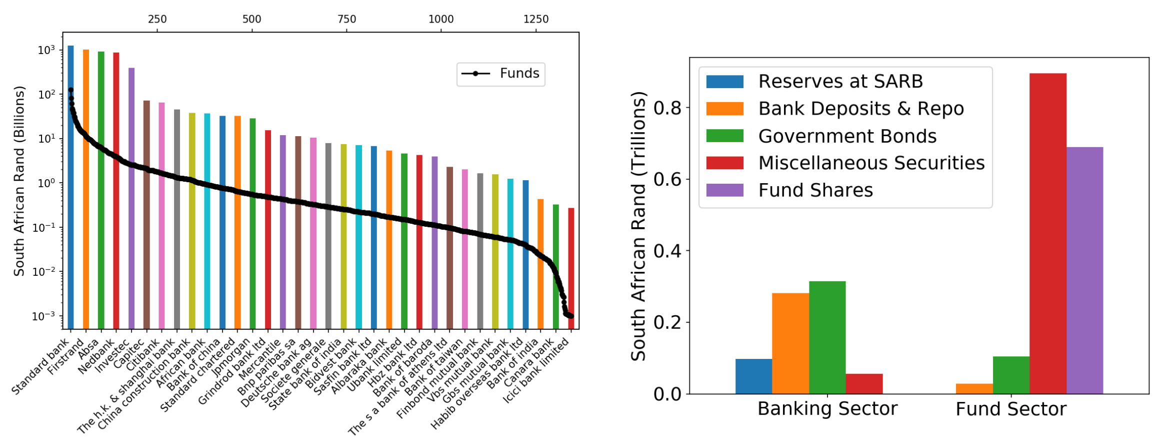

(a) Banks and funds ranked by total assets

(b) Liquid asset aggregates of both sectors

Figure 1: **Total assets and liquid assets distributions.** Figure (a) ranks the banks and funds in decreasing order of total asset size from left to right. The asset totals are listen on the y-axis (log-scale) in billions of South African Rand and the banks’ names are listed on the x-axis. The funds’ names are too numerous to list but their numbers are indicated at the top of the graph. The banking sector consist of a core of five large banks and a periphery of 29 smaller banks. The investment fund sector includes over 1300 funds and also shows a strong concentration in terms of asset size. Figure (b) shows the liquid asset holdings of the banking sector and fund sector aggregated into the categories set out in section 4.2.1. Banks do not hold fund shares, while funds do not hold reserves at the South African Reserve Bank (SARB). Miscellaneous securities include bonds issued by institutions other than the government, money market instruments, and equity shares. Note that the banks’ positions in bonds issued by the South African government greatly exceeds the funds’ position, while the funds’ hold far more miscellaneous tradable securities than the banks.

4.1 Institutions

In this study, we focus on banks and investment funds domiciled in South Africa. Insurance companies and pension funds and are not included due to data limitations, but we do not expect this to affect our results substantially as these institutions have limited contribution to the channels included in our model[15]. Furthermore, our data include tradable securities issued by the government, as well as by non-financial corporates, henceforth referred to as the corporate sector, but modeling these sectors is beyond the scope of our analysis[16].

4.1.1 Banks

The South African banking sector comprises 34 registered banks, local branches of foreign banks and mutual banks as of Q4 2016. 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 1a. Overall, the banking sector is largely funded by deposits, but banks also issue debt instruments, such as bonds and money market instruments, and equity shares. The banks' leverages (debt-to-equity ratio) vary, with a median of 7.4. The largest banks' leverages are between 11 and 13, which is not uncommon in well-developed financial systems, while the smaller South African banks typically have lower leverages.

4.1.2 Investment Funds

Investment funds pool investors' savings 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, investment funds provide investors with an opportunity to earn higher returns

[15]Leverage, i.e. institutions' debt-to-equity ratio, is the predominant source of shock amplification and main driver of liquidity spirals in our model. Because insurance companies and pension funds typically do not have high debt-to-equity ratios, they do not contribute substantially to the amplification of shocks. Insurance companies and pension funds do have substantial obligations to policyholders, respectively, participants in the fund, however, but mechanisms like profit sharing life-insurance contracts and defined contribution pensions transfer a proportion of losses to policyholders or participants. Therefore, these obligations function differently from typical debt and are beyond the scope of this paper. Nevertheless, there is some potential for pension funds and insurance companies to participate in liquidity withdrawals and firesales, which is not captured in our model. Other contagion mechanisms that may affect pension funds and insurance companies, such as surrender shocks (see e.g. Sydow et al., 2024a) are not modeled.

[16]Under the assumption that the government does not default on its debt, which is unlikely even for a developing economy, the government is not expected to cause financial contagion and can therefore be safely left out of our model. On the other hand, excluding non-financial corporates from our model may ignore instabilities driven by the feedback loop between the financial and real economy. Such interactions (e.g. increased default rates of non-financial corporates due to decreased credit provision by the financial sector) typically materialize over the span of years, however, so they are not particularly relevant when considering short-term phenomena like liquidity spirals.

than those offered by deposits, in return for taking on greater risk. There are over 1300 open-ended investment funds registered in South Africa with assets under management of about 2 trillion South African Rand in 2016. The investment sector is highly concentrated, as show in Figure 1a.

Investors invest in funds by buying fund shares, which represent ownership of a portion of the underlying portfolio. These shares are typically redeemable on a daily basis. In extreme circumstances, funds are susceptible to "runs" – i.e. large-scale redemption requests, when investors observe or anticipate a substantial drop in their fund shares' value. When a run is initiated, the investment fund may run out of liquid assets and become unable to meet redemptions. As a result, the investment fund may have to resort to fire-selling assets (Cont and Wagalath, 2013, Wiersema et al., 2025).

The value of a fund share is given by its Net Assset Value (NAV), which is equal to the the investment fund's total asset value, divided by the investment fund's total number of outstanding shares. Fund shares can be either Constant NAV-valued (CNAV) or Variable NAV-valued (VNAV). When a investment 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. 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 fund share portfolio is identical. Therefore, we assume for simplicity that all fund shares in our model are VNAV-valued.

4.2 Assets

The data used are sourced from two publicly available data sets as of Q4 2016. 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, or *SARB* (SARB, 2016). Balance sheet entries are aggregated by asset type and counterparty type (e.g. "deposits at domestic banks"). The bank data distinguish between the various asset types discussed in the next section, and between all counterparty types considered in our model (i.e. banks, funds, the corporate sector and the government). As our contagion model requires a network of bilateral exposures between banks, we generate random reconstructions of this network from banks' aggregate holdings per asset type. We explain the reconstruction method, and why it does not significantly affect our results, in appendix A.4. The bank data also cover asset and counterparty types not included in our model, such as household mortgages.

Data on investment funds' assets were sourced from Morningstar Inc and are highly granular. These data report the funds' investments per instrument type in individual counterparties (e.g. "bonds issued by Standard Bank"). The great majority of the funds' assets are covered by our model, but a small fraction (less than 5%) is excluded due to data limitations, e.g. because the counterparty or asset type is unknown. Note that neither the bank nor fund data include short positions in tradable securities, so only long positions are considered.

4.2.1 Balance Sheet Composition

The liquid assets included in our analysis cover deposits, bonds, equity shares, and fund shares, repo loans, and money market instruments, or *MMIs*. We refer to MMIs, equity shares, and bonds issued by institutions other than the government as *miscellaneous securities*. The liquid assets included in the pecking orders we consider cover a little over 15% of banks' total assets in aggregate, with mortgages and corporate loans making up the the majority of banks' illiquid assets. Furthermore, 87% of funds' assets consist of exposures to other institutions in our model, while the remaining 13% of fund assets constitute investments in foreign or non-financial corporations.

Table 2 shows where the included assets appear on the stylized balance sheets of banks and investment funds. Portfolio composition varies strongly across investment funds, and some funds in particular have a strategy of heavily investing in a specific type of asset. For example, Money Market Funds predominantly invest in deposits and MMIs issued by banks (SARB, 2017b), while so-called "Fund-of-Funds" mainly invest in shares issued by other funds (Kemp, 2017). This gives rise to a complex, interconnected asset network.

We only distinguish between the types of tradable securities that are individually listed on banks' balance sheet data: government bonds, bank bonds, corporate bonds, bank MMIs, corporate MMIs, bank equities and corporate equities. We do not distinguish between different types (e.g. bonds of different maturities) or issuers (e.g. equities issued by different banks) within these categories due to data limitations. For example, all domestic bank bonds are assumed to have the same market depth, and selling a bank bond is assumed to cause the same price impact across all bonds issued by any domestic bank. Hence, we lose some granularity in the distribution of firesale losses, but given that we still include 7 different types of securities and that in reality similar securities would be correlated anyway, we expect the impact on our results to be small.

Note that the banks also have a small position in gold, which could provide liquidity to absorb shocks (Bayram and Othman, 2022). However, as the position only makes up about 1% of banks' total assets in aggregate and due to data limitations, gold is not included in our main results. In appendix A.1, we explore the inclusion of gold as a liquid asset in the miscellaneous securities category and find that our results are not strongly affected.

Assets		**Liabilities**
Repurchase agreements		Repurchase agreements
SARB and Bank Deposits		Banks' and Funds' Deposits
Tradable securities	MMIs	MMIs
	Bonds	Bonds
	Equity Shares	Other liabilities
Other assets		**Equity**

(a) Stylised balance sheet of a bank.

Assets		**Liabilities**
Bank Deposits		None
Tradable securities	MMIs	
	Bonds	
	Equity shares	
Fund shares		
Other assets		**Fund Shares**

(b) Stylised balance sheet of an investment fund.

Table 2: **Stylised balance sheets of South-African banks and investment funds**. (a) shows the stylized balance sheet of a bank and (b) of an investment fund. Note that specific subsets of investment funds may heavily invest in one type of asset while not investing in other types.

Central Bank Reserves

Banks hold reserves at the SARB, whereas investment funds do not. We assume that central bank reserves are perfectly liquid and do not cause contagion when withdrawn. SARB reserves make up 13% of banks' liquid assets in aggregate.

Commercial Bank Deposits & Repo Loans

South African banks take deposits and issue repo, while investment funds do not (as they do not take on debt). Furthermore, both banks and funds make deposits at (other) South African banks, but only banks buy repo. The banks and funds also make deposits at foreign banks, but these only make up a small part of the banks' and funds' portfolios. We assume that both deposits and repo loans can be redeemed on a daily basis. Therefore, the dynamics in our model of repo and deposits at (commercial) banks are identical, so we group them together for simplicity. While deposits at South African banks represent only 1% of funds' liquid assets, they make up 37% of banks' liquid assets. Due to data limitations, we do not model collateral. Consequently, our model does not capture using unencumbered assets to raise liquidity from new repo funding or a central bank facility. While the repo market serves to redistribute liquidity from those with a surplus to those with a deficit, it cannot solve a general lack of liquidity in the system, so the exclusion of new repo funding from our model is not expected to strongly affect our results. Liquidity received from the central bank, however, has the potential to stabilize the system, as explored in section 4.6.1.

Bonds

Domestic bonds are issued by banks, the corporate sector and the South African government. Additionally, South African banks and investment funds own some bonds issued by foreign parties, but these position are minor. The investment fund data distinguish between bonds issued by different banks, whereas the bank data do not. Contrary to repo loans and deposits, bonds are tradable. South African government bonds make up 42% of banks' liquid assets in aggregate, and bonds issued by other South African institutions 3%. Furthermore, 5% of funds' liquid assets are made up of South African government bonds and 3% are made up of bonds issued by other South African institutions.

Money market instruments

MMIs 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 identically in our model. This means that we lose some granularity in the firesale channel but is not expected to affect our results substantially, as discussed previously. MMIs in our data are exclusively issued by domestic banks and corporates and are bought by both banks and funds. Like bonds, MMIs are tradable. MMIs make up 19% of funds's liquid assets in aggregate and 3% of banks' liquid assets.

Equity shares

Funds invest in listed equity shares issued by the South African banks and the corporate sector, while banks also invest in listed shares issued by the corporate sector but hold unlisted shares in (other) South African banks. Furthermore, South African banks and investment funds own some listed equity shares issued by foreign parties but the great majority of shares held by the banks and investment funds are domestically issued. Listed equity shares are tradable whereas unlisted shares are not, and neither are redeemable. 29% of funds' liquid assets are made up of equity shares in South African institutions in aggregate, while domestic equity shares make up less than 1% of banks' liquid assets.

Fund Shares

Fund shares are issued by investment funds and are assumed to be redeemable on a daily basis (as is almost always the case in reality). Therefore, fund shares are not traded. South African investment funds buy other funds' shares while banks do not. 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 investment fund may suffer. Fund Shares make up 36% of funds' liquid assets in aggregate.

4.2.2 Initialization Values

We do not have data on the market prices or NAVs of the securities that institutions hold, nor the number of securities they hold, but only on the value of an institution's position in a security (i.e. the market value of a position in a tradable security, or the NAV times the number of shares of a position in fund shares). We assume that the initial NAV of each fund share and the initial market price of each tradable security to one South African Rand, and that the initial market price of a listed equity share is equal to its book value (i.e. equity shares' initial market values are assumed equal to the issuer's accounting value of that share). These assumptions are only for simplicity do not affect our results.

4.3 Pecking Orders

The contagion that an institution transmits in response to a liquidity shock and its contribution to the potential emergence of a liquidity spiral is determined by the institution's pecking order. Now that we have introduced the various liquid assets included in our model, we can discuss where these appear on institutions pecking order. To distinguish between various liquidity-differentiated pecking orders, we group assets in decreasing order of liquidity as follows:

1. Central bank reserves
2. Deposits at commercial banks, repo loans, and fund shares
3. Government bonds
4. Miscellaneous tradable securities (bonds, MMIs, and listed equity shares issued by banks or non-financial corporates)

Hence, institutions with the optimistic pecking order liquidate assets in order from group 1. to 4., and the conservative pecking order is the reverse. Note that any institution that does not own any assets of the types in the group at the top of the pecking order liquidates assets from the group that is next in line. Table 3 summarizes the various pecking orders that we consider.

The limited empirical research into the pecking orders does not guarantee that the optimistic and conservative pecking orders are the only liquidity-differentiated pecking orders that institutions employ under various circumstances. We therefore consider two more pecking orders for a comprehensive picture of the system's stability. The main motivation for including these two additional pecking orders, is so that each of the four groups of liquid assets we distinguish sits at the top of an included pecking order.

The *short-term funding pecking order* is the optimistic pecking order but with group 1. moved to the bottom and consequently has group 2. at the top. Hence, institutions with the short-term funding pecking order use withdrawing deposits, not rolling over repo loans and redeeming fund shares as their first line of defense against liquidity shocks. These are indeed a common ways of liquidity for many institutions, and banks may resort to these methods whenever they wish to conserve their central bank reserves (for practical or regulatory reasons) and use their next most liquid assets to respond to shocks instead. Note that funds do not hold reserves at the central bank, so the optimistic and short-term pecking orders are identical for funds.

The *government bonds pecking order* has group 3 at the top and groups 1. and 2. moved to the bottom. Hence, institutions with the government bonds pecking order use selling these bonds as their preferred method of responding to liquidity shocks, which is also observed in practice. The government bonds pecking order is similar to the conservative pecking order, as institutions conserve their directly accessible sources of liquidity (deposits, repo, fund shares, and reserves in the case of banks) and firesell securities to raise liquidity instead. The only difference is that instead of selling their miscellaneous securities, institutions with the government bonds pecking order preferentially sell these bonds instead, which are more liquid.

	Optimistic pecking order	Short-term funding pecking order	Government bonds pecking order	Conservative pecking order	Uniform pecking order
Top	Central bank deposits	Deposits at commercial banks, repo, fund shares	Government bonds	Miscellaneous tradable securities	All liquid assets
	Deposits at commercial banks, repo, fund shares	Government bonds	Miscellaneous tradable securities	Government bonds	
	Government bonds	Miscellaneous tradable securities	Central bank deposits	Deposits at commercial banks, repo, fund shares	
Bot- tom	Miscellaneous tradable securities	Central bank deposits	Deposits at commercial banks, repo, fund shares	Central bank deposits	

Table 3: **Pecking orders**

Finally, in addition to the pecking orders discussed above, we also consider a *sector-specific pecking order*, where banks employ the conservative pecking order, while funds have the government bonds pecking order. Under Basel regulation, banks are subject to a minimum Liquidity Coverage Ratio (LCR) requirement and a maximum Risk-Weighted Assets (RWA) limit. Although we cannot model LCRs and RWAs explicitly due to data limitations, banks are likely to employ the conservative pecking order whenever they face an LCR or RWA constraint. This is because the conservative pecking order liquidates assets in increasing order of contribution to the LCRs and decreasing order of contribution to RWAs, making it the most effective pecking order to return to an LCR or RWA target[17]. Funds, on the other hand, do not face RWA or LCR regulation and are therefore more likely to employ the government bonds pecking order. In fact, funds will liquidate government bonds whenever they choose to conserve their sources of directly accessible liquidity and sell their most firesell their most liquid securities first.

An institution may have multiple assets at the top of its pecking order simultaneously. For example, institutions with the uniform pecking order have all of their liquid assets at the tops of their pecking orders simultaneously. Furthermore, an institution with e.g. the short-term funding pecking order may have various deposits, repos and fund shares issued by multiple institutions at the top of its pecking order. We assume that institutions liquidate a vertical slice across all assets at the top of their pecking orders, i.e. the institution recovers an amount of liquidity from each asset proportional to that asset's total value. Furthermore, we assume for simplicity that the vertical slice used to respond to the liquidity shock $\boldsymbol{x}_t^l$ is based on the asset values at the start of round t (i.e. before the liquidity and valuation shocks $\boldsymbol{x}_t^l$ and $\boldsymbol{x}_t^v$ are taken into account, which may reduce the value of some securities).

Because institutions liquidate a vertical slice across all assets at the top of their pecking orders, an institution with the uniform pecking order reduces each of its liquid assets by the same proportion in response to a liquidity shock. Therefore, the shock only affects the magnitude but not the composition of the institution's pool of liquid assets, leaving the institution's pecking order and, consequently, its response to liquidity shocks unchanged (which is not the case for other pecking orders when shocks exceed the top layer).

4.4 Contagion Equations

We now derive representative formulas for each of the contagion channels that act on the asset types that may sit at the top of pecking orders, under the assumption that liquidity shocks do not exceed the top layer. (We consider shocks that exceed the top layer in section 4.8.) Note that these forms are chosen for simplicity and that we do not consider liquidity hoarding, so institutions do not attempt to recover liquidity in excess of the shock incurred. More elaborate contagion models may be considered when studying individual contagion channels, but these forms suffice for our present purposes of capturing the collective stability of these interacting contagion channels.

[17] Out of the four liquid asset classes considered, miscellaneous securities generally have the lowest contribution to LCRs and highest contribution to RWAs. Ceterus paribus, liquidating these assets is therefore the most effective strategy to increasing LCRs and lowering RWAs. (Note that the cash recovered from liquidating these securities has maximal LCR contribution and zero RWA contribution.)

4.4.1 Funding Contagion

Suppose institution i has deposits at institution j and/or has bought repo or fund shares issued j with a total value of $d_{ij,t}$. Furthermore, suppose that these form the top of i's pecking order and let $T_{t,i}$ denote the total value of these assets. When these deposits, repo and/or fund shares are part of the top layer, on receiving a liquidity shock $x^l_{i,t}$, institution i withdraws a total value of $x^l_{i,t}d_{ij,t}/T_{t,i}$ of these deposits, repo and/or fund shares (i.e. i liquidates a vertical slice across the assets in its top pecking order layer). This transmits a liquidity shock $A^{ll}_{ji,t}x^l_i$ to institution j, where

$$A^{ll}_{ji,t} = \frac{d_{ij,t}}{T_{i,t}}. \tag{5}$$

In our model, withdrawing liquidity from the SARB or foreign institutions does not cause contagion, while it reduces the amount of liquidity required to be recovered from other sources (which typically do cause contagion). Therefore, withdrawing liquidity from the SARB or foreign institutions has a stabilizing effect on the system by reducing contagion from raising liquidity through other means.

4.4.2 Overlapping portfolio contagion

Suppose institution i holds $n_{\sigma i,t}$ shares of security σ with market price $p_{\sigma,t}$, which are in i's top pecking order layer. When i experiences a liquidity shock $x^l_{i,t}$, it sells shares in security σ worth $x^l_{i,t}n_{\sigma i,t}p_{\sigma,t}/T_{i,t}$ (i.e. a vertical slice) to raise liquidity. The sale depresses the price of security σ by a price impact of $\Delta p_{\sigma,t} = p_{\sigma,t} - p_{\sigma,t+1}$, which causes losses to all institutions that have a position in the security. Movements in market prices are a function of supply as well as demand. While the sale represents an increased supply, we do not explicitly model the demand side. Rather, we use the market depth $D_{\sigma,t}$ as a proxy for demand-side conditions to model the price impact of the sale.

We assume that the market price $p_{\sigma,t}$ falls by the same proportion as the fraction of the market depth $D_{\sigma,t}$ sold, such that price impacts are additive when multiple institutions i decide to sell shares in security σ at time t[18]:

$$\frac{\Delta p_{\sigma,t}}{p_{\sigma,t}} = \sum_{i\in\mathcal{S}_t} \frac{x^l_{i,t}n_{\sigma i,t}p_{\sigma,t}}{T_{i,t}D_{\sigma,t}}, \tag{6}$$

[18]We show in appendix A.5 that this assumption implies that the price impact is a concave function of the number of securities sold. Empirical evidence suggests that the price impact is indeed a concave function although the particular shape may depend on the context (Gatheral, 2010) and is challenging to quantify (Fukker et al. (2022)). We do not aim to perfectly replicate any of the empirically observed functional forms, as the current approximation suffices for our present purposes. For more accurate modeling of the overlapping portfolio contagion channel see e.g. Bouchaud and Cont (1998), Bouchaud et al. (2009).

where $\mathcal{S}_t$ gives the set of institutions that sell security σ at time t. Equation (6) implies that the market depth $D_{\sigma,t}$ represents the maximum amount of liquidity that can be extracted from the market for security σ before the price falls to zero. As we cannot measure the market depth directly due to data limitations, we assume that the market depth D_σ is proportional to the market capitalization C_σ for any security σ[19]:

$$D_{\sigma,t} \stackrel{\text{def}}{=} \frac{C_{\sigma,t}}{\delta}, \tag{7}$$

where the estimation of the market capitalization C_σ is discussed in appendix A.7, and the ratio of the market cap to market depth δ, or simply *cap-to-depth ratio*, is a non-dimensional constant for which we explore various values in the results section.

As prices cannot fall below zero even in the extreme scenario where all holders of a security liquidate their positions simultaneously (i.e. the entire market capitalization is sold), the market capitalization must thus give a conservative lower bound to market depth under the assumption of linearity. However, empirical evidence suggests that the price impact is most likely a concave function of trading volume, so the marginal price impact decreases as volume increases. Hence, the price impact of small volumes sold could theoretically exceed the linear approximation with market depth equal to the market capitalization, but such a harsh price impact would probably only materialize during highly adverse market conditions (if at all), when demand-side conditions are severely compromised[20].

Assuming that the number of outstanding securities is constant, i.e. $C_{\sigma,t}/p_{\sigma,t} = C_\sigma/p_\sigma$, where C_σ and p_σ denote the initial market capitalization and price, respectively, we can rewrite equation (6) as

$$\Delta p_{\sigma,t} = \sum_{i\in\mathcal{S}_t} \frac{n_{\sigma i,t} p_{\sigma,t}}{T_{i,t}} \frac{\delta p_\sigma}{C_\sigma} x^l_{i,t}. \tag{8}$$

We assume for simplicity that all shares sold at time t are sold against the new price $p_{\sigma,t+1}$, such that any institution, including institutions $i \in \mathcal{S}_t$, suffers a valuation shock of $\Delta p_{\sigma,t} n_{\sigma j,t}$. This implies that

$$A^{lv}_{ji} = \sum_{\sigma\in\mathcal{T}_{i,t}} \frac{n_{\sigma i,t} p_{\sigma,t}}{T_{i,t}} \frac{\delta p_\sigma}{C_\sigma} n_{\sigma j,t}, \tag{9}$$

where $\mathcal{T}_{i,t}$ denotes the set of all tradable securities σ at the top of i's pecking order at time t. Note that the first term in equation (9) corresponds to i's vertical slice across the assets at the top of its pecking order, the second term to the price impact, and the third term to the j's securities affected by the price impact.

[19] See e.g. Potters and Bouchaud, 2003, Lillo et al., 2003, Cont and Schaanning, 2019, Fukker et al., 2022 for a discussion of the challenges involved in estimating the market depth.

[20] See e.g. Potters and Bouchaud, 2003, Lillo et al., 2003, Cont and Schaanning, 2019, Fukker et al., 2022, Sydow et al., 2024b for estimates of market depths in other financial markets.

Note that our data include tradable securities issued by foreign institutions. Because these securities are predominantly held by foreign institutions and we do not have data on the market depth of these securities, we do not model the overlapping portfolio contagion caused by the sale of these securities. This is not expected to affect our results significantly, as these securities are mainly held by foreign investors, so any price impact from selling these securities would be predominantly suffered by these foreign investors, too. As selling foreign securities reduces the amount of liquidity required to be raised through other means, it has a stabilizing effect on the system.

4.4.3 Deleveraging Contagion

Finally, we discuss the contagion induced by valuation shocks. We do not model defaults, as this is not the focus of our analysis and would require a different approach than the contagion model employed here; we are interested in the emergence of a liquidity spiral, which typically takes place before defaults occur and is driven by other mechanisms than credit losses. Only once the liquidity spiral has already materialized are institutions expected to start defaulting due to the spiral. Because we do not model defaults, we focus on valuation shocks that do not exceed banks' equity.

Suppose bank i maintains a leverage target λ_i (i.e. the ratio of debt to equity). If i receives a valuation shock $x^v_{i,t}$ and its equity is reduced, i must pay off debt to return to its target. The amount by which it must reduce debt is $\lambda_i x^v_{i,t}$, so i experiences a liquidity shock $A^{vl}_{ii,t} x^v_{i,t}$ at time $t+1$, where

$$A^{vl}_{ii,t} = \lambda_i, \tag{10}$$

when i is a bank. The leverage target λ_i is given by the data and is assumed to be kept constant by bank i.

4.4.4 Shareholder Contagion

We first discuss shareholder contagion for fund shares, whose NAV is publically updated daily by the issuing fund, making the shareholder contagion channel straightforward to calculate. Suppose institution i at time t holds a number $s_{ij,t}$ shares in fund j of the total number $S_{j,t}$ of shares issued by j. A valuation shock suffered by j is distributed proportionally across its shareholders through a markdown of the shares' NAV, so

$$A^{vv}_{ji,t} = \frac{s_{ij,t}}{S_{j,t}}, \tag{11}$$

when j is a fund.

Let us now consider shareholder contagion for unlisted equity shares issued by banks. Unlisted equity shares issued by banks are marked-to-book, i.e. they are valued based on the accounting equity of the issuing bank. Therefore, unlisted equity shares propagate losses to shareholders analogously to fund shares[21]: When bank j incurs a valuation shock, the accounting equity of an unlisted share is reduced by $x^v_{j,t}/S_{j,t}$ where $S_{j,t}$ is the total number of (listed and unlisted) shares issued by bank j. Hence, when i holds $s_{ij,t}$ unlisted shares in bank j, the shareholder contagion i suffers on these shares equals $s_{ij,t}x^v_{j,t}/S_{j,t}$ and so equation (11) also holds for unlisted equity shares issued by banks.

Finally, we consider shareholder contagion for listed equity shares. Since the listed equity shares issued by banks are traded, modern accounting practices require them to be marked-to-market rather than marked-to-book. Nevertheless, we will assume for simplicity that if bank j incurs a valuation shock, the value of its issued listed shares falls by the same amount as its unlisted shares (i.e. the shares' market value and book value fall by the same amount) such that equation (11) also holds for listed shares. Although this will not be exactly true in reality, it is a reasonable approximation, as market prices should reflect the issuer's performance and losses suffered by the issuer becoming public information is indeed a common cause of falling stock prices. Note that both overlapping portfolio and shareholder contagion affect the market price of listed equity shares issued by South African banks. We assume for simplicity that the overlapping portfolio and shareholder contagion channels' impact on the market price is additive and explain in appendix A.6 why this is a reasonable approximation.

4.4.5 Redemption Contagion

Suppose that fund i suffers a valuation shock $x^v_{i,t}$ at time t, which depresses the NAV of shares issued by i and may prompt i's shareholders to withdraw liquidity from the fund by redeeming shares. Funds' shares are held by other investment funds in our data and "external holders", i.e. any party other than the banks and funds that we model. We assume for simplicity that other investment funds that hold shares in i do not withdraw liquidity from i in response to the valuation shock $x^v_{i,t}$ that i suffered (but these funds may decide to withdraw liquidity from i in response to a liquidity shock that they themselves suffer, which is discussed in the funding contagion section). Furthermore, we assume that all external holders withdraw liquidity from i proportionally to the loss i suffered at the

[21] Note that losses on unlisted equity shares may not be immediately visible to market participants, as the issuing bank does not publically update its accounting equity on a daily basis like funds update their NAVs. Our shareholder contagion channel reflects the true (book-value) losses on the unlisted equity shares, which market participants can only approximate.

redemption rate R[22], which implies that

$$A^{vl}_{ii,t} = \epsilon_{i,t} R, \tag{12}$$

where $\epsilon_{i,t}$ denotes the fraction of fund i's shares held by external holders. As the aggregate value of all i's outstanding shares equals i's total asset value, the fraction $\epsilon_{i,t}$ is found by subtracting the aggregate value of shares in i held by other funds in our data from i's total asset value (and dividing the resulting difference by i's total asset value). The redemption rate R is a nondimensional constant of order one, which we assume for simplicity to be the same across all funds i from which shares are withdrawn. Estimates for other markets[23] suggest that the redemption rate R is typically somewhat smaller than one during normal times. We will also explore more adverse redemption rates $R > 1$, which may reflect conditions of reduced investor confidence during periods of market turmoil.

4.5 Results

Here, we study the contribution of the banking and investment fund sectors to the emergence of liquidity spirals for various pecking orders. We do so by comparing the stability of the full system to the stability of each sector individually. We can evaluate the individual stability of a sector by simply omitting all institutions from the shock transmission matrix that do not belong to that sector (i.e. omitting all banks when we evaluate the stability of the fund sector and vice versa). We can then calculate the dominant eigenvalue of this sector-specific shock transmission matrix, which gives the stability of the corresponding sector in isolation of the other sector. Note that the dominant eigenvalue of the full system is lower bounded by the eigenvalues of the individual sectors, because the full system includes both sectors as well as the shock propagation between them[24].

The shock transmission matrices are calculated for the time t corresponding to the snapshot date of our dataset. All results present the means over 50 random generations of the interbank assets network reconstruction. The network reconstruction does not strongly affect our results and therefore the standard errors of the means are too small to be visible in the plots.

[22] We show in appendix A.3 that this assumption implies that the *number* of shares withdrawn is a convex function of the NAV loss, as one would reasonably expect (see e.g. Cont and Wagalath, 2013, Wiersema et al., 2025).

[23] See e.g. Goldstein et al., 2017, Aikman et al., 2019, Mirza et al., 2020, Fricke and Wilke, 2023.

[24] From the Perron-Frobenius follows that replacing zeros by positive values in a non-negative matrix (which corresponds to adding together two sectors into the full system's shock transition matrix) cannot decrease the matrix' dominant eigenvalue.

4.6 Uniform Pecking Orders

In Figure 2, we analyze stability under the assumption that all institutions have the uniform pecking order. In Figure 2a, we set the redemption rate to $R = 1$ and vary the cap-to-depth ratio δ. While market depth affects the stability of both the banking and fund sectors, the overall system stability in Figure 2a is dominated by the banking sector. This is evident from the fact that the dominant eigenvalue of the full system is equal to that of the banking sector at every point.

Our findings indicate that market depth plays a crucial role in system stability. Specifically, as soon as the cap-to-depth ratio δ increases beyond one, the dominant eigenvalue of the full system exceeds one and hence a liquidity spiral emerges. A cap-to-depth ratio $\delta > 1$ implies a market depth smaller than the market capitalization. However, comparisons with estimates from other financial markets suggest that the market capitalization likely already serves as a conservative lower bound for the true market depth[25]. We nevertheless explore smaller market depths, because this is when liquidity spirals typically emerge in our results. The reader should thus take comfort in the facts that these represent highly adverse demand-side conditions that would only materialize during severe market turmoil, if ever.

In Figure 2b we examine the impact of the redemption rate R while the market depth is set equal to the market capitalization. As the redemption rate R only impacts the rate at which liquidity is withdrawn from funds, the redemption rate R affects the system's stability only through the fund sector. The redemption rate R noticably affects the system's stability, but the impact is not as pronounced as that of the market depth seen in Figure 2a. Note that the dominant eigenvalue of the banking sector is already close to one under the conservative assumption of a market depth equal to the market capitalization. As a result, only a moderate increase in the redemption rate to about $R = 4$ already triggers a liquidity spiral in the full system. Although the full system is unstable for a redemption rate of $R = 4$, the dominant eigenvalue of the fund sector is still substantially below the instability threshold at that point. This emphasizes the importance of capturing the interactions between banks and investment funds for a comprehensive evaluation of financial stability.

4.6.1 Central Bank Liquidity Provision

In Figure 2c and Figure 2d we examine how the SARB, acting as a lender of last resort, may attempt a liquidity injection to mitigate a spiral. Specifically, we analyze the effect of increasing banks' central bank reserves. Because banks use the uniform pecking order and therefore liquidate proportionally to asset values, raising the banks' SARB reserves

[25]Potters and Bouchaud, 2003, Lillo et al., 2003, Cont and Schaanning, 2019, Fukker et al., 2022, Sydow et al., 2024b

decreases the banks' reliance on other assets to meet liquidity shocks. As withdrawing SARB reserves does not cause contagion, while the liquidation of (most) other assets does, increasing the banks' SARB reserves stabilizes the system.

In Figure 2c we set the redemption rate to $R = 1$ and choose a severely shallow market depth, represented by a cap-to-depth ratio of $\delta = 2$. This leads to a liquidity spiral driven by the banking sector. For comparison's sake, we consider in Figure 2d the scenario where market depths are equal to the corresponding market capitalizations (i.e. $\delta = 1$), representing a conservative estimate of typical demand-side conditions, and the where the redemption rate is increased to $R = 5$. This combination of parameters results in a liquidity spiral driven by the interaction between both sectors. In both figures, the increase in banks' SARB reserves as a fraction of the banking sector's total liquid assets is presented on the x-axis.

In both figures, the liquidity spiral dissipates when the liquidity injection into the banking sector is sufficiently large. However, the effectiveness of this intervention differs between the two cases, because the system's dominant eigenvalue is lower bounded by the dominant eigenvalues of the individual sectors (as explained previously); in Figure 2d, the stabilizing effect of the liquidity injection is minimal because the dominant eigenvalue of the full system quickly becomes lower-bounded by the fund sector, whose individual stability is not affected by a liquidity injection into the banking sector. In contrast, Figure 2c demonstrates that the impact of the liquidity injection is stronger when the liquidity spiral is purely driven by the banking sector. Hence, the effectiveness of a liquidity injection into banking sectors depends on whether the instability is driven by the banking sector or by (interactions with) the fund sector.

Speaking more generally, interventions that target one sector, while another sector is the main cause of instability, will have little effect. Furthermore, when the instability is driven by the interaction between sectors, the effectiveness of targeting one sector will be more effective, but still limited by the individual stability of the other sector (as we saw in Figure 2d). Hence, effective regulatory intervention requires a clear understanding of which sector(s) primarily drive the instability. Consequently, methods that elucidate each sector's contribution to the instability, such as the framework proposed in this paper, are invaluable tools for designing targeted and appropriate policy responses.

It is important to note that the more effective liquidity injection of the two considered, i.e. the one in Figure 2c, still has a relatively modest impact when considering the substantial increase in banks' SARB reserves. The reason for this is that banks respond to liquidity shocks by liquidating a vertical slice across all liquid assets, rather than prioritizing their SARB reserves. As a result, for liquidity provision to be effective, it must constitute a significant portion of banks' total liquid asset pool. In section 4.8, we demonstrate that the potential of the intervention increases significantly when banks preferentially liquidate SARB reserves.

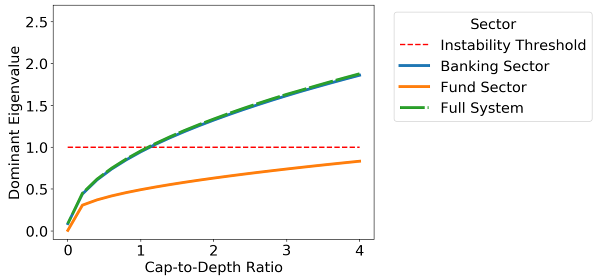

(a) Eigenvalue vs. Cap-to-Depth Ratio δ

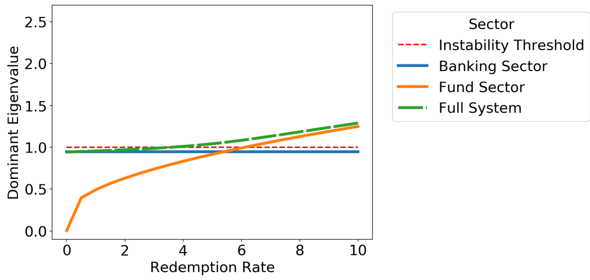

(b) Eigenvalue vs. Redemption Rate R

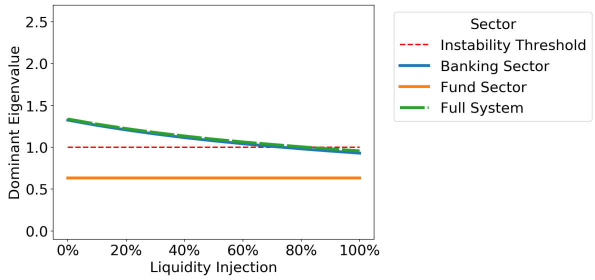

(c) Liquidity Provision for $\delta = 2$

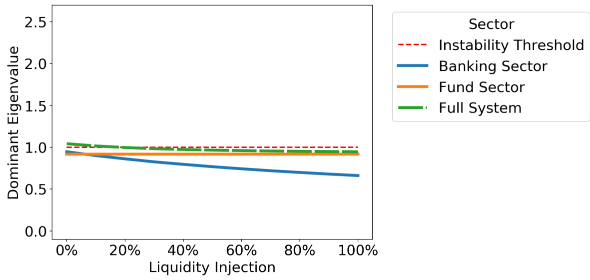

(d) Liquidity Provision for $R = 5$

Figure 2: **Stability of Uniform Pecking Orders.** All institutions are assumed to have the uniform pecking order. We explore when the dominant eigenvalue exceeds one and a liquidity spiral emerges for various cap-to-depth ratios δ in (a) and for various redemption rates R in (b). (c) shows the impact of a central bank cash injection into the banking sector when $\delta = 2$ and $R = 1$, and (d) shows this impact when $R = 5$ and $\delta = 1$.

4.7 Liquidity-Differentiated Pecking Orders

We now examine the stability of the liquidity-differentiated pecking orders, assuming that shocks do not exhaust the assets at the top of any institution's pecking order. Under this assumption, the stability of the system is only affected by the assets at the top of each institution's pecking order. Figure 1b presents the aggregate values of all assets that may occupy the top layer of institutions' pecking orders. This gives an indication of the maximum shock size that can be absorbed before these assets are exhausted. For instance, funds' government bonds holdings are depleted relatively quickly, whereas for banks, this is the case for miscellaneous securities. After analyzing how the liquidation of each asset type impacts stability, we extend our analysis in section 4.8 to consider shocks that exceed the assets at the top of institutions' pecking orders, forcing institutions to liquidate the assets that are next in line.

4.7.1 Market Depths' Impact on Stability

In Figure 3, we plot the dominant eigenvalue of the banking sector, the fund sector, and the full system for various pecking orders. We set the redemption rate to $R = 1$ and vary the cap-to-depth ratio δ to explore the impact of the market depth on stability. In Figure

3a, all institutions are assumed to have the optimistic pecking order, in Figure 3b the government bonds pecking order and in Figure 3c the conservative pecking order. Figure 3d shows the sector-specific pecking order, i.e. where all banks have the conservative pecking order and all funds the government bonds pecking order. As the stability of the full system is equivalent under the short-term funding and optimistic pecking orders, we primarily discuss the optimistic pecking order and show the results for the short-term funding pecking order in Figure 7a in Appendix A.1.

The results in Figure 3a show that the optimistic pecking order is very stable, because the liquidation of the assets at the top of this pecking order do not cause high levels of contagion. No liquidity spiral emerges even under highly adverse demand-side conditions, represented by cap-to-depth ratios of $\delta \geq 1$. The banking sector's dominant eigenvalue remains constant at zero because the banks' SARB reserves absorb liquidity shocks regardless of the market depth. Similarly, the assets at the top of banks' short-term funding pecking orders are also unaffected by the market depth (See Figure 7a in Appendix A.1). However, under the short-term funding pecking order, the dominant eigenvalue of the banking sector is somewhat larger than zero (~ 0.3), indicating that liquidity shocks propagate in the interbank deposits and repo network before dissipating.

Since funds do not hold any SARB reserves, they liquidate assets from the next (populated) layer of the optimistic pecking order instead. As a result, the optimistic and short-term funding pecking orders function identically for funds. Additionally, some funds do not hold deposits or fund shares either, forcing them to liquidate tradable securities from subsequent layers of their pecking orders. Through these funds that sell tradable securities, the market depth affects stability, as shown in the Figure 3a.

When government bonds are at the top of the pecking order, as shown in Figure 3b, the banking sector is highly sensitive to the market depth. The banking sector individually drives a liquidity spiral to emerge for market depths close to or smaller than the corresponding market capitalizations. As noted previously (see section 4.6), such market depths represent demand-side conditions likely to materialize only during periods of severe market distress, if at all. Furthermore, we explained in section 4.3 that LCR and RWA regulation generally requires banks to prioritize the conservative over the government bonds pecking order. Finally, banks may be, at least partially, hedged to market risk, but we do not include hedges due to data limitations. Hence, these results are not an immediate cause for concern, but they nevertheless identify a potentially dangerous feedback loop: Since banks hold the majority of the government bonds in our dataset, they bear the brunt of the price impact when these bonds are sold. Crucially, the banks strongly amplify the valuations shocks they incur through the high leverages that the banks maintain, which deteriorates the system's stability (Wiersema et al., 2023). Hence, the banking sector fireselling government bonds exemplifies the risk of a sector amplifying shocks through high leverages and propagating these shocks back onto itself.

Figure 3c shows that the market depth also influences the stability of the banking and fund sectors when they follow the conservative pecking order. However, the impact on the banking sector is significantly less pronounced than under the government bonds pecking order, as the miscellaneous securities at the top of the conservative pecking order are predominantly held by funds. Notably, unlike the previously discussed pecking orders, under the conservative pecking order the dominant eigenvalue of the full system is substantially higher than those of the individual sectors. This effect is driven by the interactions between the two sectors, which are not captured by the sectors' individual stability. Figure 3d illustrates that this effect is amplified under the sector-specific pecking order, where banks and funds intensify their interaction by preferentially fireselling securities in which the other sector has a larger position (i.e. funds sell government bonds while banks sell miscellaneous securities). Note that the individual stability of the banking and fund sectors is similar between Figures 3c and 3d, while system stability is much worse under the sector-specific pecking order. Hence, the increased interaction between the sectors under the sector-specific pecking order exacerbates instability in the system.

In both Figure 3c and Figure 3d, a liquidity spiral emerges even though the dominant eigenvalues of both individual sectors remain below one. This illustrates a critical risk: Ignoring cross-sector interactions could mask the onset of a liquidity spiral. The risk of overestimating stability by overlooking interactions between banks and NBFIs is likely to be present in many financial systems; the only reason is does not manifest more prevalently in our analysis of the South African financial system is because banks and funds predominantly invest in different securities. Whenever institutions sell securities in which another sector has a substantial position, stability is overestimated when the interaction between sectors is not captured. The risk of ignoring interactions underscores the importance of models that account for the interplay between different financial sectors. A comprehensive approach that captures these interactions, such as the one presented here, is essential for accurately assessing financial stability.

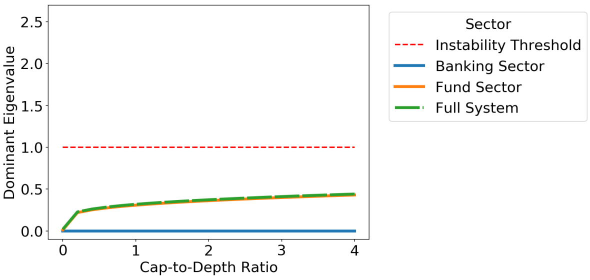

(a) Optimistic Pecking Order

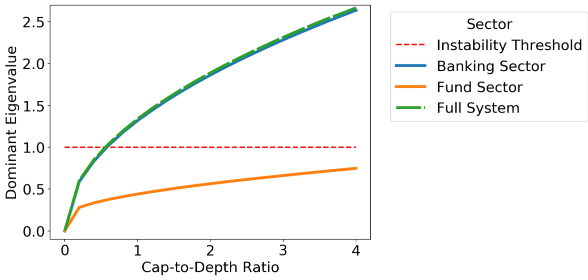

(b) Government Bonds Pecking Order

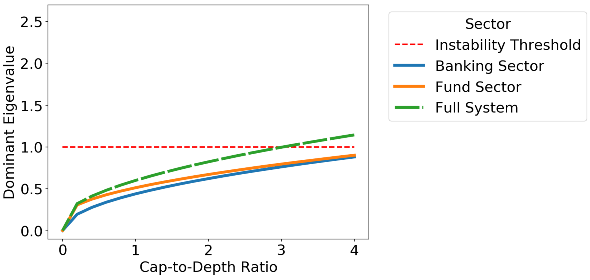

(c) Conservative Pecking Order

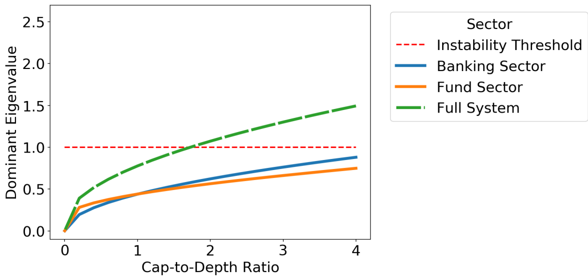

(d) Sector-Specific Pecking Order

Figure 3: **Stability of Liquidity-Differentiated Pecking Orders for Varying Market Depths.** We explore when the dominant eigenvalue exceeds one, signaling the onset of a liquidity spiral, across various market depths by varying the cap-to-depth ratio δ (representing severely stressed market conditions for $\delta \geq 1$) and liquidity-differentiated pecking orders. The redemption rate is fixed to $R = 1$. Panel (a) assumes all institutions follow the optimistic pecking order, (b) the government bonds pecking order, and (c) the conservative pecking order. In (d), institutions adopt the sector-specific pecking order, with banks following the conservative pecking order and funds the government bonds pecking order.

4.7.2 Redemption Rates' Impact on Stability

Figure 4 is analogous to Figure 3, but instead of varying the market depth, we vary the redemption rate R and use the securities' market capitalizations as a conservative estimate of their market depth. Since R is the rate at which fund shares are redeemed, it does not affect the individual stability of the banking sector, so its dominant eigenvalue remains constant across the plots in Figure 4. The results once again show the optimistic pecking order to be very stable as no liquidity spiral emerges in Figure 4a even for very high redemption rates. Under the short-term funding pecking order, the dominant eigenvalue of the fund sector remains identical to that of the optimistic pecking order and the dominant eigenvalue of the banking sector is again constant at a value slightly greater than zero ($\sim$0.3; see Figure 7b in Appendix A.1).

Similar to the results in the previous section, Figure 4b shows that under the government bonds pecking order, the banking sector individually causes a liquidity spiral when securities' market depths equal their market capitalizations, or shallower. As previously noted, this is not necessarily a cause for concern because the market capitalization is likely to be a conservative estimate of the true market depth, and because banks are likely to hedge (some) market risks and to prioritize the conservative over the government bonds pecking order. Nevertheless, note that higher redemption rates further exacerbate the instability. Moreover, higher redemption rates drive the system's dominant eigenvalue to exceed the banking sector's dominant eigenvalue, indicating that interactions between the banks and funds intensify the spiral.

Figure 4c illustrates that, while the conservative pecking order is generally more stable than the government bonds pecking order, it makes the system more sensitive to the redemption rate R. The main reason for this is that the miscellaneous tradable securities at the top of the conservative pecking order are predominantly held by the funds. As a result, funds bear the brunt of the price impact from selling these securities, causing these funds to play a larger role in determining the system's stability. As the redemption rate strongly affects whether funds dampen or amplify the shocks they receive, the redemption rate in turn has a larger impact on system stability under the conservative pecking order than under the government bonds pecking order.

The sector-specific pecking order in figure 4d triggers a system-wide liquidity spiral for redemption rates of about $R = 3$ or more. (As noted previously, redemption rates larger than one are nevertheless unlikely to materialize during times of benign market conditions[26].) Furthermore, while the individual stability of the fund sector is worse under the conservative pecking order, the sector-specific pecking order shows a wide a range of redemption rates that drive a liquidity spiral to emerge while neither of the individual sectors is unstable. As in Figure 3, we thus also see here that the increased interactions between the sectors due to the sector-specific pecking order worsen the stability of the system. This further underscores the critical importance of capturing these interactions.

[26]See e.g. Goldstein et al., 2017, Aikman et al., 2019, Mirza et al., 2020, Fricke and Wilke, 2023.

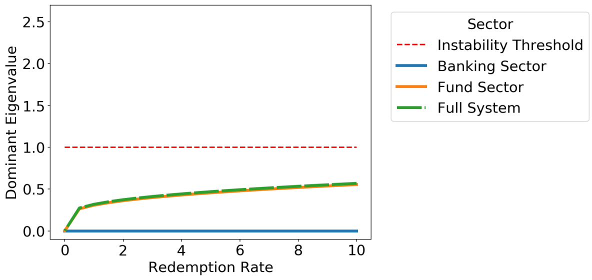

(a) Optimistic Pecking Order

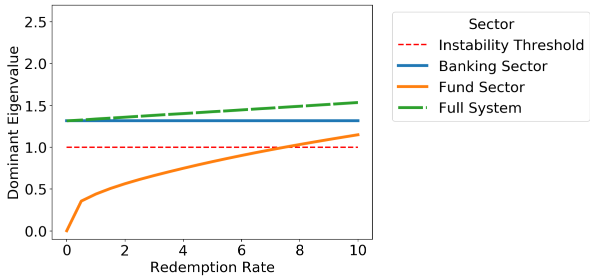

(b) Government Bonds Pecking Order

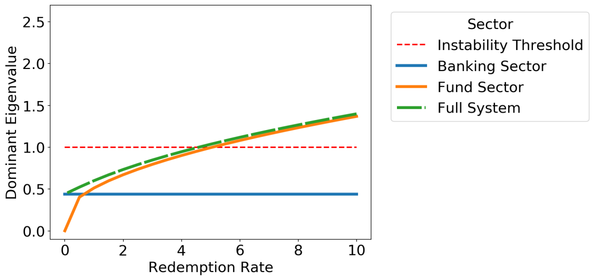

(c) Conservative Pecking Order

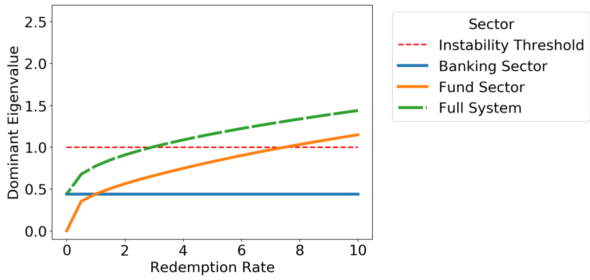

(d) Sector-Specific Pecking Order

Figure 4: **Stability of Liquidity-Differentiated Pecking Orders for Varying Redemption Rates.** We analyze when the dominant eigenvalue exceeds one, indicating a liquidity spiral, across various redemption rates R and liquidity-differentiated pecking orders. We use a security's market capitalization as a conservative estimate of its market depth (i.e. $\delta = 1$). Panel (a) assumes all institutions follow the optimistic pecking order, (b) the government bonds pecking order, and (c) the conservative pecking order. In (d), institutions adopt the sector-specific pecking order, with banks following the conservative pecking order and funds the government bonds pecking order.

4.8 Large Liquidity Shocks

In Figure 5, we examine how one or more large liquidity shocks impact stability when all institutions follow the optimistic pecking order (empirically the most common pecking order[27]). As shown earlier, the system is highly resilient to the liquidation of the assets at the top of this pecking order. However, when liquidity shocks deplete these assets, forcing institutions to liquidate those next in line, stability may deteriorate. Figures 5a and 5b illustrate how the dominant eigenvalue increases as large shocks increasingly deplete institutions' pecking orders, while Figures 5c and 5d depict the asset types that subsequently move to the top of institutions' pecking orders.

For simplicity, we reduce all institutions' liquidity buffers by the same proportion. While an institution's liquid asset holdings are likely calibrated, at least to some extent, to the liquidity outflows the institution anticipates (as dictated by LCR regulation for banks), real-world shocks will not be distributed perfectly proportional to institutions' liquidity buffers. Such a heterogeneous distribution of liquidity shocks across institutions could give rise to instabilities not captured by the approach used here. Future research should

[27]See e.g. Kim (1998), van den End and Tabbae (2012), Ma et al. (2020).

therefore examine financial stability under a broader range of liquidity shock scenarios, that vary not only in magnitude, but also in distribution across institutions.

In Figure 5a, the redemption rate $R = 1$, and we use securities' market capitalizations as an estimate of the market depths during the stressed demand-side conditions under which a liquidity spiral may emerge. The x-axis represents the proportion of an institution's total liquid assets depleted by large liquidity shocks. As explained previously, our model is not suitable to model defaults and becomes inaccurate as institutions approach default, because that introduces dynamics not captured by our model. To avoid the region where institutions are close to becoming illiquid, we consider depletions up to 80% of institutions' liquidity buffers. The figure demonstrates large shocks' destabilizing potential, as a liquidity spiral emerges once approximately half of institutions' liquid assets are exhausted and institutions are forced to liquidate assets lower on their pecking orders.

The liquidity spiral that emerges in Figure 5a is driven by the banking sector. Figure 5c illustrates the depletion levels at which the five largest banks transition to each of the subsequent layers of their pecking orders, as these correspond to clearly recognizable increments in the dominant eigenvalue of the banking sector in Figure 5a. The initial, gradual increase to a little over 0.4 corresponds to the banks transitioning from their SARB reserves to the short-term funding layers of their pecking orders. The eigenvalue then rises in two "steps" — first to approximately 0.8 as Firstrand transitions to the government bonds layer, and then to around 1.2 when Nedbank and Absa make the same transition. Finally, it rises to 1.3 as Investec and Standard Bank make the same transition.

These distinct increments in the dominant eigenvalue, which correspond to specific banks transitioning to subsequent layers of their pecking orders, potentially enable targeted interventions to counter the onset of a liquidity spiral by preventing these transitions. This is discussed in more detail in Section 4.8.1. However, similar to Figure 2d, Figure 5b demonstrates that at higher redemption rates, the banking sector no longer solely dictates the stability of the full system, diminishing the effectiveness of bank-targeted interventions.

For instance, Figure 5b shows that at a redemption rate of $R = 5$, which represents investors fleeing from the investment fund sector at severely high rates, the dominant eigenvalue of the fund sector and, accordingly, of the full system, already surpasses 0.7 when the shock size reaches just 4%. Figure 5d reveals that this increase corresponds to a significant proportion of funds transitioning from the short-term funding layer to the miscellaneous securities layer. As more funds make this transition, the system's dominant eigenvalue continues to climb, reaching approximately 0.9 purely due to the impact of the fund sector on the system's stability. The banking sector only begins to affect the system's stability when Firstrand transitions to liquidating government bonds. However, at that depletion level, the fund sector's contribution to the instability is already that strong that the effectiveness of a potential intervention targeting the banking sector is limited.

Finally, we compare the optimistic pecking order in Figure 5a to the uniform pecking order in Figure 2. For a redemption rate of $R = 1$ and market depths equal to the corresponding market capitalizations (i.e. $\delta = 1$), we find that the optimistic pecking order makes the system more resilient to small shocks than the uniform pecking order. However, unlike the uniform pecking order, which remains unaffected by liquidity shocks (as explained in Section 4.4), the optimistic pecking order may cause a liquidity spiral to emerge when a large shocks occurs. Thus, relative to the uniform pecking order, the optimistic pecking order creates a "robust-yet-fragile" system: highly resistant to small shocks but vulnerable to a single large shock. This mirrors the robust-yet-fragile network structures identified by Gai and Kapadia (2010), but here it arises from pecking order configurations. These findings underscore the importance of evaluating the resilience of financial systems against a wide range of shocks.

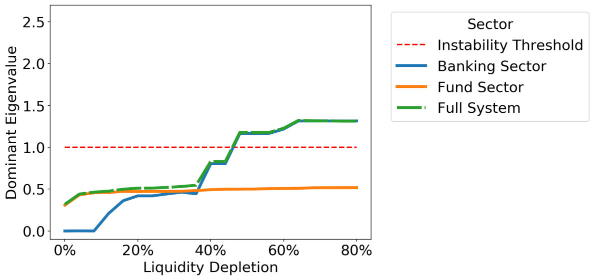

(a) Eigenvalue vs. Liquidity Shock ($R = 1$)

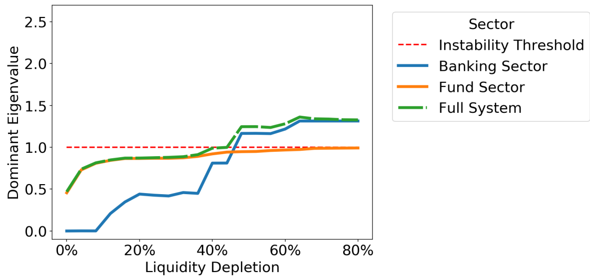

(b) Eigenvalue vs. Liquidity Shock ($R = 5$)

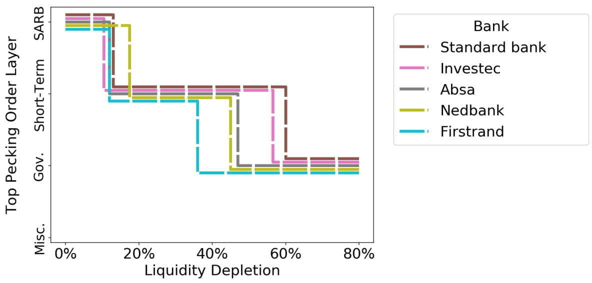

(c) Banks' Top Layer of Pecking Order

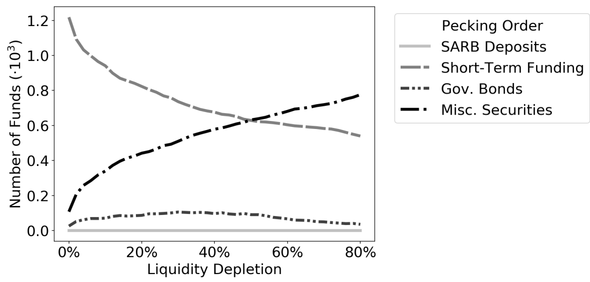

(d) Funds' Top Layer of Pecking Order

Figure 5: **Impact of Large Shocks.** We investigate how large liquidity shocks impact stability and whether they could push the dominant eigenvalue above one, triggering a liquidity spiral. All institutions are assumed to have the optimistic pecking order. The x-axis represents the proportion of each institution's total liquid asset pool depleted by large liquidity shocks and the y-axis gives the dominant eigenvalue of the system with the corresponding depletion level. In panel (a) the redemption rate is set to $R = 1$ and in (b) to $R = 5$. We use securities' market capitalizations as a conservative estimate of the securities' market depths (i.e. $\delta = 1$) in both plots. Panel (c) shows the depletion levels at which the five largest banks transition to each of the subsequent layers of their pecking orders and (d) shows, for each asset category, the number of funds that have it at the top of their pecking order depending on the depletion level.

4.8.1 Targeted Interventions

Figure 5 highlighted that the two sharpest increases in the system's dominant eigenvalue correspond to specific banks transitioning to the next layer of their pecking orders. To explore how the SARB may attempt to stabilize the system, we consider a set of targeted interventions in which the SARB provides one or more of these banks with an unlimited lender of last resort (LoLR) facility to absorb liquidity shocks, preventing them from liquidating assets lower on their pecking order. Because we rely on a stylized shock scenario and limited portfolio data, our model is subject to a margin of error. Therefore, the results in this section are illustrative and should not be interpreted as policy recommendations. Rather, they demonstrate the potential of the framework to inform policy when sufficiently detailed data on institutions' portfolios, contagion dynamics and propagating liquidity shocks are available.

Figure 6 illustrates the effectiveness of granting selected banks access to an unconditional LoLR facility, which allows them to avoid liquidating assets lower on their pecking order. Each intervention involves providing such a facility to one or more specific banks. Figure 6a, similar to Figure 5a, shows how the dominant eigenvalue of the system increases as banks' liquidity buffers are progressively depleted. However, for banks with access to the LoLR facility, the x-axis reflects the *hypothetical* depletion level they would experience in the absence of the facility, since they draw on this facility instead of their liquidity buffers. The figure focuses on depletion levels between 30% and 70%, as this is the range where banks' transitions to lower pecking-order layers most strongly affect the dominant eigenvalue. Each line in the figure reflects how the dominant eigenvalue evolves under a different intervention scenario. Note that the line corresponding to no intervention ("No Intervent.") is identical to the "Full System" result in Figure 5a.

Figure 6a shows that providing the LoLR facility exclusively to Nedbank is ineffective in preventing the system from crossing the instability threshold and entering a liquidity spiral. In contrast, extending the LoLR facility to Absa or Firstrand successfully prevents the onset of a liquidity spiral for depletion levels up to nearly 60%. Furthermore, the Firstrand intervention also substantially improves stability around depletion levels of about 40%, potentially averting further depletions by dampening early-stage shocks. (As long as the dominant eigenvalue remains low, liquidity shocks are quickly absorbed, reducing the risk of further depletion in institutions' liquidity buffers.) Importantly, all interventions that include Firstrand along with at least one other of th considered banks prevent instability across the entire range of examined range of depletion levels; while the dominant eigenvalue of the Nedbank+Absa intervention reaches just above the instability threshold for depletion levels around 65% or more, it remains just below the threshold for the Nedbank+Firstrand intervention. Furthermore, the Absa+Firstrand intervention has an even stronger stabilizing impact. Unsurprisingly, the intervention of granting all three considered banks access to the LoLR facility is most effective at stabilizing the system.

Table 6b presents each intervention's "cost", defined as the amount of liquidity provided by the SARB through the LoLR facility. The costs is expressed as a percentage of the banking sector's total liquid asset holdings, to facilitate comparison to the liquidity injections in Figure 2. Since banks draw on the facility in response to accumulated liquidity shocks, the intervention cost depends on the (hypothetical) depletion level reached. The second and third columns of the table report costs under depletions of 55% and 70%, respectively. Grayed-out cells indicate interventions under which the system becomes unstable before reaching the respective depletion level. The tipping point in the fourth column of the table marks the shock size at which the system crosses the instability threshold, triggering a liquidity spiral.

Table 6b indicates that the Absa intervention is slightly more cost-effective than the Firstrand intervention in stabilizing the system against depletion levels up to 55%. However, this comparison does not account for the added stability that the Firstrand intervention provides around depletion levels of 40%, where it significantly dampens early-stage contagion. Given the small cost difference between the two, the Firstrand intervention appears to be preferable. That said, under both interventions, the dominant eigenvalue is very close to the instability threshold at depletion levels near 55%. Considering model uncertainty and the ineffectiveness of either intervention beyond 60% depletion, a more prudent approach would be to extend the LoLR facility to both Firstrand and Absa. Finally, comparison with the uniform pecking order in Figure 2 highlights the enhanced effectiveness of central bank interventions when institutions follow the optimistic pecking order. Taking into account differences in market depth and redemption rates, Figure 6a nevertheless clearly shows that substantially larger reductions in the dominant eigenvalue can be achieved at a fraction of the cost when institutions use the optimistic pecking order.

In sum, these results illustrate that targeted liquidity interventions are particularly effective when institutions employ the optimistic pecking order, especially when informed by detailed data on institutions' portfolios and contagion dynamics. Crucially, the findings highlight that not all institutions are equally systemically relevant under stress, making interventions that target the right banks significantly more (cost) effective than others. While the analysis is illustrative, it demonstrates the framework's potential to support more effective and efficient policy responses once richer data become available.

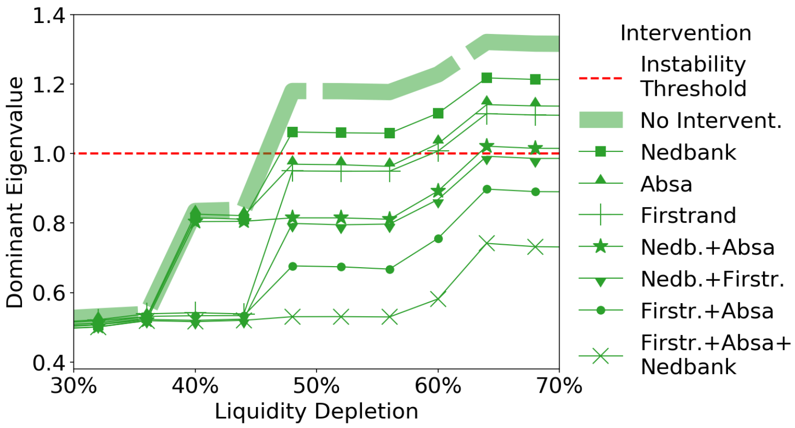

(a) Full System Stability per Intervention

Intervention	Cost 55%	Cost 70%	Tipping Point
No intervention	-	-	45%
Nedbank	4.9%	6.9%	47%
Absa	8.5%	11.5%	58%
Firstrand	9.3%	12.6%	59%
Nedbank+ Absa	13.4%	18.4%	63%
Nedbank+ Firstrand	14.2%	19.5%	>70%
Firstrand+ Absa	17.8%	24.0%	>70%
Firstrand+ Absa+ Nedbank	22.7%	31.0%	>70%

(b) Cost and Tipping Point

Figure 6: **Bank-Targeted Interventions.** The figure shows the effect of providing one or more banks with an unlimited Lender-of-Last-Resort facility to absorb liquidity shocks, preventing the banks from liquidating assets lower on their pecking order. Panel (a) shows the system's dominant eigenvalue for the corresponding depletion level applied to all institutions' liquidity buffers. The figure zooms in on depletion levels between 30% and 70%, where banks' transitions to lower pecking-order layers have the largest impact. Each line in Panel (a) corresponds to an intervention of providing the liquidity facility to one or more specific banks. The depletion level for which the system crosses the instability threshold is also shown in the fourth column ("Tipping Point") of panel (b). The middle columns of Panel (b) presents each intervention's "cost", i.e. the liquidity drawn from the SARB, as a percentage of the banking sector's total liquid assets. The amount drawn is determined by the liquidity depleted by accumulated shocks. The second and third columns in Panel (b) show intervention costs when depletion levels are 55% or 70%, respectively. Grayed-out cells indicate interventions under which the system becomes unstable before reaching the respective depletion level.

5 Discussion

Liquidity spirals progressively worsen market and funding liquidity (Brunnermeier and Pedersen, 2009). Our study examines liquidity spirals that arise from multiple interacting contagion channels and institutional types. To accurately assess these dynamics, models that take these interactions into account are required. We employ the framework developed by Wiersema et al. (2023), which enables the early identification of liquidity spirals before significant declines in market and funding liquidity occur. Wiersema et al. (2023) demonstrate that financial stability may be significantly overestimated when interactions between contagion channels are ignored. We extend this analysis to shocks of larger size, which may penetrate below the top layers of institutions' pecking orders, and show that neglecting the interplay between banks and investment funds (NBFIs), as well as between different contagion channels may lead to liquidity spirals being completely overlooked.

This framework allows us to assess how institutions' pecking orders influence the likelihood of liquidity spirals emerging without relying on any specific, subjective stress scenario. Our findings reveal that pecking orders play a critical role in financial stability, with some configurations resulting in a "robust-yet-fragile" system. This concept, initially observed by Gai and Kapadia (2010) in the context of certain network topologies, emerges here in relation to specific pecking order configurations. The identification of robust-yet-fragile tendencies across multiple dimensions underscores the risks of optimizing financial stability for frequent small shocks. A system highly resilient to minor shocks may prove exceptionally fragile when exposed to a large, unexpected shock. This finding highlights the necessity of stability measures that assess resilience across a broad spectrum of potential shocks, such as the eigenvalue-based approach developed in our study.

We apply our methodology to a granular dataset on the South African financial system, capturing the interdependent dynamics of the banking and investment fund sectors. Wiersema et al. (2025) show that exposures within the South African financial system are underestimated when interactions between these sectors are ignored. Our analysis identifies market conditions under which a liquidity spiral emerges that would be undetectable without accounting for these interactions. Furthermore, we demonstrate that certain liquidity spirals are driven predominantly by one sector, which has important implications for policy interventions, such as central bank liquidity provisions. Our model can thus serve as a valuable tool for policymakers in determining the most effective strategies for mitigating liquidity spirals, by targeting the right sector(s).

We have explored how the system's stability changes in response to a large liquidity shock. Our findings indicate that when institutions prioritize selling their most liquid assets first, the system remains resilient to small shocks. However, once a significant portion of institutions' pecking orders is depleted, a liquidity spiral may emerge. This robust-yet-fragile tendency is a systemic risk inherent in financial structures where institutions liquidate assets in order of decreasing liquidity, as contagion effects intensify when institutions are forced to sell increasingly illiquid assets. This emphasizes the importance of assessing financial stability across all layers of institutions' pecking orders.

The evolution of the system in response to the shock depends strongly on the distribution and magnitude of the shock. In this study, we have investigated liquidity shocks that are distributed proportionally to institutions' total liquid assets pools. Future research should develop realistic stress scenarios and investigate how our financial stability model evolves over time in response to these. Given the critical role that pecking orders play in the emergence of liquidity spirals, and the impact they have on the effectiveness of central bank interventions, further empirical research into the pecking order strategies that financial institutions employ under different market conditions is also warranted. Finally, we have identified market depths to play a crucial role in financial stability. Given the challenges in measuring market depths[28], substantial efforts should be directed towards better

[28]See e.g. Potters and Bouchaud, 2003, Lillo et al., 2003, Cont and Schaanning, 2019, Fukker et al., 2022.

understanding the price impact that results from extracting liquidity from the market.

References

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.

T. Adrian and H. S. Shin. Procyclical leverage and value-at-risk. *Review of Financial Studies*, 27(2):373–403, 2014.

D. Aikman, P. Chichkanov, G. Douglas, Y. Georgiev, J. Howat, and B. King. System-wide stress simulation. 2019.

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

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

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

C. Aymanns, F. Caccioli, J. D. Farmer, and V. W. Tan. Taming the basel leverage cycle. *Journal of financial stability*, 27:263–277, 2016.

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.

K. Bayram and A. H. A. Othman. The suitability of gold as a high-quality liquid asset: Empirical evidence from volatility structure analysis. In *Future of Organizations and Work After the 4th Industrial Revolution: The Role of Artificial Intelligence, Big Data, Automation, and Robotics*, pages 125–144. Springer, 2022.

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.

A. E. Bernardo and I. Welch. Liquidity and financial market runs. *The Quarterly Journal of Economics*, 119(1):135–158, 2004.

BIS. Basel iii: The liquidity coverage ratio and liquidity risk monitoring tools. Technical report, Bank of International Settlements, 2013.

BIS. Basel iii: the net stable funding ratio. Technical report, Bank of International Settlements, 2014.

F. S. Board. Board notice 90 of 2014. Technical report, 2014.

C. Borio, M. Drehmann, and K. Tsatsaronis. Stress-testing macro stress testing: does it live up to expectations? *Journal of Financial Stability*, 12:3–15, 2014.

J.-P. Bouchaud and R. Cont. A langevin approach to stock market fluctuations and crashes. *The European Physical Journal B-Condensed Matter and Complex Systems*, 6 (4):543–550, 1998.

J.-P. Bouchaud, J. D. Farmer, and F. Lillo. How markets slowly digest changes in supply and demand. In *Handbook of financial markets: dynamics and evolution*, pages 57–160. Elsevier, 2009.

G. Brandi, R. Di Clemente, and G. Cimini. Epidemics of liquidity shortages in interbank markets. *Physica A: Statistical Mechanics and its Applications*, 507:255–267, 2018.

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 2007-08. Technical report, National Bureau of Economic Research, 2008.

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

F. Caccioli, J. D. Farmer, N. Foti, and D. Rockmore. How interbank lending amplifies overlapping portfolio contagion: A case study of the austrian banking network. *arXiv preprint arXiv:1306.3704*, 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.

J. Choi, S. Hoseinzade, S. S. Shin, and H. Tehranian. Corporate bond mutual funds and asset fire sales. *Journal of Financial Economics*, 138(2):432–457, 2020.

R. Cont and E. Schaanning. Fire sales, indirect contagion and systemic stress testing. Technical report, Norges Bank Working Paper, 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. Kotlicki, and L. Valderrama. Liquidity at risk: Joint stress testing of solvency and liquidity. *Journal of Banking & Finance*, 118:105871, 2020.

L. De Haan and J. W. van den End. Banks' responses to funding liquidity shocks: Lending adjustment, liquidity hoarding and fire sales. *Journal of International Financial Markets, Institutions and Money*, 26:152–174, 2013.

R. A. De Santis. The euro area sovereign debt crisis: Identifying flight-to-liquidity and the spillover mechanisms. *Journal of Empirical Finance*, 26:150–170, 2014.

N. Detering, T. Meyer-Brandis, K. Panagiotou, and D. Ritter. An integrated model for fire sales and default contagion. *Mathematics and Financial Economics*, 15(1):59–101, 2021.

D. W. Diamond and P. H. Dybvig. Bank runs, deposit insurance, and liquidity. *The Journal of Political Economy*, 91(3):401–419, 1983.

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

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

J. D. Farmer, A. M. Kleinnijenhuis, P. Nahai-Williamson, and T. Wetzer. Foundations of system-wide financial stress testing with heterogeneous institutions. *Bank of England Staff Working Paper No. 861*, 2020.

A. Fostel and J. Geanakoplos. Leverage cycles and the anxious economy. *The American Economic Review*, 98(4):1211–1244, 2008.

D. Fricke and H. Wilke. Connected funds. *The Review of Financial Studies*, 36(11): 4546–4587, 2023.

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.

J. Gatheral. No-dynamic-arbitrage and market impact. *Quantitative finance*, 10(7):749–759, 2010.

J. Geanakoplos. The leverage cycle. *NBER macroeconomics annual*, 24(1):1–66, 2010.

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.

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

G. Hałaj. System-wide implications of funding risk. *Physica A: Statistical Mechanics and its Applications*, 503:1151–1181, 2018.

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

H. Jiang, D. Li, and A. Wang. Dynamic liquidity management by corporate bond mutual funds. *Journal of Financial and Quantitative Analysis*, 56(5):1622–1652, 2021.

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

C. E. Kim. The effects of asset liquidity: Evidence from the contract drilling industry. *Journal of financial Intermediation*, 7(2):151–176, 1998.

N. Kiyotaki and J. Moore. Credit cycles. *Journal of political economy*, 105(2):211–248, 1997.

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

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

Y. Ma, K. Xiao, and Y. Zeng. Mutual fund liquidity transformation and reverse flight to liquidity. *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*, 2020.

H. Mirza, D. Moccero, S. Palligkinis, and C. Pancaro. Fire sales by euro area banks and funds: What is their asset price impact? *Economic Modelling*, 93:430–444, 2020.

S. Poledna, J. L. Molina-Borboa, S. Martínez-Jaramillo, M. Van Der Leij, and S. Thurner. The multi-layer network nature of systemic risk and its implications for the costs of financial crises. *Journal of Financial Stability*, 20:70–81, 2015.

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.

J.-C. Rochet and X. Vives. Coordination failures and the lender of last resort: was bagehot right after all? *Journal of the European Economic Association*, 2(6):1116–1147, 2004.

T. Roukny, H. Bersini, H. Pirotte, G. Caldarelli, and S. Battiston. Default cascades in complex networks: Topology and systemic risk. *Scientific reports*, 3:2759, 2013.

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

SARB. Annual report 2016. Technical report, 2017a.

SARB. Financial stability review first edition. Technical report, 2017b.

A. Shleifer and R. W. Vishny. The limits of arbitrage. *The Journal of finance*, 52(1): 35–55, 1997.

M. Sydow, G. Fukker, T. Dubiel-Teleszynski, F. Franch, S. Gallet, H. Gründl, S. Kotronis, D. Miccio, M. Pellegrino, S. Schlütter, et al. Banks and non-banks stressed: liquidity shocks and the mitigating role of insurance companies. 2024a.

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, 2024b.

S. Thurner, J. D. Farmer, and J. Geanakoplos. Leverage causes fat tails and clustered volatility. *Quantitative Finance*, 12(5):695–707, 2012.

J. W. van den End and M. Tabbae. When liquidity risk becomes a systemic issue: Empirical evidence of bank behaviour. *Journal of Financial Stability*, 8(2):107–120, 2012.

G. Wiersema, A. M. Kleinnijenhuis, T. Wetzer, and J. D. Farmer. Scenario-free analysis of financial stability with interacting contagion channels. 2019.

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.

G. Wiersema, A. M. Kleinnijenhuis, E. Kemp, and T. Wetzer. *Higher-Order Exposures*. Number 3091. ECB Working Paper, 2025.

A Appendix

A.1 Additional Results

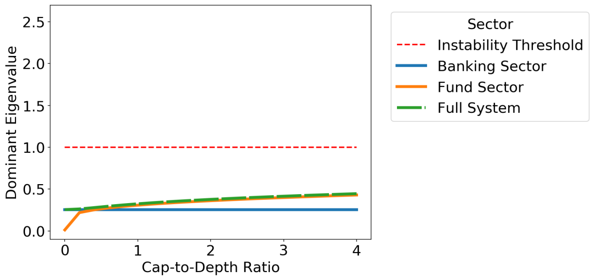

(a) Short-Term Funding Pecking Order

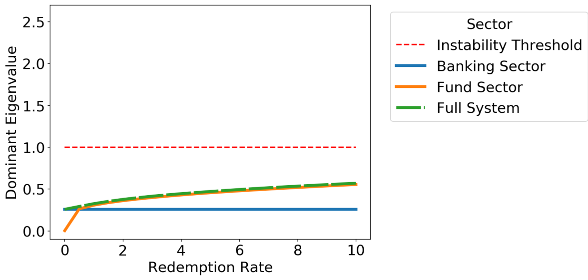

(b) Short-Term Funding Pecking Order

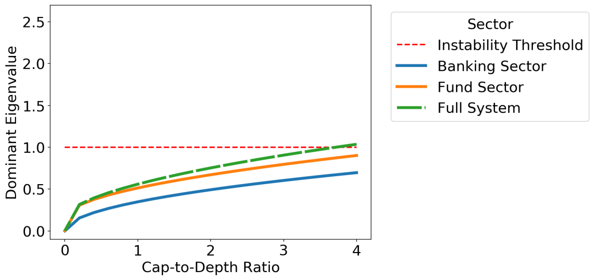

(c) Conservative Pecking Order

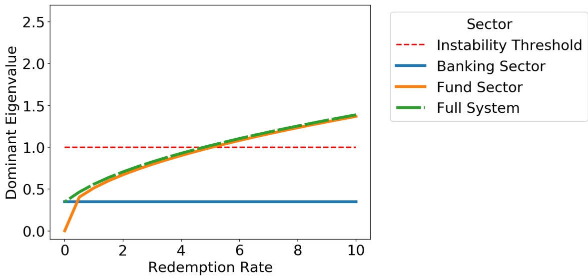

(d) Conservative Pecking Order

Figure 7: **Additional Results.** Panels (a) and (b) show the stability of the short-term funding pecking order for various market depths and redemption rates, respectively. The graphs are very similar to the results discussed in section 4.5. Panels (c) and (d) illustrate the stability of the conservative pecking order for various market depths and redemption rates when gold is included as a miscellaneous asset. Given that gold is traded globally, we assume it has infinite market depth for simplicity. Consequently, liquidating gold does not cause contagion, and the system in Figures 7c and 7d is slightly more stable than the system in Figures 3c and 4c. For other pecking orders, the stabilizing effect of gold is too small to be visible in the plots.

A.2 Liquidity Spirals Driven by Funding Contagion

Here, we explain that the funding contagion channel, given by the upper-left quadrant of the shock transmission matrix, is the only contagion channel out of all channels considered in this analysis that can cause a liquidity spiral in the absence of other contagion channels (i.e. when the other quadrants are zero): From the properties of block matrices, we know that the largest eigenvalue is zero when the only non-empty contagion quadrant is an off-diagonal quadrant. Hence, neither deleveraging contagion (upper-left quadrant) nor overlapping portfolio contagion (lower-right quadrant) can individually drive the emergence of a liquidity spiral. Furthermore, when the lower-right quadrant, i.e. the share redemption and shareholder contagion quadrant, is the only non-empty quadrant, the largest eigenvalue is positive but only valuation shocks propagate (while liquidity shocks dissipate immediately, as can be seen from eq. 4). Hence, the only quadrant of the shock transmission matrix that can cause a liquidity spiral when all other quadrants are zero is the upper-left quadrant, i.e. the funding contagion channel.

Furthermore, in the absence of other contagion channels, the funding contagion channel can only cause a liquidity spiral when banks hoard liquidity in response to shocks; when a bank does *not* hoard liquidity, it only withdraws deposits to meet the liquidity shock it incurred up to the magnitude of the shock. Hence, the aggregate liquidity that the bank withdraws from other banks does not exceed the liquidity shock incurred, so the bank does not amplify the propagating shock. The sum of a column in the funding contagion quadrant of the shock transmission matrix gives the aggregate liquidity that a bank withdraws from other banks as a fraction of the liquidity shock that the bank incurred (see eq. 4). Hence, absent liquidity hoarding, column sums in the funding contagion quadrant cannot exceed one. Furthermore, the Perron-Frobenius theorem guarantees that the largest eigenvalue of shock transmission matrix does not exceed the largest column sum. Therefore, absent other contagion channels, i.e. the largest eigenvalue is given by the funding contagion quadrant, and absent liquidity hoarding, i.e. the column sums in the funding contagion quadrant cannot exceed one, the largest eigenvalue is upper-bounded by one and no liquidity spiral can emerge. Hence, when banks do not hoard liquidity, studying contagion channels in isolation, and thus ignoring their interactions, as is often done, will overlook any liquidity spiral even in the most unstable of systems.

A.3 Withdrawal of Fund Shares

When we discussed the share redemption contagion channel in section 4.4, we assumed that the amount of liquidity withdrawn by external investors from the investment fund is linear in the NAV loss of the fund's shares. Here, we show that this assumption implies that the number of shares withdrawn is a convex function of the fund's NAV loss.

When investment fund i suffers a loss $x^v_{i,t}$, the NAV of the fund's shares falls by

$$\Delta NAV_{i,t} = x^v_{i,t}/S_{i,t}, \tag{13}$$

where $S_{i,t}$ denotes i total number of outstanding shares at time t. Furthermore, the amount paid out per share that investors redeem is given by

$$NAV_{i,t+1} = NAV_{i,t} - \Delta NAV_{i,t}, \tag{14}$$

and we denote the number of shares withdrawn by the external investors in response to the NAV loss $\Delta NAV_{i,t}$ as $\Delta S_{i,t}$. The amount paid out for the redeemed shares gives the liquidity withdrawn from the fund;

$$(NAV_{i,t} - \Delta NAV_{i,t})\,\Delta S_{i,t} = \epsilon_{i,t} R x^v_{i,t}, \tag{15}$$

where we have used the factor $\epsilon_{i,t}R$ from equation (12) to express the amount of liquidity withdrawn in terms of the loss $x^v_{i,t}$. Using (13), we find that

$$\Delta S_{i,t} = \min\left\{\frac{\epsilon_{i,t}RS_{i,t}\Delta NAV_{i,t}}{NAV_{i,t} - \Delta NAV_{i,t}}, \epsilon_{i,t}S_{i,t}\right\}, \tag{16}$$

where we have used that $\Delta S_{i,t} \leq \epsilon_{i,t}S_{i,t}$ (as only the fraction of i's shares held by external holders can be withdrawn through shareholder contagion). Hence, the number of shares withdrawn in response to the NAV loss is linear for small NAV losses and convex for larger NAV losses (until the upper bound of $\Delta S_{i,t}$ is fixed).

A.4 Interbank Asset Allocation

As the contagion channels that we model require the counteparties of interbank deposits, repo, and unlisted equity shares to be specified, we reconstruct the interbank networks for these assets from the aggregate holdings on banks' balance sheets. Note that our model does not distinguish between more granular types of tradable securities than what is listed in the banks' balance sheet data, so banks' positions in these securities are not reconstructed.

The technique used for the reconstruction of the banks' investments is similar to Kok and Montagna (2013) and Wiersema et al. (2025), and aims to reproduce the high heterogeneity and fat-tailed distributions of interconnections observed in real financial networks. We randomly assign banks' deposits, repo and unlisted equity holding to banks with significant liabilities of the corresponding type remaining after subtracting liabilities to funds (as given by the fund data). As our results show the system to be resilient against the withdrawal of liquidity from deposits or repo loans, and because interbank equity holding are minor, the impact of the specific algorithm used to approximate the true interbank network is expected to be small.

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
$\beta \in \{$deposits, repo, (un)listed equity shares$\}$:

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.

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.[29]

4. The investment is added to the balance sheets of y and z, and the value of the investment is 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.

After step 5, the counterparties of all (relevant) assets and liabilities are defined.

A.5 Price Impact of Number of Shares Sold

When we discussed the overlapping portfolio contagion channel in section 4.4, we assumed that the price impact $\Delta p_{\sigma,t}$ is linear in the liquidity recovered from the sale. Here, we show that this assumption implies that the price impact is concave in the number of shares sold.

Let us assume for simplicity that institution i is the only institution that sells shares in security σ. The derivation generalizes straightforwardly to the case when multiple institutions sell shares in σ at the same time. For notational convenience, we assume that security σ is the only asset at the top of institution i's pecking order (and that the shock $x^l_{i,t}$ does not exhaust the asset), such that the price impact (8) reduces to

$$\Delta p_{\sigma,t} = \frac{x^l_{i,t}}{D_\sigma} = \frac{(p_{\sigma,t} - \Delta p_{\sigma,t})\,\Delta n_{\sigma i,t}}{D_\sigma}, \tag{17}$$

where $\Delta n_{\sigma i,t}$ denotes the number of shares in security σ that i sells at time t to raise liquidity $x^l_{i,t}$ and we have used the assumption that all shares are sold against the new price $p_{\sigma,t+1} = p_{\sigma,t} - \Delta p_{\sigma,t}$. Rewriting equation (17), we find the price impact as a function of the number of shares sold:

$$\Delta p_{\sigma,t} = \frac{p_{\sigma,t}\Delta n_{\sigma i,t}}{D_\sigma + \Delta n_{\sigma i,t}}, \tag{18}$$

which is linear in the number of shares sold when $\Delta n_{\sigma i,t}$ is small (similar to e.g. Cont and Schaanning, 2019 and Wiersema et al., 2019) and concave in the number of shares sold when $\Delta n_{\sigma i,t}$ is large (see e.g. Gatheral, 2010). Furthermore, note from equation (18) that $\Delta p_{\sigma,t} < p_{\sigma,t}$ so the price cannot become negative.

[29] 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 South African Rand.

A.6 Market Price of Listed Equity Shares

The overlapping portfolio and shareholder contagion mechanisms in section 4.4 both affect the market price of listed equity shares issued by banks. Therefore, both contagion channels should take into account losses to the market value caused by the other contagion channel. For the overlapping portfolio contagion channel, this achieved automatically: Because the shareholder contagion channel decreases both the market price and depth of the shares by the same proportion, it is immediately obvious from equation (6) that the price impact remains unaffected. The reason is that, when shareholder contagion has decreased the market price, the price impact per share sold decreases accordingly, but, due to the decreased price, the number of shares one need to sell to raise the required amount of liquidity increases by the same factor, so they cancel out.

When overlapping contagion drives the market value of shares below their book value, we can realistically assume that the impact of the shareholder contagion channel is diminished, as there is simply less market value left to be marked down in response to a decline in book value. Assuming that losses in the market value of the shares are proportional to book value losses, we can reflect the impact of the overlapping portfolio contagion channel on the shareholder contagion channel by multiplying equation (11) by the shares' market-to-book ratio; let $E_{\sigma,t}$ denote the equity at time t of the institution that issued the shares, and $S_{\sigma,t}$ the total number of shares that the institution issued, such that a share's book value is given by $E_{\sigma,t}/S_{\sigma,t}$ and market-to-book ratio by $p_{\sigma,t}S_{\sigma,t}/E_{\sigma,t}$. Multiplying equation (11) with the market-to-book ratio yields

$$A^{vv}_{\sigma i,t} = s_{i\sigma,t}\frac{p_{\sigma,t}}{E_{\sigma,t}}, \tag{19}$$

such that when institution σ suffers a loss $x^v_{\sigma,t} = E_{\sigma,t}$, the shareholder contagion suffered by i equals $s_{i\sigma,t}p_{\sigma,t}$. However, as the results in this paper are derived for time $t = 1$, when the shares' market values are assumed to be equal to their book values, the market-to-book ratio $p_{\sigma,t}S_{\sigma,t}/E_{\sigma,t} = 1$, so we have omitted it from equation (11) for simplicity.

A.7 Market Capitalization Estimates

We estimate market capitalizations C_σ for six different categories σ of domestic tradable securities: Government bonds, listed equity shares and bonds issued by the non-financial corporate sector, and MMIs, listed equity shares and bonds issued by the banking sector. As explained, we do not distinguish between different types of tradable securities within these categories due to data limitations.

For (domestic) government bonds, we use the market capitalization of South African government bonds at the end of 2016[30]. Note that, in our data, the banks and funds

[30]The market capitalization of South African bonds as of Q4 2016 is sourced from the Q1 2017

collectively own 21% of the total market capitalization of the South African government bonds. Due to data limitations, we estimate other market capitalizations by assuming that the banks and funds own the same fraction of 21% of other securities' market capitalizations. Hence, the market capitalization of any security category is given by banks' and funds' collective holdings of that security category divided by 21%. This has the analytical advantage that securities' market depths are proportional to the banks' and funds' holdings of these securities. Consequently, differences in contagion dynamics across these securities are caused by the distribution of these securities across institutions (which are known from our data), rather than the market depths of these securities (which are difficult to estimate)[31].

SARB Quarterly Bulletin; https://www.resbank.co.za/content/dam/sarb/publications/quarterly-bulletins/quarterly-bulletin-publications/2017/7718/07Statistical-tables—Public-Finance.pdf.

[31]When market depths across all securities are proportional to banks' and funds' holding of these securities, and security α has double the market depth of β, then the price impact of raising an amount of liquidity x from selling security α causes half the price impact of raising it from β. The price impact of α, however, affects twice the amount of banks' and funds' positions in the security as that of β, so banks' and funds' collective losses are the same regardless whether x liquidity is raised from α or β.

Acknowledgements

The views expressed are those of the authors and do not necessarily reflect those of the ECB. We thank Rama Cont, Fabio Caciolli and the participants of seminars at the Institute of New Economic Thinking at the Oxford Martin School, ETH Zurich, and the South African Reserve Bank for their helpful discussions.
The authors take full responsibility for the contents of this paper and report limited use of Generative AI to improve readability. This paper should not be reported as representing the views of the European Central Bank (ECB).
The views expressed in this paper are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

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

Esti Kemp
International Monetary Fund, Washington D.C., United States; email: esti.kemp@imf.org

J. Doyne Farmer
Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, United Kingdom; Smith School of Enterprise and the Environment, University of Oxford, Oxford, United Kingdom; Macrocosm Inc, Brooklyn New York, United States; Santa Fe Institute, Santa Fe, New Mexico, United States; email: dorothy.nicholas@inet.ox.ac.uk

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

PDF ISBN 978-92-899-7619-0 ISSN 1725-2806 doi:10.2866/2539834 QB-01-26-006-EN-N