Scenario-Free Analyses of Financial Stability with Interacting Contagion Channels

Garbrand Wiersema
Somerville College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2022

Statement of Originality

I certify that I am the first author on all papers presented in this thesis.

Candidate: Garbrand Wiersema

Signature:

Date: October 7, 2022

“Everything should be made as simple as possible, but no simpler.”

Albert Einstein

This thesis is dedicated to:

Willy Wiersema (1928-2020)

&

Annie Vermaas (1928-2018)

Acknowledgements

I am grateful to Professor J. Doyne Farmer for encouraging me to pursue my research interests and for providing the guidance and inspiration needed to excel in this project. I would also like to thank my parents Edwin and Trudy Wiersema and sister Sharon IJpma for their unconditional support of my endeavors. My research greatly benefited from the intellectually stimulating environment provided by the Mathematical Institute at the University of Oxford and the Institute for New Economic Thinking at the Oxford Martin School, and would not have been possible without the generous support of Rabobank, the Prins Bernhard Cultuurfonds, the Henrik Mullerfonds, the KAS BANK studiefonds, and the Ketel 1 studiefonds.

Abstract

The Great Financial Crisis of '07/'08 highlighted the dangers of instabilities in financial systems. In unstable financial systems, an initially small and localized financial shock may be amplified and spread throughout the system in a process referred to as *financial contagion*. Various channels of contagion exist, and propagate shocks through a deterioration of the liquidity and solvency of institutions, decreases in the market prices of tradable securities, etc. Contagion channels interact and amplify one another, such that the channels' collective impact may far exceed the sum of the individual channels' contributions.

The current state-of-the-art for studying the (endogenous) *systemic risk* generated by interacting contagion channels are so-called "system-wide stress tests". These stress tests subject a financial system to an initial stress scenario and simulate how interacting contagion channels propagate the stress across the system. Stress tests provide valuable insights into systemic risk and the stability of financial systems. These insights are limited in scope, however, as a stress test's results are conditional on the (often subjectively imposed) initial stress scenario. We show that measures of systemic risk due to interacting contagion channels that forego the reliance on an initial stress scenario contribute greatly to the understanding of financial stability.

Exposures are an important measure of risk that does not rely on an initial stress scenario. When assessing an institution's exposure to the default of a counterparty, traditional exposure measures focus on direct exposures. Since the Global Financial Crisis, indirect exposures via common asset holdings are increasingly recognised too. Yet direct and indirect exposures fail to capture the losses that result from shock propagation and amplification following the counterparty's default. We refer to those spill-over losses as *higher-order exposures* and contribute to the literature on financial exposures by proposing a way to formalize and

quantify them. Using granular data of the South African financial system and a contagion model that captures the most commonly studied contagion channels and their interactions, we show that higher-order exposures make up a significant part of exposures – particularly during times of financial distress when exposures matter most. We also show that higher-order exposures cannot simply be extrapolated from direct or indirect exposures, since they depend strongly on the network structure and the robustness of individual institutions.

As the higher-order exposures we propose capture various contagion channels and their interactions without reliance on an initial stress scenario, they are an objective measure of systemic risks that may go unnoticed when using other measures. Like scenario-based stress tests, however, as a measure of systemic risk higher-order exposures are limited in scope because they only capture the losses that follow from the (idiosyncratic) default of an institution in the system. Yet, financial systems may be subjected to various other kinds of shocks.

Eigenvalue-based approaches study the inherent tendency of a financial system to dampen or amplify shocks, providing a holistic measure of systemic risk. Current eigenvalue-based approaches, however, only handle a single contagion mechanism. We develop an eigenvalue-based approach that gives the best of both worlds, allowing analysis of multiple, interacting contagion channels without the need to impose a subjective stress scenario. This allows us to demonstrate that the instability due to interacting contagion channels can far exceed that of the sum of the individual channels acting alone, which highlights the importance of capturing these interactions. We also derive an analytic formula in the limit of a large number of institutions that elucidates the mechanisms through which interactions between contagion channels amplify instabilities.

We apply the developed eigenvalue-based approach to the data of the South African financial system to study the risk of liquidity spirals emerging that consist of various contagion channels and/or span multiple sectors. We refer to these as complex liquidity spirals and show that the intensity of these liquidity spirals may be severely underestimated when interactions between contagion channels or sectors are overlooked. We capture the collective stability of the banking sector and investment fund sector and show that the stability of the South African financial system strongly depends on how institutions choose to respond to a liquidity shock, with some choices yielding a "robust-yet-fragile" system. We also show that liquidity spirals are exacerbated when the liquidity of institutions worsens, and that central bank-provided liquidity can greatly dampen liquidity spirals. We study the banking and investment fund sectors' individual contributions to the liquidity spiral and find that market conditions determine which of the two sectors is the main driver of the spiral. The approach developed here can be used to formulate interventions that specifically target the sector that is causing the liquidity spiral.

Contents

List of Figures

Chapter 1

Introduction

Systemic risk refers to the endogenous risk in a financial system that emerges from the interactions between the institutions that make up the system. The Great Financial Crisis of '07/'08 highlighted the dangers of systemic risk by revealing large exposures between seemingly unconnected parts of the financial system and amplifying initially localized shocks to global proportions (Adrian et al., 2008). Systemic risk is relevant to the institutions that make up the system, as prudent strategy would seek to limit exposure to systemic risk, but even more important to the regulators and policymakers whose responsibility is to control system risk. To properly understand and control systemic risk, the interactions through which risk spreads must be modeled accurately. This problem is complicated by the fact that the financial system is heterogeneous, with many different types of actors that interact through a potentially even greater number of different mechanisms. This makes the financial system a prime example of a *complex system*.

The hallmark of complex systems is *emergence*, which refers to the tendency of a complex system to show properties that arise through interactions between the system's components and that are not immediately obvious from individual mechanisms that operate at the microscopic level in the system (Holland, 2000). Through emergence, risk controls that make individual institutions more robust in isolation may exacerbate financial distress when institutions interact. Those that seek to understand the risks in a financial system, and in particular those that seek to model the emerging systemic risks accurately, should therefore model the financial system as a complex system.

The interactions through which financial risks spread are referred to as *contagion channels* (Allen and Gale, 2000, Battiston and Martinez-Jaramillo, 2018). We consider the most important and commonly studied contagion channels (Aymanns et al., 2018): In a financial system where institutions have large investments in one another, the default of single institution may set off a chain of defaults, i.e. a "default cascade". *Counterparty default contagion* refers to the propagation of losses through write-downs that institutions suffer on investments in a counterparty when that counterparty defaults. Even if the counterparty remains solvent when it incurs a shock, the expected value of the counterparty's liabilities falls (which causes losses to the creditors) because its probability of default has increased. This probabilistic propagation of losses is referred to as *counterparty risk contagion. Funding contagion* describes how an institution that incurs a liquidity shortage may withdraw (short-term) funding from other institutions to raise liquidity, which propagates the liquidity shock. Financial distress may also drive an institution to liquidate its position in a tradable security, which depresses the security's market price. This *price impact* causes other institutions that have a position in the same security suffer mark-to-market losses. The propagation of losses through (fire)sales of tradable securities is referred to as *overlapping portfolio contagion.* When an institution incurs a loss, margin calls, leverage targets, and/or other risk controls may force the institution to raise liquidity in order to deleverage. Hence, the shock is propagated *within* the institution, from its solvency to its liquidity. We refer to this as *deleveraging contagion.*

When deleveraging contagion forces institutions to liquidate assets, this liquidation may cause further contagion. In particular, when the institutions liquidate their tradable securities, the resulting overlapping portfolio contagion may force further deleveraging. This is referred to as a liquidity spiral and provides an important example of a positive feedback loop in financial systems (Brunnermeier and Pedersen, 2009). In general, positive feedback loops are an important driver of instability in the system. Like liquidity spirals, most positive feedback loops in financial systems are made up of multiple, interacting contagion channels. Consequently, when studying individual contagion channels in isolation from other channels, such feedback loops are overlooked. This yields an optimistic picture of stability, with potentially very dangerous consequences. To assess financial stability accurately, the collective impact of all (relevant) contagion channels and their interactions must be considered.

The most commonly used tool for evaluating and safeguarding financial stability are so-called *stress tests.* The term stress test describes a wide range of analyses that study the consequences of subjecting individual institutions or (components of) a financial system to a stress scenario. Nowadays, banks (and increasingly other financial institutions too) are required to meet regulatory stress tests, which subject the balance sheets of financial institution to a macroeconomic stress scenario and evaluate the post-stress balance sheets against a set of predetermined requirements (Foglia, 2018). These regulatory stress tests serve to make institutions more resilient to macroeconomic adversity. Crucially, they do not consider systemic risk, as the stress scenario is applied to the balance sheet in isolation from other financial institutions' responses to the stress. More advanced so called "system-wide" stress tests model (large components of) financial systems and the interactions of institutions therein to gain a more holistic view of financial stability that reflects the risk of contagion.

System-wide stress tests are the current state-of-the-art for studying the (endogenous) systemic risk generated by interacting contagion channels and yield potentially highly valuable insights into the stability of financial systems (Aymanns et al., 2018). By collectively subjecting a set of financial institutions to a stress scenario and simulating how contagion channels propagate the stress between institutions, system-wide stress tests may assess the impact of various contagion channels and their interactions. (All institutions that make up the financial system, i.e. banks, funds, insurers, brokers, etc., should ideally be included in the analysis, although typically only a subset of these institutions is modeled due to data and methodological limitations.) The downside of this approach is, however, that the results of any stress test are conditional on the stress scenario. Typically, stress tests consider only one or a small number of stress scenarios, limiting the scope of the analysis and raising questions about the inherently subjective nature of the chosen of scenario(s). This limitation can be partially overcome by considering a larger set of scenarios. However, the set of potential future stress scenarios is virtually infinite and axiomatically includes scenarios that drive the entire system to collapse, whereas the probability of any scenario to occur not immediately obvious. Furthermore, when an unstable scenario is identified, tools for unraveling what dynamics drive the instability are limited, which makes formulating effective interventions challenging. Hence, system-wide stress tests should be complemented by alternative measures of financial stability to better understand systemic risk. Such measures, which capture multiple interacting contagion channels without reference to an initial stress scenario, are explored in this thesis.

1.1 Outline

The overarching research question of this thesis is how scenario-independent measures of financial stability with interacting contagion channels can contribute to the understanding systemic risk. In chapter 2, we investigate to what extent the literature answers this question and describe how chapters 3-5 contribute to this literature.

Chapter 3 is titled Higher-Order Exposures. Exposures are an important measure of the financial risks faced by individual institutions and can be quantified without reference to any stress scenario. An institution's exposure to a counterparty is typically understood to capture the maximal loss the institution stands to suffer when the counterparty defaults. When assessing the institution's exposure to the counterparty, traditional exposure measures focus on direct exposures. Since the Global Financial Crisis, indirect exposures via common asset holdings are increasingly recognized too (Cont and Wagalath, 2013, Caballero and Simsek, 2013). Yet direct and indirect exposures fail to capture the spill-over losses that result from shock propagation and amplification following the counterparty's default. In chapter 3, we introduce the concept of "higher-order exposures", which captures these spill-over losses that result from shock propagation and amplification following the counterparty's default. We contribute to the literature on financial exposures by proposing a way to formalize and quantify higher-order exposures. Only when including higher-order exposures as part of exposure do financial exposures properly reflect the (systemic) risk of contagion.

Using granular data of the South African financial system, we show in chapter 3 that higher-order exposures make up the vast majority of the exposures in a substantial part of the South African financial system – particularly during times of financial distress, when exposures matter most. Hence, recognizing higher-order exposures is essential for controlling institutions' exposure to (systemic) risk. We also show that higher-order exposures depend strongly on the network structure and the robustness of individual institutions. This highlights the importance of the complex systems approach that models the structure of the financial network and the characteristics of the institutions therein as accurately as possible. It also emphasizes the importance of collecting granular data on the contracts between institutions, as the network cannot be modeled accurately without.

Higher-order exposures are a substantial improvement over traditional exposures and provide an objective measure of systemic risk, but are nevertheless limited in scope; although they capture various interacting contagion channels and do not rely on an initial stress scenario, they only consider the losses that would follow the (idiosyncratic) default of an institution in the system. Yet, financial systems may be subjected to various other kinds of shocks; even a system in which all institutions' exposures to one another are very low is not guaranteed to be resilient in general, as the system may include shock-amplifying dynamics not captured by the exposures. Furthermore, the method used for quantifying (higher-order) exposures in chapter 3 focuses on institutions' solvency, while contagion channels that propagate through liquidity shocks may also contribute substantially to systemic risk (Hałaj, 2018, Cont et al., 2020). For these reasons, we develop a measure of systemic risk in chapter 4 that elucidates a system's underlying shock-amplifying and/or dampening dynamics and considers both shocks to institutions' solvency and the institutions' liquidity.

Chapter 4 is titled Scenario-Free Analysis of Financial Stability with Interacting Contagion Channels and develops an eigenvalue-based measure of financial stability. Eigenvalue-based approaches study the inherent tendency of a financial system to dampen or amplify shocks without reference to any (subjectively imposed) stress scenario. The eigenvalue-based approaches that have been proposed so far, however, only handle a single contagion mechanism (Caccioli et al., 2014, Bardoscia et al., 2017, Cont and Schaanning, 2019). We show that the most commonly studied contagion channels and their interactions can be expressed in terms of the propagation of liquidity shocks and valuation shocks, and the conversion of liquidity shocks into valuation shocks and vice versa. We develop an eigenvalue-based stability measure in chapter 4 that captures this *Solency-Liquidity Nexus* of interacting contagion channels without reference to any stress scenario and show that the instability due to interacting channels may far exceed the sum of the individual channels acting alone. This highlights the importance of capturing these interactions.

We derive an analytic formula in the limit of a large number of institutions that gives the instability threshold as a function of the relative size and intensity of contagion channels, providing valuable insights into financial stability whilst requiring very little data to be calibrated to real financial systems. We generate random financial systems of various sizes and densities and show that the analytic solution accurately approximates the stability even of the systems with very few institutions. We also find that the sparsest systems are generally less stable than their denser counterparts.

Chapter 5 is titled Complex Liquidity Spirals. In chapter 5, we apply the eigenvalue-based approach developed in chapter 4 to the data set on the South African financial system studied in chapter 3. Chapter 5 complements the insights into the stability of the South African financial system derived in chapter 3 by focusing on liquidity risk, and by providing a general measure of stability rather than focusing only on exposures. In chapter 5, we study the risk of liquidity spirals emerging in the South African financial system that consist of various contagion channels and/or span multiple sectors. We refer to these as complex liquidity spirals and show that the intensity of these liquidity spirals may be severely underestimated when interactions between contagion channels or sectors are overlooked.
We capture the collective stability of the banking sector and investment fund sector and identify conditions for which a liquidity spiral emerges that is overlooked when considering the stability of each sector individually. This highlights the importance capturing the interactions between sectors. Conversely, we also identify conditions for which one of the two sectors predominantly drives the liquidity spiral. In particular, we find that central bank liquidity provision to the banking sector can greatly dampen a liquidity spiral caused by the banking sector, but has little effect when the spiral originates in the fund sector. The approach developed here can be used to identify whether a liquidity spiral is predominantly driven by an individual sector, or by the interactions between sectors. Furthermore, the approach can be used to formulate interventions that specifically target the sector(s) that is (are) causing the liquidity spiral.
We explore how the stability of the South African financial system depends on how institutions choose to respond to a liquidity shock. We show that the stability of the system greatly depends on what assets institutions choose to liquidate to meet the shock. In particular, we find that when institutions liquidate assets in order of liquidity, the system shows a "robust-yet-fragile" tendency; i.e. it may be very resilient to small liquidity shocks, in response to which institutions only liquidate their most liquid assets, but become highly unstable when institutions have liquidated all of their most liquid assets in response to a large liquidity shock and subsequently must rely on their assets of lesser liquidity to meet liquidity demands. This highlights the importance of stability measures that assess a system's resilience to a wide range of shocks, such as the eigenvalue-based approach developed here.
Chapter 6 concludes by summarizing our findings and discussing avenues for further research into the topic of systemic risk.

Chapter 2

Literature

This chapter reviews the existing literature on the systemic risk of financial contagion and discusses how this thesis contributes to the literature.

2.1 Risk in Financial Systems

The future is uncertain. Uncertainty combined with exposure to future outcomes creates risk. The financial industry is particularly involved in dealing with risk; banks charge interest rates based on risks, insurers collectivize risks, hedge funds construct low-risk portfolios (whilst aiming for higher returns than other investments with similar risk profiles), etc. Financial risk is therefore an area of great interest that has been studied extensively over the years, and a wide range of measures of financial risk have been proposed (see e.g. Altman and Saunders [1997], Crouhy et al. [2000] for an overview).

2.1.1 Exposures

One measure of financial risk is (counterparty) exposure. An institution's exposure to a counterparty is typically understood to refer to the maximum loss that the institution stands to suffer when that counterparty defaults. Hence, exposures forgo the probabilistic question of how likely the counterparty is to default and simply quantify the losses that the institution is exposed to *if* the counterparty defaults. Traditional measures of exposure focus on *direct exposures*, where an institution's direct exposure to a counterparty is measured as the current value of the institution's investments in (financial instruments issued by) the counterparty. Direct exposures have been studied extensively as part of counterparty credit risk[1], where

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

the expected loss is calculated as the product of the exposure at default (EAD) to the counterparty, the probability of default (PD) of the counterparty and the % loss given default (LGD). Direct exposures feature prominently in financial regulation, playing an important role in e.g. large exposure limits (BIS, 2018b), minimum risk-based capital requirements (BIS, 2019), and the identification of Global Systemically Important Banks or GSIBs (BIS, 2018a).
Since the Great Financial Crisis, *indirect exposures* are increasingly recognized as an important component of exposure too.[2] Two institutions are indirectly exposed to one another when both have a position in the same tradable security. When either liquidates its position, the sudden influx of selling orders in the order book drives down the market price, which causes mark-to-market losses to the other institution. (This is explained in more detail in sec. 2.2.3 Overlapping Portfolio Contagion.) Through indirect exposures, substantial exposures may exist between institutions that have little or no *direct* exposures to one another, with the institutions that generate the largest indirect exposures not necessarily having the largest balance sheets (Cont and Schaanning, 2019). This shows that indirect exposures are qualitatively different from direct exposures (as direct exposures are very strongly correlated with balance sheet size). Direct and indirect exposures to a counterparty do not, however, capture the spill-over losses that result from propagation and amplification of the financial distress following the counterparty's default.

2.1.2 Systemic Risk

The endogenous risk of propagation and amplification of distress in financial systems is referred to as *systemic risk*. Systemic risk emerges from the interconnections in the system.[3] A comprehensive overview of research on systemic risk can be found in e.g. Fouque and Langsam [2013] or Aymanns et al. [2018]. The Global Financial Crisis impressively highlighted the dangers of systemic risk. Losses that were initially localized to the US subprime mortgage market were amplified by several orders of magnitude over the course of the crisis and spread across the globe and across various sectors of the financial and real economy. Haldane [2013] likens the crisis to the consequences of the '02/'03 SARS outbreak, when uncertainty led to panic and overreaction causing major harm to the Asian economy, while

[2]See e.g. Cont and Wagalath [2013], Caballero and Simsek [2013], Greenwood et al. [2015], Clerc et al. [2016], Cont and Schaanning [2017], Calimani et al. [2017], Aymanns et al. [2018], Cont and Schaanning [2019], Aldasoro et al. [2020].

[3]See e.g. Gai and Kapadia [2010], Gai et al. [2011], Cont et al. [2010], Glasserman and Young [2016].

the SARS outbreak, by epidemiological standards, turned out to be quite modest. Adrian et al. [2008] estimate that the initial losses in the US subprime mortgage market were on the same order of magnitude of only a one percent loss in the US stock market – which often occurs on a daily basis. To understand why some shocks fizzle out almost immediately while other (initially small) shocks may set off a chain of events that eventually brings the entire system to its knees, we must recognize that the financial system is a *complex system*.

2.1.3 Complex Systems

Complex systems are large systems consisting of many different components that interact through various mechanisms. The financial system is a prime example of a complex system (Aymanns et al., 2018); it is a dynamical system that includes many different actors, such as (shadow) banks, funds, insurers, retail investors, etc. Financial systems are highly interconnected, arising from many market activities besides the typically studied interbank exposures (Battiston and Martinez-Jaramillo, 2018); actors interact through the various types of contracts between them, which represent loans, bonds, equities, derivatives, as well as other financial instruments and other contractual requirements. As these contracts may vary widely in nature, the corresponding interactions may vary widely too.
Complex systems have a tendency to show *emergence*. Emergence refers to properties of a complex system that are not immediately obvious from the system's building blocks and mechanisms at the microscopic level, but instead emerge from the interactions between the various components of the system (Holland, 2000). Emergent phenomena can be observed across complex systems in various disciplines (de Haan, 2006). In financial systems, emergence explains how some initially small and localized shocks may cause a global crisis while others do not. Emergence also explains how risk controls that make individual institutions more robust in isolation, may exacerbate financial stress when institutions interact (see e.g. Aymanns and Farmer [2015], Aymanns et al. [2016]). The complex nature of financial systems should thus be taken into account both when studying systemic (i.e. emergent) risk, as well as when seeking to control it.

2.1.4 Multiplex Networks

Capturing the complex nature of the financial system is highly non-trivial and requires both advanced modeling techniques and highly detailed data. Financial institutions invest in assets of many different types and the networks of institutions' investments of a specific type can vary widely across asset types (Poledna et al., 2015, Hüser and Kok, 2019). All investment networks corresponding to the various asset types should ideally be captured such that the systemic risk due to the interactions between asset types is not underestimated. For example, Poledna et al. [2015] find that systemic risk is underestimated by up to 90% when focusing on a single asset type. Similarly, the various types of institutions in the system and the heterogeneity among them should be captured too (Farmer et al., 2020).

To capture the various types of institutions and assets, the financial system can be represented as a multiplex network (i.e. a multilayer network in which each layer has the same set of nodes), where the nodes represent the institutions and each network layers represents a single asset type. The state of each node is given by the corresponding institution's balance sheet and a (weighted, directed) edge in a layer of the network gives the value of an institution's investment of the layer's type. As this network representation only includes nodes corresponding to (financial) institutions, its edges only represent investments *between* the institutions. Therefore, the financial system is sometimes modeled as bipartite network, where a second set of nodes is added to the network to represent the tradable securities that institutions invest in. For example, Caccioli et al. [2014], Cont and Schaanning [2019] use a bi-partite network representation of the financial system to capture the indirect exposures between institutions.

Although some qualitative results may be derived without explicit calibration, to derive meaningful quantitative results, this network representation must be calibrated to the financial system. Accurate calibration requires contract-level bilateral data on institutions' individual investments, specifying, at a minimum, both counterparties, the type of the investment, and its current value. Until recently, investment networks were typically reconstructed from aggregate exposures[4], as disaggregate, contract-level data were rarely available. Now, disaggregate data are becoming more widely available and, although data collection efforts should continue to expand into the future, studies are already starting to take advantage of them.[5]

Although having detailed network data on the structure of the financial system is necessary for understanding systemic risk it is not sufficient, as models of how financial distress spreads across and between the asset networks is also required.

[4]See e.g. Musmeci et al. [2013], Mastrandrea et al. [2014], Mazzarisi and Lillo [2017], Squartini et al. [2017], Di Gangi et al. [2018], Gandy and Veraart [2019].

[5]See e.g. Battiston et al. [2016a], Bookstaber et al. [2016], Lux [2016], Aldasoro and Alves [2018], Berndsen et al. [2018].

2.2 Contagion Channels

Financial distress spreads across assets through the interactions between financial institutions. This is referred to as (financial) contagion and has been extensively studied. Although a comprehensive review of this literature is beyond the scope of this thesis, we note that important foundations for understanding the mechanisms that drive contagion were laid around the turn of the century.[6]. More specific contributions are discussed below.

Contagion can spread through decreases in the solvency and liquidity of institutions, through depreciation of tradable securities, etc. The literature typically distinguishes between three channels of contagion (Aymanns et al., 2018): Counterparty default (risk) contagion, overlapping portfolio contagion, and funding contagion. Financial distress may also be exacerbated by margin calls, leverage requirements and other risk controls that force institutions to deleverage during crises. We discuss these contagion channels now.

2.2.1 Counterparty Default Contagion

When an exogenous shock exhausts an institution's equity, causing the institution to become insolvent and default, the value of the institution's liabilities is written-down to their recovery value. Hence, the institution's creditors suffer write-downs on their investments in the institution, where the write-down on an investment is given by the product of the investment's EAD and LGD. The propagation of losses through the default of counterparties is referred to as counterparty default contagion and has been extensively studied.[7]

When institutions are highly exposed to one another, the initial default may set off a cascade of defaults, spreading losses throughout the system. Eisenberg and Noe [2001] assume zero bankruptcy costs, so the recovery value of an institution's liabilities is equal to its asset value that remains after incurring the shock. (Note that, regardless of bankruptcy costs, the value of an institution's issued equity shares is by definition completely written-off when the institution defaults through insolvency.) In this case, there is no amplification and the shock is eventually absorbed by the equity of the institutions in the system, bringing the default cascade to a

[6]See e.g. Kiyotaki and Moore [1997], Shleifer and Vishny [1997], Allen and Gale [2000], Eisenberg and Noe [2001].

[7]See e.g. Gai and Kapadia [2010], Battiston et al. [2012], Elliott et al. [2014], Acemoglu et al. [2015], Amini et al. [2016].

halt. The aggregate loss of the institutions' shareholders is equal to the exogenous shock (as there is no amplification) and Eisenberg and Noe [2001] show that a unique solution to the final distribution of losses across the shareholders exists. The assumption of zero bankruptcy costs, however, is (almost always) optimistic. The framework proposed by Eisenberg and Noe [2001] is expanded to include bankruptcy costs in e.g. Rogers and Veraart [2013] and Cont et al. [2010]. More specifically, the time taken to resolve a default is much longer than the typical timescales over which the contagion process unfolds, so in the context of contagion recovery rates may realistically be assumed to be zero (Elsinger et al., 2006, Cont et al., 2010). When bankruptcy costs are included, amplification may occur, causing the aggregate losses of the institutions' shareholders to exceed the exogenous shock.

2.2.2 Counterparty Risk Contagion

Even when the exogenous shock does not exhaust the equity of the institution and it remains solvent, contagion may still occur. As the value of the institution's liabilities would be written-off upon default, the expected value of the liabilities of any institution with non-zero default risk is below the liabilities' nominal value; the higher the institution's probability of default, the lower the expected value of its liabilities. This is reflected in capital regulation (e.g. BIS [2019]), modern accounting standards (e.g. IFRS [2021]), and recent contagion literature (see e.g. Battiston et al. [2012, 2016b], Bardoscia et al. [2015, 2017]). When an institution incurs a loss, its probability of default rises and the expected value of its liabilities falls, causing losses to its creditors. This is referred to as counterparty risk contagion. In accounting and capital calculation, sophisticated models are used to calculate expected losses on counterparty exposures. Bardoscia et al. [2017] focus on banks and argue that a bank's probability of default must be a convex function of the (equity) loss it suffers. They derive (the majority of) their results assuming zero recoveries and a linear probability of default function (as a conservative upper bound to the assumed convex function), which implies that when a bank's equity falls, the expected value of its liabilities falls by the same proportion. Bardoscia et al. [2017] define the interbank leverage as the bank's total debt to other banks over its equity and find that when no bank's interbank leverage exceeds one, the propagating shock is not amplified and is eventually damped out, bringing the contagion process to a halt (but the final aggregate losses exceed the initial exogenous shock because recoveries are assumed to be zero).

2.2.3 Overlapping Portfolio Contagion

As was already alluded to when discussing indirect exposures, overlapping portfolio contagion describes the mark-to-market losses that institutions suffer when an institution with which they have a portfolio overlap is forced to liquidate (part of) its portfolio of tradable securities.[8] An institution may be forced to sell tradable securities to raise liquidity (as discussed in sec. 2.2.5 Deleveraging Contagion) or as part of the resultion process when it defaults (Burrows et al., 2012, Caccioli et al., 2015).

The sale of the tradable securities depresses the securities' market price, which is referred to as the sale's price impact and depends, among other things, on the average daily trading volume and daily volatility of the security.[9] The price impact may be assumed to be linear in the sales volume (e.g. Cont and Schaanning [2019]) or log-linear (Bouchaud and Cont, 1998, Farmer, 2002, Caccioli et al., 2014). Empirical evidence suggests that under normal conditions (when execution is slow enough for the order book to replenish between successive trades) the price impact is concave in the sales volume and may be approximated accurately by a square-root function (Bouchaud et al., 2009). Yet, when sales volumes become very large (as they do in a firesale), the price impact function may become more linear or even convex (Gatheral, 2010). The appropriate choice of price impact function is thus somewhat ambiguous and depends on the conditions that are being studied.

2.2.4 Funding Contagion

When an institution relies on short-term funding to finance its activities, it is vulnerable to that funding being withdrawn. In particular, when one of the institution's financiers needs to raise liquidity, it may do so by withdrawing (part of its) short-term funding from the institution. Hence, the financier propagates its "liquidity shortage" to the institution. We refer to this is as funding contagion. Relative to the counterparty default and overlapping portfolio contagion channels, the funding contagion channel has received little attention in the literature so far, with notable exceptions being Gai et al. [2011], Anand et al. [2015] and Hałaj [2018]. Note that the financier may also withdraw funding from the institution as a risk mitigation strategy when it anticipates the institution's default. Such risk-mitigating

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

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

strategies may cause an institution that suffers losses to default through illiquidity before it becomes insolvent. This highlights that liquidity and solvency are not independent but in fact are closely connected, so models of financial stability should reflect this (Hałaj, 2018, Cont et al., 2020).

2.2.5 Deleveraging Contagion

Margin calls, leverage requirements and other risk controls may force an institution to deleverage when it incurs losses. Deleveraging requires liquidity to pay off debt so when an institution is forced to deleverage, the institution effectively incurs a liquidity shortage (unless it has surplus liquidity that may be used to deleverage). Through the margin call or leverage requirement (or potentially another risk control), a shock to an institution's solvency is thus propagated to the institution's liquidity. We refer to this mechanism as deleveraging contagion.

In the literature, deleveraging contagion is generally considered as part of another contagion channel. Treating it as a separate contagion channel, however, may be more insightful as discussed in chapter 4. When an institution is forced to deleverage, it may raise liquidity by withdrawing short-term funding from other institutions, such that the deleveraging causes funding contagion. This mechanism is studied in Gai et al. [2011]. Alternatively, the institution may raise liquidity by selling securities, such that the deleveraging causes overlapping portfolio contagion. The resulting losses due to the price impact may demand further deleveraging from this and other institutions, resulting in additional sales (and so on). This positive feedback loop is referred to as a liquidity spiral (Brunnermeier and Pedersen, 2009). Positive feedback loops are an important source of emergence in complex systems and the risk of liquidity spirals makes the complex systems approach to modeling financial stability vital. The particular liquidity spiral explained here highlights how a leverage requirement or other risk control may be prudent in isolation but drives instability when institutions interact. The feedback loop between leverage and market liquidity has been extensively studied.[10] Modeling of liquidity spirals has allowed for reproducing some empirically observed features of financial time series, such as fat tails in the distribution of returns and clustered volatility (Cont et al., 2010), and realistic pre- and post-shock volatility dynamics (Poledna et al., 2014), that cannot be reproduced by models that do not take the complex nature of the financial system into account.

[10] See e.g. Fostel and Geanakoplos [2008], Brunnermeier and Pedersen [2009], Geanakoplos [2010], Thurner et al. [2012], Gorton and Metrick [2012].

Financial institutions are empirically observed to control their leverage; mark-to-market losses due to falling market prices would increase institutions' leverages if not controlled, yet institutions' leverages typically do not to rise during periods of falling prices (Adrian and Shin, 2010). In particular, Adrian and Shin [2010] observe that security broker dealers increase leverage over periods of rising market prices and decrease leverage over periods of falling market prices. This is referred to as a procyclical leverage strategy as it exacerbates the liquidity spiral; progressively withdrawing market liquidity accelerates the downward trend in markets (while the opposite is also true as procyclical leverage also leads to progressively injecting market liquidity when prices are rising, which accelerates the upward trend in the prices.) Aymanns and Farmer [2015] and Aymanns et al. [2016] show that such procyclical leverage dynamics may result from a Value-at-Risk (VaR) leverage constraint (as was required under the Basel II regulatory framework).

The overarching takeaway from the literature related to deleveraging contagion is that leverage is an important source of instability in financial systems. Ceterus paribus, the higher the leverage, the more unstable the system. This implies that high levels of leverage must be offset by sources of damping of financial shocks for financial systems to remain stable. In Aymanns and Farmer [2015] and Aymanns et al. [2016], the dampening force is provided by "fundamental value investors", that inject market liquidity when liquidity spirals have driven market prices to fall below their fundamental value. Yet, sources of damping of financial shocks in complex systems, however vital, have not been studied more generally.

2.2.6 Interacting Contagion Channels

As systemic risk emerges from the interactions in the financial system, it emerges from interactions between contagion channels too. Models of financial stability should thus include the interactions between contagion channels to yield a comprehensive picture. Yet, only few studies include multiple interacting contagion channels and those that do are typically limited to two contagion channels.[11]

Using data from the Austrian banking system, Caccioli et al. [2015] show that the systemic risk due to interacting channels of counterparty default and overlapping portfolio contagion greatly exceeds the sum of the risks due to the individual contagion channels. Kok and Montagna [2016] reach a similar conclusion by combining counterparty default, overlapping portfolio, and funding contagion in a single

[11] See e.g. Arinaminpathy et al. [2012], Kok and Montagna [2016], Caccioli et al. [2015], Poledna et al. [2015].

model. As interactions between contagion channels strongly amplify risks, ignoring these interactions may greatly overestimate stability.

2.3 Stress Testing

The current state-of-the-art for analyzing systemic risk due interacting contagion channels is to perform a "system-wide stress test".[12] The term *stress test* refer to a wide range of methods that assess financial stability and resilience to crises by subjecting (part of) the financial system to a stress scenario and evaluating the consequences. Stress scenarios are designed to be "severe yet plausible"; severe enough to meaningfully challenge the system yet plausible enough to be relevant (Quagliariello, 2009). The choice of stress scenario should represent potential future crisis conditions, and is of great consequence to the outcomes of the stress test. The scenario should therefore be designed very carefully. Yet some degree of subjectivity is always present in the choice of scenario as the nature of future crises cannot be know in advance.

Before the Great Financial Crisis, the International Monetary Fund (IMF) already subjected countries' financial systems to stress tests, and banks performed in-house stress tests for internal risk management under the Market Risk Amendment of the Basel I Capital Accord. Since the crisis, however, stress tests have become decidedly more prominent, as central banks made safeguarding financial stability one of their core objectives and stress tests one of their main tools for doing so (Armour et al., 2016). Nowadays, central banks regularly subject the banks under their supervision to regulatory stress tests, and stress tests have become an essential element of the Basel II framework (Foglia, 2018). Stress tests are increasingly performed on non-banks too, including insurers (Jobst et al., 2014, BoE, 2019, EIOPA, 2021), pension funds (EIOPA, 2019), and CCPs (ESMA, 2015, CFTC, 2016, IMF, 2016).

2.3.1 Regulatory Stress Tests

The regulatory stress tests that banks are subjected to are *microprudential* in nature. That is, the bank is subjected to the stress test in isolation from other banks' (and other financial institutions') responses to the stress scenario. Hence, the risk

[12]See e.g. Burrows et al. [2012], Kok and Montagna [2016], Caccioli et al. [2015], Cont and Schaanning [2017], Dees and Henry [2017], Aikman et al. [2019b].

of contagion is not captured by the stress test. Knowing this, central banks tend to design stress scenarios that are, arguably, unrealistically severe (Borio et al., 2014). The macroeconomic stress scenario is typically designed using a macroeconometric model (Foglia, 2018). The stress scenario consists of a deterioration of macroeconomic variables such as GDP, unemployment rates, housing prices, etc. and of financial variables such as interest rates and credit spreads. A "satellite model" is used to translate the macroeconomic variables into quantities that directly affect the bank's balance sheet, with the most important one being credit risk. The designed scenario is applied to the bank's balance sheet and the impact on its capital (and liquidity) ratios is calculated and compared to required levels. (Note that liquidity is not always considered.) When the bank's capital ratios do not meet required levels, the bank will be obligated to raise additional capital.

Regulatory stress tests have proven to be effective at recapitalizing banks (Armour et al., 2016). As leverage is an important driver of systemic risk, increasing capital improves financial stability. Regulatory stress tests may also raise confidence in the resilience of banks by providing insight into the otherwise opaque balance sheets of banks (Bookstaber et al., 2014). Lack of confidence in the banking sector may in itself be an important driver of instability (Diamond and Dybvig, 1983, Brunnermeier, 2009). Furthermore, stress tests should, at least in theory, incentivize risk averting or mitigating strategies. However, regulatory stress tests have two distinct shortcomings: The first is that stress tests rely on a subjectively imposed stress scenario. By its very nature, the scenario cannot capture the full spectrum of potential stress events that may materialize in the future and hence stress tests can only assess financial stability very selectively. Moreover, the subjective nature of the scenario raises debate about its realism and incentives banks to "pass the test" rather than to increase their resilience to stress, which potentially explains why stress tests' outcomes seem to have become less informative over time (Borio et al., 2014, Candelon and Sy, 2015, Glasserman and Tangirala, 2015). The second and arguably most important shortcoming of regulatory stress tests is that they do not capture system risk due to their microprudential nature and hence overlook the most important driver of financial crises (Aymanns et al., 2018).

2.3.2 System-wide Stress Tests

We refer to stress tests that do, to some extent, capture systemic risk as "system-wide" stress tests. They do so by simulating how financial distress propagates across institutions following an initial stress scenario through the contagion channels discussed above. As the name suggests, a system-wide stress test ideally covers the entire financial system (including not only banks, but also funds, insurers, brokers, etc.). In practice, however, system-wide stress tests only capture part of the financial system due to data and methodological limitations. System-wide stress tests are a valuable tool in the policy maker's arsenal for assessing systemic risk.

Like microprudential stress tests, a system-wide stress test begins by subjecting institutions to a stress scenario and evaluating how their balance sheets are affected. While the microprudential stress test ends there, a system-wide stress test continues by simulating how institutions propagate the experienced financial distress to other institutions through the aforementioned contagion channels. (Note that the losses that result from this distress propagation are sometimes referred to as second round effects; e.g. in Hałaj and Kok [2013].) This iterative simulation of how institutions propagate the last round's losses into the next round may continue indefinitely until no further contagion takes place, either because the distress propagation has been damped out (see e.g. Burrows et al. [2012]) or because the entire financial system has collapsed. Alternatively, the simulation may be cut-off after a fixed number of rounds, which neglects some of the contagion materializing in higher rounds but is also less affected by the uncertainties that compound over the rounds.

System-wide stress tests have the important advantage over microprudential stress tests of capturing systemic risk (to, at least, some extent). For example, consider the microprudential stress test performed as part of the Canada Financial Sector Stability Assessment (IMF, 2014). When expanding the stress test to capture systemic risk losses using the MacroFinancial Risk Assessment Framework, capital levels of the Canadian banks fall by an additional 20% (Anand et al., 2014). Other system-wide stress tests, capturing multiple, interacting contagion channels, also show that systemic risk is significant and substantial in various financial systems.[13]

System-wide stress tests are typically used as an early-warning tool for identifying potential instabilities in the financial system and understanding the conditions and mechanisms that drive instabilities. Another application is to use system-wide stress tests as a crisis management tool. For example, they may be used during a crisis to determine the right level of capital injection to prevent a credit crunch

[13]See e.g. Burrows et al. [2012], Kok and Montagna [2016], Caccioli et al. [2015], Cont and Schaanning [2017], Dees and Henry [2017], Aikman et al. [2019b].

(Greenlaw et al., 2012) or to decide whether to bail-in, bail-out, or resolve a failing bank (Kleinnijenhuis et al., 2021). Borio et al. [2014] argue that system-wide stress tests are particularly valuable as a crisis management tool, as during a crisis the relevant stress scenario has already revealed itself. As an early warning tool, on the other hand, system-wide stress tests share the shortcoming of microprudential stress tests that their outcomes are conditional on a subjective stress scenario that is far from guaranteed to be representative of the next crisis. This is particularly problematic given the complex nature of the financial system, which may cause two scenarios that seem equally severe to yield completely different contagion dynamics. Hence, scenario-independent measures of financial stability that capture multiple interacting contagion channels should be developed to complement existing system-wide stress tests.

2.4 Contribution to the Literature

This thesis complements existing methods by proposing measures of financial stability that take into account multiple interacting contagion channels without relying on an initial stress scenario. More specifically, based on the discussed literature we identify the potential for improving the current understanding of systemic risk through methods that meet the following requirements:

1. Allow for calibration to microstructure or real financial systems (e.g. Battiston and Martinez-Jaramillo [2018], see section 2.1).

2. Capture all relevant asset types (e.g. Poledna et al. [2015], see section 2.1), all relevant types of institutions (e.g. Farmer et al. [2020], see section 2.1), and all relevant contagion channels and their interactions (e.g. Caccioli et al. [2015], see section 2.2).

3. Assess systemic risk without reliance on any specific, subjective stress scenario (e.g. Borio et al. [2014], see section 2.3).

We contribute to the literature by developing measures of financial stability that meet these requirements, and in doing so we find additional evidence for the importance measures that meet these criteria.

2.4.1 Higher-Order Exposures

In chapter 3, we introduce a measure of exposure that captures systemic risk. Exposures have the advantage over stress tests that they avoid the subjectivity inherent in relying on a stress scenario. Traditional measures of exposure focus on direct exposures to a counterparty, which arise from investments in that counterparty and are written-downs upon the default.[14] Indirect exposures arise from portfolio overlaps between institutions are increasingly recognized too (see e.g. Cont and Wagalath [2013], Caballero and Simsek [2013]). However, if exposure is understood to be the maximum loss that an institution stands to suffer when a counterparty defaults, then exposure includes (losses due to) systemic risk too. We refer to the component of exposure generate by systemic risk as "higher-order exposure" and propose a method for quantifying higher-order exposures that captures the most commonly studied contagion channels and their interactions.

We measure higher-order exposures in the South African financial system using a highly detailed data set that includes various types of institutions and assets. We contribute to the financial exposure literature by demonstrating that higher-order exposures are significant, substantial, and particularly large during financial crises (which is when exposures matter most, as defaults are more likely during crises and hence exposures more likely to materialize into losses). This highlights the importance of capturing systemic risk as part of exposure. We also contribute to the literature on risk in financial systems with heterogeneous institutions (see e.g. Farmer et al. [2020]), as we identify substantial exposures between the various sectors of the financial system. Finally, our results show that higher-order exposures cannot be extrapolated from traditional measures of exposure and are amplified by interactions between contagion channels and institutions of different types, which emphasizes the importance of the complex systems approach to measuring systemic risk.

2.4.2 Eigenvalue-Based Stability Measures

Although higher-order exposures provide an objective measure of systemic risk, they are, like stress tests, still narrow in scope as they only consider systemic risk due to the default of an institution in the system. Yet, systemic risk may emerge from various kinds of shocks so a system in which exposures are low may still generate severe systemic risks. Furthermore, the measure of higher-order exposures

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

we propose does not consider the liquidity of financial institutions, even though contagion may spread rapidly through a deterioration in banks' liquidity positions (Hałaj, 2018).

Eigenvalue-based approaches characterize the stability of a financial system by its eigenvalue spectrum and corresponding eigenvectors (Caccioli et al., 2014, Bardoscia et al., 2017, Cont and Schaanning, 2019). In particular, the tendency of the financial system to amplify or damp shocks of any nature may be characterized by the system's largest eigenvalue. Hence, eigenvalue-based approaches provide a more comprehensive measure of stability of the financial system than (higher-order) exposures or (system-wide) stress tests. Current eigenvalue-based approaches, however, do not capture interactions between contagion channels. For example, Caccioli et al. [2014] and Cont and Schaanning [2019] focus on overlapping portfolio contagion and show how it is amplified by leverage, and Bardoscia et al. [2017] show how leverage exacerbates counterparty risk contagion.

In chapter 4, we develop an eigenvalue-based stability measure that captures the interactions between multiple contagion channels, and between liquidity and solvency (whilst also meeting the other aforementioned criteria). Our results in chapter 4 contribute to the literature by elucidating how interactions between contagion channels and between institutions of different types can amplify shocks. We also provide insight into the balance between the shock-amplifying forces of leverage and various sources of damping. Although the amplification due to leverage has been widely studied (see Geanakoplos [2010], Thurner et al. [2012], Bardoscia et al. [2017] and sec. 2.2.5), the sources of damping required to offset leverage and keep the system stable have not (with the fundamental value investors in Aymanns and Farmer [2015] and Aymanns et al. [2016] being a notable exception).

We generate random financial systems of varying densities in chapter 4 and evaluate their stability. We find the sparsest systems to be less stable than their denser counterparts. This opposes the finding of Bardoscia et al. [2017], who find the sparsest systems to be the most stable. The underlying cause is that Bardoscia et al. [2017] require the sparsest systems to be acyclic, whereas we do not. These opposing findings highlight the importance of calibrating financial stability models to accurate network data of real financial systems, as findings derived from reconstructed networks carry the inherent risk of being artifacts of the specific reconstruction method chosen. Such calibration requires both comprehensive data collection efforts, as well as models that allow for calibration to the detailed, multiplex nature of real financial networks.

In chapter 5, we demonstrate that the stability measure developed in chapter 4 can be straightforwardly calibrated to real financial systems. We apply the measure to the data set on the South African financial system to provide an objective and comprehensive assessment of the system's stability. We show the stability measure may be used to inform strategies such as liquidity injections to stabilize the financial system. Policymakers may readily use the developed stability measure as an early-warning tool, in policy experiments, and potentially even to inform regulation directly, making it a valuable complement to (system-wide) stress tests.

We explore how stability is affected by liquidity shocks of various sizes in chapter 5 and show that, depending on how institutions choose to respond to liquidity shocks, the system may be "robust-yet-fragile"; that is, the system may be highly resilient to small liquidity shocks, but a single, large liquidity shock may cause the system to become highly unstable. This complements the findings of Gai and Kapadia [2010], who identify specific robust-yet-fragile network topologies, while we identify strategies for responding to liquidity shocks that yield a robust-yet-fragile financial system. The identification of robust-yet-fragile tendencies of financial systems across dimensions highlights the importance of stability measures that assess a system's resilience to a wide range of shocks, such as the eigenvalue-based approach developed here.

Chapter 3

Higher-Order Exposures

Based on the eponymous paper by Garbrand Wiersema, Alissa M. Kleinnijenhuis, Esti Kemp and Thom Wetzer. This work has been presented in the 2022 Directorate General Macroprudential Policy and Financial Stability Seminar Series at the ECB, in the 2021 Financial Stability Research Seminar Series at the South African Reserve Bank, and in the MT 2021 INET Oxford Seminar Series.

3.1 Summary

Traditional exposure measures focus on direct exposures to evaluate the losses an institution is exposed to upon the default of a counterparty. Since the Global Financial Crisis of '07/'08, indirect exposures via common asset holdings are increasingly recognized too. Yet direct and indirect exposures fail to capture the losses that result from shock propagation and amplification following the counterparty's default. In this paper, we introduce the concept of "higher-order exposures" to refer to these spill-over losses and propose a way to formalize and quantify them. Using granular data of the South African financial system and a contagion model that captures the most commonly studied contagion channels and their interactions, we show that higher-order exposures make up a significant part of exposures – particularly during times of financial distress when exposures matter most. We also show that higher-order exposures cannot simply be extrapolated from direct or indirect exposures, since they depend strongly on the network structure and the robustness of individual institutions. Our findings suggest that higher-order exposures should inform the design and calibration of those tools in the regulators' arsenal where exposures matter – including large exposure limits, capital requirement calibration, stress test design and resolution. Failure to do so may result in both lax *ex-ante* regulation and ill-informed *ex-post* handling of financial crises.

3.2 Higher-Order Exposures

An institution's exposure to a counterparty is typically understood to be the maximum loss the institution stands to suffer when the counterparty defaults. The most commonly studied exposures to a counterparty are *direct exposures*. Direct exposures arise from loans, bonds, stocks and other investments in a counterparty that are written-downs upon the default of the counterparty.[1] More recently, *indirect exposures* have garnered attention (Cont and Wagalath, 2013, Caballero and Simsek, 2013, Cont and Schaanning, 2019). Indirect exposures arise when the portfolios of tradable securities of two or more institutions' overlap. When one of these institutions decides to liquidate (part of) its portfolio, the market price of the securities falls, causing mark-to-market losses to the other institution(s).[2]

Institutions' exposures are, however, not limited to direct and indirect exposures. When institution j defaults, other institutions in the system may propagate the losses caused by the default. The propagation of losses by financial institutions is referred to as contagion.[3] Consequently, when j defaults institution i may suffer losses on other assets than its investments in and portfolio overlap with j. Put differently, these losses are not captured by i's direct and indirect exposures to j. We refer to the losses that were caused by the default of j, and propagated by at least one intermediate institution k before being suffered by i, as i's "higher-order exposure" to j. The exposure that i has to j (i.e. the total losses it suffers when j defaults) thus has a higher-order component:

$$\mathit{Exposure}_{ij} \stackrel{\text{def}}{=} \mathit{Direct\ Exposure}_{ij} + \mathit{Indirect\ Exposure}_{ij} + \mathit{Higher\text{-}Order\ Exposure}_{ij} \tag{3.1}$$

By contrast, we refer to the direct and indirect component of the exposure as the "first-order exposure". We further introduce the "higher-order share of exposure" (HSE), which expresses the share of an exposure made up by higher-order exposures. The HSE of the exposure of i to j is given by

$$\mathit{HSE}_{ij} \stackrel{\text{def}}{=} \frac{\mathit{Higher\text{-}Order\ Exposure}_{ij}}{\mathit{Exposure}_{ij}}. \tag{3.2}$$

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

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

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

Hence, the HSE expresses the fraction of an exposure that is overlooked when only considering direct and indirect exposures. When the HSE is high, conventional methods will dangerously underestimate exposures.

Stylized Example. Consider three financial institutions, i, j and k, which are part of a larger economic system. Institution i has extended a loan l_{ik} to institution k. Furthermore, i has external assets a_i and external debt d_i (i.e. investments in and debt to institutions other than i, j and k). Institution i's equity is denoted as e_i. Institution j holds a number s_j shares in stock S and has external assets a_j, external debt d_j, and equity e_j. Lastly, institution k holds a number s_k shares in stock S and has debt $d_{ki} = l_{ik}$ to i. Furthermore, k has external assets a_k, external debt d_k, and equity e_k. The balance sheets are shown in panel 3.1 and our notation is summarized in table A.1 in the Appendix. We assume for simplicity that the institutions' external assets a_q and external debt d_q, $q \in \{i, j, k\}$, do not generate exposures.

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

(a) Balance sheet institution i

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

(b) Balance sheet institution j

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

(c) Balance sheet institution k

Table 3.1: **Balance sheets of institutions i, j and k.**

Let us consider the direct exposures in this system first. The loan from institution i to k gives rise to a direct exposure E_{ik}; the value of the loan l_{ik} is written-down to its recovery value when k defaults. The propagation of losses through the write-down of contracts with a counterparty when the counterparty defaults is typically referred to as *counterparty default contagion* (Jorion and Zhang, 2009) and the corresponding direct exposure is equal to the write-down.

We assume throughout this paper that the Loss Given Default (LGD) is 100% (although this assumption is easily modified). The time taken to resolve a default is much longer than the typical timescales over which contagion materializes so short-term recovery may be realistically assumed to be zero (Elsinger et al., 2006, Cont et al., 2010). The loan l_{ik} is written-off when k defaults, so i's corresponding direct exposure to k is given by

$$E_{ik} = l_{ik}. \tag{3.3}$$

This direct exposure is visualized by the red arrow in Figure 3.1.

Let us now discuss the indirect exposures in this system. The portfolio overlap between institutions j and k generates an indirect exposure; when j defaults, its assets are liquidated as part of the default resolution, which includes the sale $\Delta s = s_j$ of j's shares in stock S. Such a fire sale typically depresses the market price P^s of stock S, which is referred to as the sale's *price impact* (Dufour and Engle [2000], Lillo et al. [2003], Potters and Bouchaud [2003]). Following Cont and Schaanning [2017], we assume a linear price-impact function of the form

$$\Delta P^s = \frac{\Delta s}{D^s}, \tag{3.4}$$

where Δs is the number of shares in stock S sold, and the market depth D^s expresses the number of shares sold per unit change in price P^s. The market depth D^s is a measure of the liquidity of stock S and approximates the sensitivity of the price P^s to changes in the supply of shares in stock S. D^s depends, among other things, on the average daily trading volume and daily volatility of the security.[4]

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

$$E_{kj} = \Delta P^s s_k = \frac{s_j s_k}{D^s}. \tag{3.5}$$

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

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

$$E_{jk} = \frac{s_j s_k}{D^s}. \tag{3.6}$$

The right-hand side of equation (3.6) is referred to as the *liquidity-weighted portfolio overlap* of j and k in stock S (Cont and Schaanning, 2017, Poledna et al., 2021). These indirect exposures are visualized by the blue arrows in Figure 3.1.
Finally, we are ready to discuss the higher-exposures in this system. Assume that k's indirect exposure to j exceeds its equity buffer, i.e.

$$E_{kj} > e_k. \tag{3.7}$$

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

$$E_{ij} = l_{ik} \tag{3.8}$$

This higher-order exposure is visualized by the purple arrow in Figure 3.1. Because i has no direct or indirect exposures to j, its higher-order share of exposure to j is given by

$$HSE_{ij} = \frac{l_{ik}}{0 + 0 + l_{ik}} = 100\%. \tag{3.9}$$

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

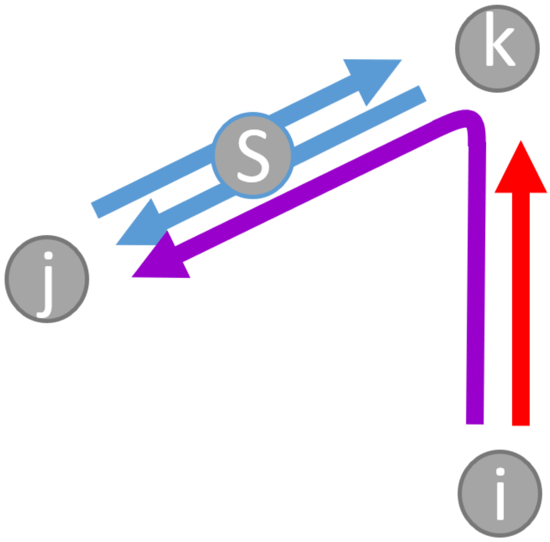

Figure 3.1: **Stylized example of a higher-order exposure.** Bank i lends to bank k. This *direct exposure* is depicted by the red arrow. Banks j and k have a large portfolio overlap through their position in stock S. When either defaults, its position in s is liquidated, which results in mark-to-market losses to the other bank, causing it to default too. This *indirect exposure* is depicted by the blue arrows. Bank i has neither a direct nor an indirect exposure to bank j. Yet, when bank j defaults, k defaults as a consequence, causing k's debt to i to be written-off. Hence, bank i has a *higher-order exposure* to bank j, which is depicted by the purple arrow. The magnitude of i's higher-order exposure to j is equal to the loss i stands to suffer when j defaults, i.e. the value of i's loan to k.

This example illustrates that a regulator needs granular data in order to understand the exposures of regulated institutions, as well as models that use these data to simulate how financial distress propagation results in higher-order exposures. Furthermore, the example also highlights the importance of capturing all relevant contagion channels and their interactions in the same model. Studying either the counterparty default contagion channel or overlapping portfolio contagion channel individually would result in overlooking the higher-order exposure of i to j completely. In the example, i's higher-order exposure to j was generated by the knock-on default of k when j defaults. As discussed below, higher-order exposures occur in the absence of knock-on defaults too, as institutions may propagate losses through various contagion mechanisms without defaulting. For example, when an institution incurs a loss, its shareholders suffer losses on the value of their shares (as the shares represent part ownership of the institution's portfolio). Furthermore, as the loss reduces the institution's solvency, making default more likely, the expected value of the institution's liabilities falls and its creditors suffer losses on the (expected value of) their assets accordingly.

Contribution. In this paper, we introduce the concept of higher-order exposures and propose a way to formalize and quantify them, using South-Africa as a case study. We calibrate a multi-layered network model to the South-African financial system, covering banks as well as various investment funds, and their interconnections through loans, deposits, various tradable securities and fund shares. We study the losses that a financial institution i is exposed to when another institution j defaults and j's tradable securities are liquidated. The first-order exposure is calculated by summing the direct exposure (i.e. i's investments in [securities issued by] j) and indirect exposure (i.e. liquidity-weighted overlap between i and j), both of which are measured in accordance with the literature. To simulate how losses propagate throughout the system following j's default, we model the most commonly studied contagion channels and their interactions. The higher-order exposures are then found by summing the additional asset losses that i incurs from the contagious spill-overs that ensue following j's default. Our results show that higher-order exposures make up the vast majority of the exposures in a significant part of the system. We also show that higher-order exposures are particularly pronounced during times of financial crises, when exposures matter most (because defaults are more likely occur and hence exposures more likely to materialize into losses). Furthermore, we find that higher-order exposures cannot simply be extrapolated from direct and/or indirect exposures, as the higher-order exposures depend strongly on the network structure and the robustness of individual institutions. The results suggest that higher-order exposures should inform the design and calibration of those tools in the regulators' arsenal where exposures matter, including large exposure limits, capital requirements, stress testing and resolution. Failure to do so may result in both lax *ex-ante* regulation and ill-informed *ex-post* handling of financial crises.

Link to the Literature. We contribute to the literature on financial exposures. A large body of work exists, dating back many years, that measures direct exposures between counterparties.[5] Since the Great Financial Crisis, the importance

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

of indirect exposures has been emphasized and measures thereof have been introduced.[6] The idea that higher-order interconnections between one institution i and another j could pose a financial risk to institution i if institution j defaults (through the process of financial contagion) is well-understood and has often been modeled.[7] Yet, these "higher-order exposures" have neither been conceptually introduced nor formally measured before. We propose that an institution's exposure to another should be understood as the sum of its direct, indirect and higher-order exposures, rather than just its direct exposures alone. Currently, indirect exposures are still frequently not taken into account in exposure measurements. Our case study of the South-African financial system shows that even when both direct and indirect exposures are considered, exposure is still significantly underestimated as higher-order exposures are often substantial, even when direct and indirect exposures are minimal or completely absent. This finding reinforces the importance of taking all components of exposure into account when calculating institutions' exposures for prudential regulatory purposes, thereby contributing to the literature on (macro)prudential regulation.[8] To the best of our knowledge, we are the first to point out that large exposure limits may be misguided if based solely on direct exposures – which is the current practice – as they are only part of the exposure. The remainder of this paper is structured as follows: We first describe the South-African financial system and discuss our data and methodology for measuring higher-order exposures. Section 3.4 presents our results. We analyze system-wide, sectoral, individual and stressed exposures in the South-African system, with particular focus on the higher-order component of these exposures. Section 3.5 discusses the implications of our findings.

3.3 Exposures in the South African Financial System

We apply the concept of higher-order exposures to the South African financial system. South Africa is a small open economy with a relatively well-developed financial market compared to other African or emerging-market economies. The South

[6] See e.g. Cont and Wagalath [2013], Caballero and Simsek [2013], Greenwood et al. [2015], Clerc et al. [2016], Cont and Schaanning [2017], Calimani et al. [2017], Aymanns et al. [2018], Cont and Schaanning [2019], Aldasoro et al. [2020].

[7] See e.g. Allen and Gale [2000], Gai and Kapadia [2010], Kok and Montagna [2016], Elliott et al. [2014], Glasserman and Young [2015], Wiersema et al. [2019], Farmer et al. [2020].

[8] See e.g. Persaud [2009], Freixas et al. [2015], Aikman et al. [2019a], Jeanne and Korinek [2020].

African debt market is liquid and well-developed in terms of the number of participants and their daily activity, and its equity market dominates the region in terms of capitalisation (Andrianaivo and Yartey, 2010). The local currency is South African Rand, which will henceforth be referred to as ZAR.
Banking sector assets exceed GDP in aggregate terms, but are smaller than the assets held by the non-bank financial intermediation sector, which includes entities such as insurers, pension funds and collective investment schemes, which are henceforth referred to as "funds". Since the Global Financial Crisis, the share of assets held by banks has decreased, as the growth of assets held by the non-bank financial sector – in particular funds - has outpaced that of banks (Kemp, 2017).
Non-bank financial intermediaries are an important source of funding for banks and direct linkages among banks and non-bank financial intermediaries other than pension funds and insurers are relatively high, amounting to 15% of bank assets (FSB, 2018).

3.3.1 Institutions

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

3.3.1.1 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 3.2a. Overall, the banking sector is largely funded by deposits, but banks also issue debt instruments, such as bonds and money market instruments (MMIs), and equity shares.
We calculate higher-order exposures of banks and funds to each of the six largest banks, as they form the core of the South African financial system and generate the largest exposures. The six largest banks (by total assets) are the Standard Bank of South Africa Ltd (Standard Bank), FirstRand Bank Ltd (Firstrand), Absa Bank

Ltd (Absa), Nedbank Ltd (Nedbank), Investec Bank Ltd (Investec) and Capitec Bank Ltd (Capitec). While the assets held by Capitec are significantly smaller than the assets held by the other five large banks, it is included in our analysis given that it is the second largest retail bank based on the number of customers.

3.3.1.2 Funds

Funds pool investors' money and purchase a portfolio of securities, thereby offering investors the opportunity to obtain exposure to a diverse portfolio of underlying securities, without having to purchase and trade securities directly. From the investor's perspective, funds provide investors with an opportunity to earn higher returns than those offered by deposits, in return for taking on greater risk. There are over 1200 open-ended funds registered South African funds with assets under management of about 2 trillion ZAR in 2016. We divide funds into three categories according to the instruments they invest in: money market funds (MMFs), fund of funds (FoFs) and other funds (OFs). Participation bond schemes and hedge funds are not within the scope of this study given data limitations and their relatively small size (see Kemp [2017]).

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

In the case of a default of, e.g., one of the five largest banks (which borrow substantially from funds), funds could act as shock absorbers by spreading losses across a diverse set of investors. In extreme circumstances, however, open-ended funds involved with credit intermediation together with leverage, liquidity or maturity transformation could be susceptible to "runs" – i.e. large-scale redemption requests, when investors anticipate or observe a substantial drop in their fund shares' value. When a run is initiated, the fund may run out of liquid assets and become unable to meet redemptions. As a result, the fund may have to resort to fire-selling assets and suspending redemptions, and could eventually be wound up.

Money Market Funds

MMFs are formally designated according to legal requirements (Board, 2014): Board Notice 90 of 2014 (Board, 2014) restricts the money-market instruments that a fund manager may invest in, both in terms of maturity of the investments and in terms of the maximum counterparty exposure (inclusion limits).
For the restriction on maturity transformation, the weighted average legal maturity of the fund may not exceed 120 days, while the weighted average duration of the money-market instruments may not exceed 90 days. No single instrument that MMFs invest in may have a maturity exceeding 13 months. The regulations also limit the exposure in terms of the maximum percentage of the aggregate market value of the portfolio. These limitations include a maximum exposure of 30% of total fund assets to (MMF instruments issued by) local or foreign banks (registered in South Africa) of which the holding company is listed on the exchange if the market capitalisation of the listed group holding company exceeds R20 billion, and to 20% if the market capitalisation of the listed group holding company is between R2 billion and R20 billion. MMFs can also invest in money market instruments issued by any local or foreign entity that is listed on an exchange. This exposure is limited to 10% per issuer.
The MMF industry in South Africa is relatively small, amounting to 2% of total financial assets in Q42016. The sector is concentrated - of the 49 MMFs in South Africa, 82% of assets are held by the 10 largest MMFs as illustrated in Figure 3.2b. Money-market instruments issued by banks and deposits placed with banks make up 90% of the overall portfolio of MMFs, with the remainder of the holdings made up of instruments issued by non-banks (SARB, 2017b). Amid these regulatory restrictions on instruments that MMFs are allowed to invest in, MMFs in South Africa have large exposures to banks making the portfolios of MMFs exceptionally similar (Kemp, 2017).

Fund of Funds

Although there is no formal distinction among non-MMF funds, a specific subset of these funds can be singled-out because of their particular investment strategy (Kemp, 2017). These "Fund-of-Funds" (FoFs) invest predominantly in other funds' shares and give investors exposure to multiple fund schemes and, consequently, diversification across management styles. Due to their investments in other funds, FoFs are highly exposed to instabilities in the fund sector. This potentially generates significant higher-order exposures, which is why this particular subset of funds is singled-out. We classify funds that invest more than 80% of their portfolio in the shares of other funds as FoFs. Our data include over 400 FoFs, and the distrubtion of their total asset sizes is show in Figure 3.2c.

Other Funds

The remaining (i.e. non-MMF, non-FoF) "Other Funds" (OFs) include equity funds (which predominantly invest in equity shares), non-MMF fixed income funds (which typically invest in debt instruments with longer maturities), multi-asset funds (whose investments include both equity shares and debt instruments), and real-estate investment trusts. As these OFs invest in a mixture of instruments issued both domestically and off-shore, their portfolios are highly variable. Our data include over 800 OFs and the distribution of their total asset sizes is show in Figure 3.2d.

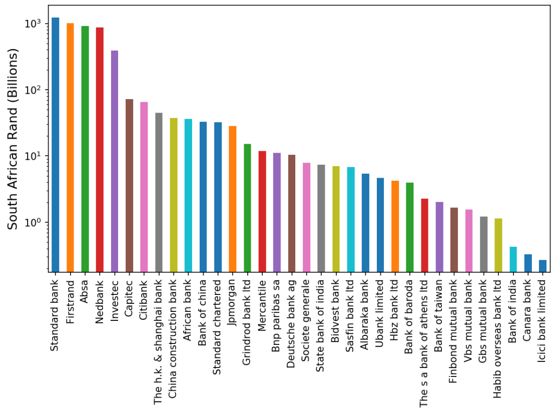

(a) Banks ranked by total assets

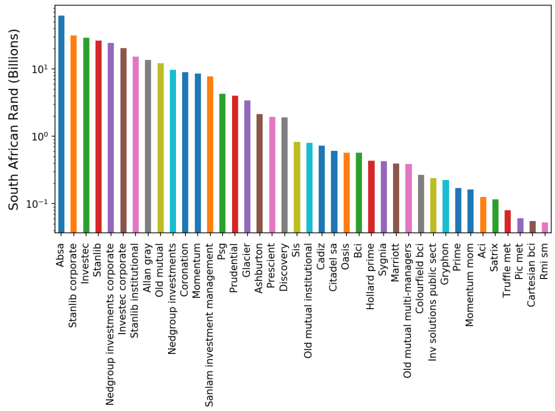

(b) MMFs ranked by total assets

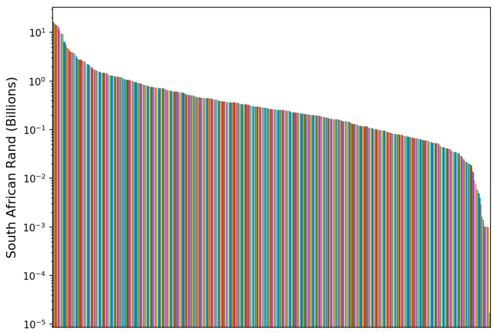

(c) FoFs ranked by total assets

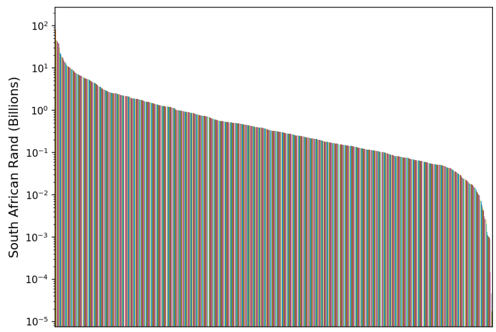

(d) OFs ranked by total assets

Figure 3.2: **Distribution of South African financial institutions by asset size.** The institutions are listed on the x-axis in decreasing order of total assets size and their total assets in billions of South African Rand are on the y-axis (log-scale). Note that the FoFs' and OFs' names are not listed because they are too numerous. The banking sector consist of a core of six large banks – the Standard Bank of South Africa Ltd (Standard Bank), FirstRand Bank Ltd (Firstrand), Absa Bank Ltd (Absa), Nedbank Ltd (Nedbank), Investec Bank Ltd (Investec) and Capitec Bank Ltd (Capitec) – and a periphery of 28 smaller banks. The FoF and OF sectors also show a strong concentration in terms of asset size, whereas the concentration in the MMF sector is less pronounced.

3.3.2 Network Data

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 (SARB, 2016b). Balance sheet entries are aggregated by asset type and counterparty type (e.g. "loans and deposits to domestic banks"). Data on funds' assets were sourced from Morningstar Inc and are highly granular. These data report funds' investments per instrument type in individual counterparties (e.g. "bonds in Absa"). The data do not explicitly report funds' shareholders. However, we observe that banks do not invest funds based on the banks' balance sheet data, and from the funds' asset data we know funds' holdings of shares in other funds. Hence, all of a fund's shareholders that are included in the model are given by the data at the level of individual counterparties. The remainder is assumed to be held by external parties.

3.3.2.1 Balance Sheet Composition

We focus on five types of domestic assets (foreign assets are not modeled): loans and deposits (l), bonds (b), money market instruments, or *MMIs* (m), equity shares (e) and fund shares (f). Figure 3.2 shows where these assets appear on the stylised balance sheet of a bank, MMF and FoF/OF. The data also include investments in gold, repo, foreign institutions and the real economy. We do not model these assets, however, as they are not expected to cause substantial contagion through the channels included in this paper.[9]

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

<table>
<tr><th colspan="2">Assets</th><th>Liabilities</th></tr>
<tr><td colspan="2">Loans + Deposits (l)</td><td>Loans + Deposits (l)</td></tr>
<tr><td rowspan="3">Tradable securities</td><td>Bonds (b)</td><td>Bonds (b)</td></tr>
<tr><td>MMIs (m)</td><td>MMIs (m)</td></tr>
<tr><td>Equity Shares (e)</td><td>Other liabilities</td></tr>
<tr><td colspan="2">Other assets</td><td>Equity (e)</td></tr>
</table>

(a) Stylised balance sheet of a bank.

<table>
<tr><th colspan="2">Assets</th><th>Liabilities</th></tr>
<tr><td colspan="2">Deposits (l)</td><td rowspan="4">Fund shares (f)</td></tr>
<tr><td>Tradable securities</td><td>MMIs (m)</td></tr>
<tr><td colspan="2">Fund shares (f)</td></tr>
<tr><td colspan="2">Other assets</td></tr>
</table>

(b) Stylised balance sheet of a MMF.

<table>
<tr><th colspan="2">Assets</th><th>Liabilities</th></tr>
<tr><td colspan="2">Deposits (l)</td><td rowspan="6">Fund shares (f)</td></tr>
<tr><td rowspan="3">Tradable securities</td><td>Bonds (b)</td></tr>
<tr><td>MMIs (m)</td></tr>
<tr><td>Equity shares (e)</td></tr>
<tr><td colspan="2">Fund shares (f)</td></tr>
<tr><td colspan="2">Other assets</td></tr>
</table>

(c) Stylised balance sheet of FoFs and OFs.

Table 3.2: **Stylised balance sheets of the types of modeled financial institutions in the South-African financial system**. We consider: (a) banks; (b) MMFs; and (c) FoFs and OFs. FoFs have the same balance sheet structure as OFs, but FoFs invest more than 80% of their assets in fund shares.

Loans & Deposits (l)

Only banks receive loans and deposits, because funds do not have debt. The bank data do not distinguish between deposits and loans (of any maturity), so they are all treated as one and the same and we only distinguish between deposits/loans to different counterparties.

Bonds (b)

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

MMIs (m)

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

Equity shares (e)

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

Fund Shares (f)

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

3.3.2.2 Initialization Values

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

3.3.3 Network Construction

The data can be represented as a weighted, directed network, with the nodes representing the institutions and the edges their assets (a node's out-edges are given by its assets and its in-edges by its liabilities and/or issued shares). We refer to this as the "asset network". The network is multiplex, consisting of five layers, which each layer representing one of the asset types $\alpha \in \{l, b, m, e, f\}$. Each layer includes the same set of nodes, made up by the (individual) banks and funds, the node $\mathcal{G}$ representing the government, and a single "corporate" node $\mathcal{C}$ representing the domestic

non-financial corporate sector.[10] As noted before, we do not model the government and corporate sector.[11] We use $\mathcal{B}$ to denote the set of banks, $\mathcal{F}$ the set of funds, $\mathcal{I} = \mathcal{B} \cup \mathcal{F}$ denotes the financial institutions, $\mathcal{H} = \mathcal{B} \cup \mathcal{G} \cup \mathcal{C}$ the set of securities-issuing institutions, and $\mathcal{A} = \mathcal{B} \cup \mathcal{F} \cup \mathcal{G} \cup \mathcal{C}$ the set of all institutions.
Let w^{α}_{ik} denote the weight of an edge pointing from node $i \in \mathcal{I}$ to node $k \in \mathcal{A}$ in the layer of the asset network corresponding to investments of type $\alpha \in \{l, b, m, e, f\}$. For loans and deposits (l), we set the weight w^{l}_{ik} equal to the sum of the principals of i's loans to and deposits at k. For securities $\sigma \in \{b, m, e, f\}$, we set the weight w^{σ}_{ik} equal to the *initial* value of i's position in securities of type σ issued by k. (Note that the values of securities may change when we simulate the losses following a default, by the weights w^{α}_{ik} do not.)
We only model the banks' and funds' investments of type $\alpha \in \{l, b, m, e, f\}$ in domestic institutions $k \in \mathcal{A}$. Therefore, each of a bank's or fund's modelled assets can be represented by a weighted, directed edge in the corresponding layer of the asset network. The directed edges corresponding to funds' assets are given directly by our detailed data on the funds' investments in individual banks and funds. As each bank's balance sheet aggregates the banks' assets and liabilities by instrument and counterparty type, the directed edges corresponding to banks' investments in individual counterparties are randomly generated based on the aggregate data per asset type. We explain the algorithm to reconstruct the interbank network next.

3.3.3.1 Interbank Network Reconstruction

The technique used for the reconstruction of the banks' investments is is similar to Kok and Montagna [2016] and aims to reproduce the high heterogeneity of interconnections observed in financial networks. Based on the banks' balance sheet data, we observe that banks do not invest in funds, and we know each bank's investments, per asset type, in the government $\mathcal{G}$ and corporate sector $\mathcal{C}$. Therefore, only the interbank investments require reconstruction. We assume that the (initial) market value of any security that a bank has issued is equal to the book value of that liability or equity share on the banks' balance sheets, and perform the following steps for each of the asset types $\beta \in \{l, b, m, e\}$ in which banks invest:

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

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

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.

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

5. Steps 2-4 are repeated until all banks' assets of type β are allocated.[12]

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

3.3.3.2 Direct Exposure Network

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

$$\hat{w}_{ik}^{\delta} = \sum_{\alpha \in \{l,b,m,e,f\}} w_{ik}^{\alpha}, \tag{3.10}$$

where $\hat{w}_{ik}^{\delta}$ gives the weight of the edge from i to k in the direct exposure network denoted by the superscript δ. Note that $i, k \in \mathcal{I}$, so the banks and funds are included in the direct exposures network, but the government and corporate nodes are not. The resulting direct exposure network is visualized in Figure 3.3.

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

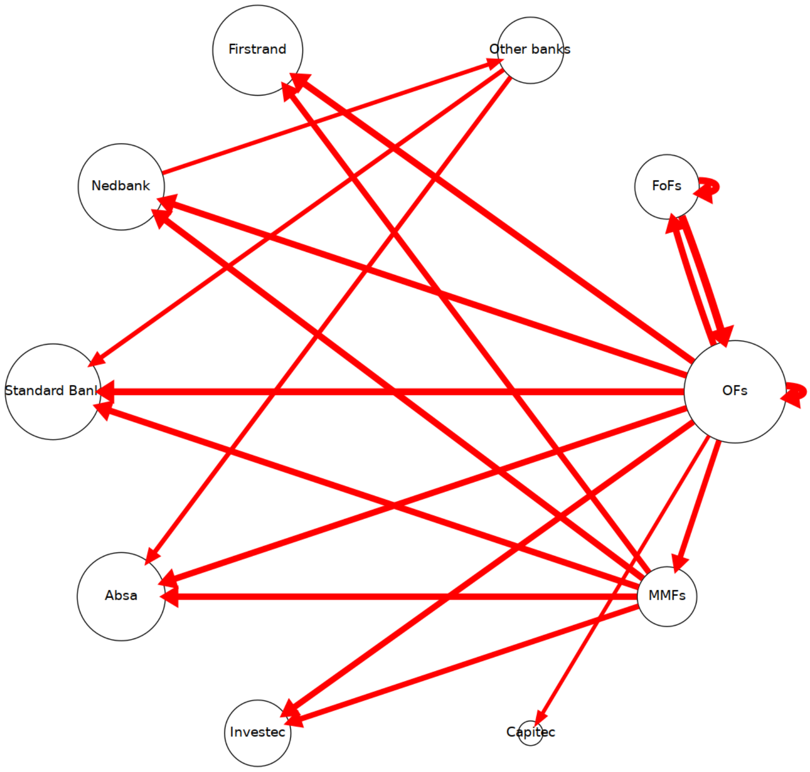

Figure 3.3: **Network of the largest direct exposures in the South African financial system**. We plot the largest direct exposures between the six largest banks, the rest of the banking sector, and the three fund sectors. Edge widths visualize the size of the exposure, varying between 5 bln. to 196 bln. South African Rand, and node sizes visualize the total asset size of the node, ranging from 717 bln. to 1.6 trn. South African Rand. The scaling is logarithmic for both edge widths and node sizes. Other than the six largest banks, all institutions of the same type are aggregated into a single node. Edges point in the direction of the exposures. For example, the edge from the MMFs to Absa denotes the sum of all MMFs' direct exposures to Absa. The OF and FoF sectors' self-loops denote the sums of all direct exposures between OFs or FoFs. The figure shows that the largest exposures are those of the MMF and OF sectors to the six large banks, and the OF and FoF sectors' exposures to itself.

3.3.3.3 Indirect Exposure Network

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

$$\hat{w}^{\phi}_{ik} = \sum_{\tau \in \{b,m,e\}} \sum_{q \in \mathcal{H}} \frac{w^{\tau}_{iq} w^{\tau}_{kq}}{D^{\tau}_{q}}, \tag{3.11}$$

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

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

$$D^{\tau}_{q} \stackrel{\text{def}}{=} \sum_{k \in \mathcal{I}} \frac{w^{\tau}_{kq}}{\mu} \tag{3.12}$$

as a proxy of the market depth throughout this paper, in line with Wiersema et al. [2019]. The market depth divisor μ is a nondimensional constant of order one which, due to data limitations, is assumed to be the same for all tradable securities. As our baseline value choose $\mu = .21$, because this produces a market depth for South African government bonds equal to the bonds' market capitalization at the end of 2016.[13]. We explore various values of the market depth modifier μ in the results section to understand the sensitivity of our findings to the market depth. Like the direct exposure network, the indirect exposure network only includes banks and funds. Figure 3.4 visualizes the indirect exposure network.

Comparing equations (3.10) and (3.11) reveals that the same tradable securities that generate indirect exposures are also included in the direct exposure network.

[13]The market capitalization of South African bonds as of Q4 2016 is sourced from the Q1 2017 SARB Quarterly Bulletin; https://www.resbank.co.za/content/dam/sarb/publications/quarterly-bulletins/quarterly-bulletin-publications/2017/7718/07Statistical-tables—Public-Finance.pdf. The total value $\sum_{k \in \mathcal{I}} w^{\tau}_{kr}$ of government bonds held held by institutions $k \in \mathcal{I}$ included in our data is equal to 21% of the market capitalization of South African government bonds'. A value of $\mu = .21$ implies that the price of South African government bonds falls to zero when all holders of the bonds (including parties not included in our model) sell their entire position in the bonds, as is assumed in Wiersema et al. [2019].

The literature that studies contagion via tradable assets typically models the indirect linkages only, while the direct linkages are ignored (Farmer et al., 2020). We improve on this approach by recognizing that the same tradable asset can generate both direct and indirect exposures.

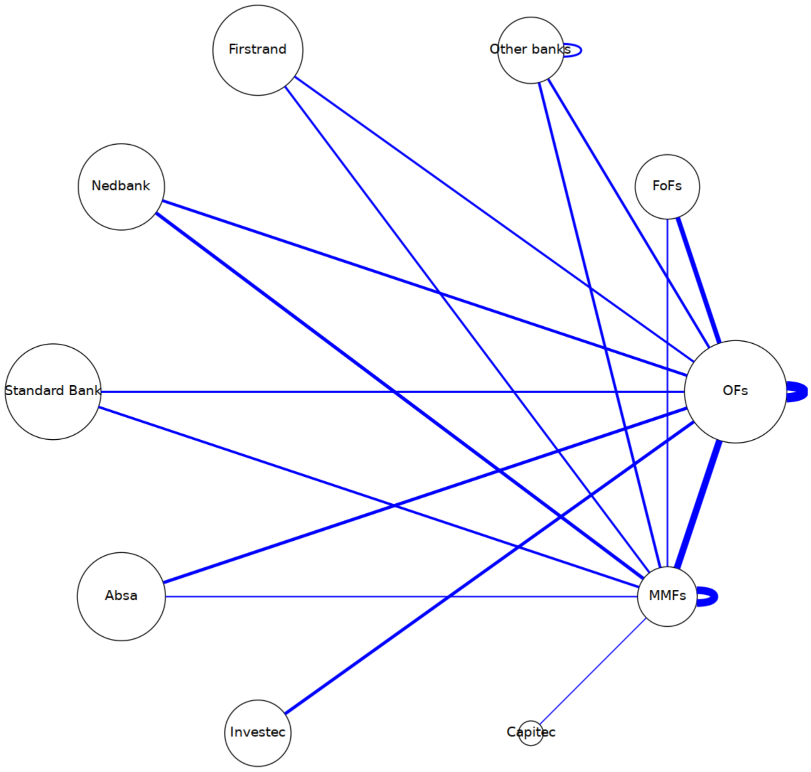

Figure 3.4: **Network of the largest indirect exposures in the South African financial system**. We plot the largest indirect exposures (i.e. institutions' liquidity-weighted portfolio overlaps) between the six largest banks, the rest of the banking sector, and the tree fund sectors. Edge widths visualize the size of the exposure, varying between 2 bln. and 187 bln. South African Rand, and node sizes visualize the total asset size of the node, ranging from 717 bln. to 1.6 trn. South African Rand. The scaling is logarithmic for both edge widths and node sizes. Other than the six largest banks, all institutions of the same type are aggregated into a single node. Indirect exposures are symmetric, so they are drawn as an undirected network. The MMF and OF sectors' self-loops denote the sums of all indirect exposures between MMFs or OFs. These self-loops make up the system's largest indirect exposures. Note that the MMF sector is much smaller than the OF sector in total asset size. Hence, the large indirect exposures between MMFs highlight their exceptionally similar portfolios.

3.3.3.4 First-Order Exposure Network

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

$$\hat{w}_{ik} = \hat{w}_{ik}^{\delta} + \hat{w}_{ik}^{\phi}, \tag{3.13}$$

where $\hat{w}_{ik}$ gives the weight of the edge from i to k in the first-order exposure network. The first-order exposure network is visualized in Figure 3.5, where red edges indicate exposures that are predominantly made up of direct exposures, blue edges exposure that are predominantly indirect, and shades of purple, a mix between red and blue, indicate exposures with both a substantial direct and indirect component (the more red the shade is, the larger is the direct component and the more blue, the larger the indirect component). Note that the exposures to the six largest banks are all predominantly direct exposures, because these banks have relatively small liquidity-weighted portfolio overlaps with other institutions in the South African financial system.

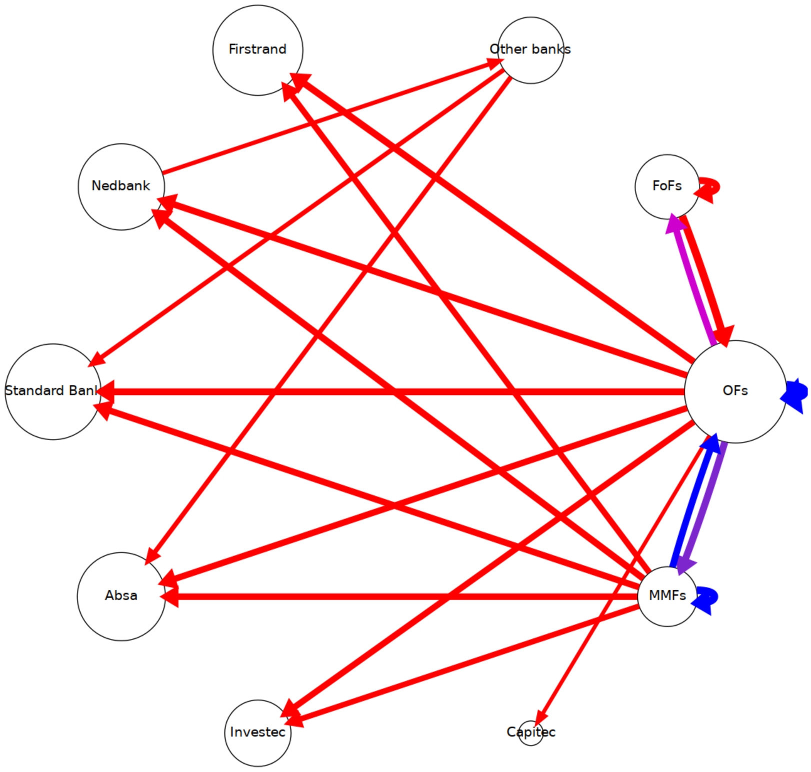

Figure 3.5: **Network of the largest first-order exposures in the South African financial system.** We plot the largest first-order exposures (i.e. direct + indirect exposures) between the six largest banks, the rest of the banking sector, and the tree fund sectors. Edge widths visualize the size of the exposure, varying between 5.7 bln. and 383 bln. South African Rand, and node sizes visualize the total asset size of the node, ranging from 717 bln. to 1.6 trn. South African Rand. The scaling is logarithmic for both edge widths and node sizes. Other than the six largest banks, all institutions of the same type are aggregated into a single node. Edges point in the direction of the exposures. For example, the edge from the MMFs to Absa denotes the sum of all MMFs' first-order exposures to Absa. The color of an arrow indicates the composition of that first-order exposure: A red edge indicates an exposure that is predominantly direct and a blue edge indicates an exposure that is predominantly indirect. Shades of purple, a mix between red and blue, indicate exposures with both a substantial direct and indirect component (the more red the shade is, the larger is the direct component and the more blue, the larger the indirect component). The figure shows that all large exposures to banks are direct and that OFs and MMFs have very large indirect exposures between them.

In Figure 3.6, we summarize our network construction by illustrating the correspondence between the layers of the asset network in the left column and the exposure networks in the right column. As each of the three tradable securities layers $\tau \in \{b, m, e\}$ (as well as the corresponding exposures) can be visualized identically, we only plot one tradable securities layer for simplicity. The figure shows on the one hand that (exposures corresponding to) multiple asset types can appear in the same exposure network. On the other hand, and more importantly, the figure visualizes that the same asset can generate exposures in both the direct and indirect network. Such assets appear as both direct and indirect exposures in the first order exposure network.

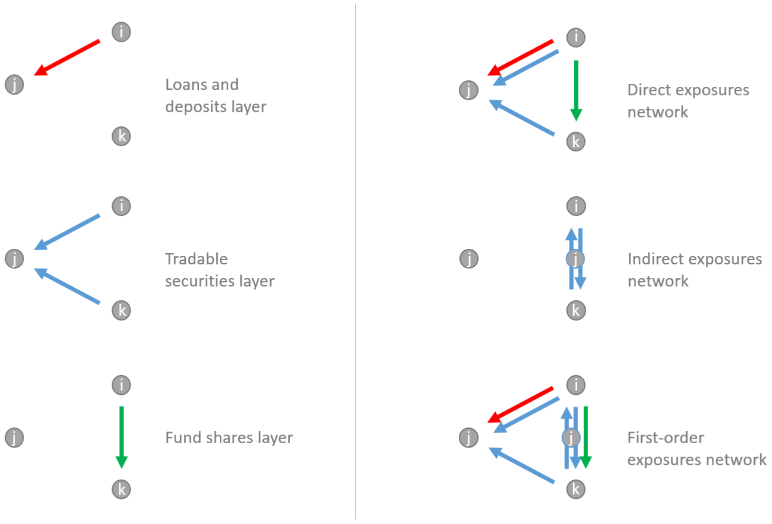

Figure 3.6: **Stylized correspondence between the asset network and exposure networks.** The figure summarizes our explanation of the network construction by illustrating the correspondence between the exposure networks and layers of the asset network. The left column shows the layers of the asset network and the right column the corresponding exposure networks. Because each of the three tradable securities layers $\tau \in \{b, m, e\}$ (as well as the corresponding exposures) can be visualized identically, we only plot one of the tradable securities layers for simplicity. Arrows in the left column point in the direction of the (principal) cashflow: FoF i deposits at bank j and buys fund shares in MMF k, so i suffers a write-off when either j or k defaults. The deposits are visualized by the red arrow and the fund shares by the green arrow. Both i and k buy MMIs in j, so either suffers a mark-to-market loss when the other liquidates their MMIs, and both suffer a write-off on the MMIs when j defaults. The MMIs are visualized by the blue arrows. Arrows in the right column point in the direction of the exposure and their color reflects the asset that generates the exposure. For example, an arrow from i to k represents an exposure of i to k and the color of the arrow is equal to that of the asset that generates the exposure. Exposures in the direct network are made-up of the write-offs that an institution would suffer on investments in a counterparty when that counterparty defaults. Exposures in the indirect network are made-up of institutions' liquidity-weighted portfolio overlaps, i.e. the mark-to-market losses suffered by an institution that has a portfolio overlap with another institution, when that other institution liquidates its portfolio. The first-order network is the sum of the direct and indirect networks. Note that the same MMIs generate both direct and an indirect exposures, causing exposures to multiple counterparties for both i and k in the first-order exposures network.

3.3.4 Measuring Higher-Order Exposures

Having set-up the asset and exposure networks, we are ready to compute the higher-order exposures. In the stylised example in Section 3.2, we calculated the exposure to an institution by recording all losses that would follow from the default of that institution. We did so by simulating how losses would propagate upon the default of the institution, where the propagation was dictated by the contagion mechanisms we formulated. Here, we develop a more comprehensive model of contagion in order to quantify the higher-order exposures more realistically and take into account the various types of assets included in our model. Yet, the approach is conceptually identical to the example and in fact the same approach may straightforwardly be generalized to calculate exposures in any financial system:

1. We formulate a model of how contagion channels propagate losses across the assets $\alpha \in \{l, b, m, e, f\}$ in the network.
2. We reconstruct the interbank investments from the banks' aggregate balance sheet data and initialize the asset network, based on the fund data and reconstructed interbank investments.
3. We use the contagion model to simulate how losses propagate throughout the network upon the (idiosyncratic) default of one of the six large banks. All losses incurred over the course of the simulation are recorded as exposures to this bank.
4. We repeat steps 2-3 1000 times for each of the six largest banks to average institutions' exposures to these banks over the random realizations of the reconstructed interbank network.

Hence, each simulation starts with assuming the (idiosyncratic) default of one of the six large banks, to calculate the exposures to that bank. We use j to denote the bank that defaults at the start of the simulation. Furthermore, we assume discrete time dynamics and use n to denote the n^{th} round following j's default. (Note that any parameter that lacks the subscript denoting the n^{th} round after j's default is a constant.) The losses $L_{ij,1}$, $i \in \mathcal{I}$, in the first round following j's default constitute the first-order exposures $E_{ij,1}$ and are given by the first-order exposure network (eq. 3.13):

$$E_{ij,1} = L_{ij,1} = \hat{w}_{ij} \tag{3.14}$$

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

$$E_{ij,n} = L_{ij,n}, \tag{3.15}$$

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

$$L_{ij,n} = \sum_{\alpha \in \{l,b,m,e,f\}} L^{\alpha}_{ij,n}. \tag{3.16}$$

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

$$HSE_{ij,n} \stackrel{\text{def}}{=} \frac{\sum_{r=2}^{n} E_{ij,r}}{\sum_{r=1}^{n} E_{ij,r}}. \tag{3.17}$$

Table 3.3 summarizes our terminology.

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

Table 3.3: **Terminology.**

3.3.4.1 Contagion Channels

The contagion channels introduced below determine the losses $L^{\alpha}_{ij,n+1}$ per investment type α in the $(n+1)^{th}$ round from the losses $L_{ij,n}$ in the n^{th} round. Once we have computed these, we can determine the higher-order exposures using equation (3.16). Assets are generally affected by multiple, interacting contagion channels. We first explain each contagion channel individually, before we discuss how the combination of interacting contagion channels determines the losses per investment type $L^{\alpha}_{ij,n+1}$.

Established methods for evaluating exposures assume that institutions' portfolios are fixed in the run-up to a default; direct exposures are calculated under the assumption that institutions do no liquidate their investments in a counterparty in the run-up to that counterparty's default and indirect exposures are calculated assuming fixed portfolio overlaps. We assume, in-line with these methods, that an institution does not liquidate any of its assets before it defaults, and that only upon

its default are all of its tradable securities liquidated as part of the resolution process. (Assuming liquidation of tradable securities upon default is in line with contagion literature such as Burrows et al. [2012] and Caccioli et al. [2015], and practice.) This assumption restricts our model to the following three contagion channels: Counterparty (default) risk contagion, overlapping portfolio contagion, and what we refer to as "shareholder contagion". (Shareholder contagion simply distributes an institution's losses across its shareholders, as explained below.)
In reality, the strategic decisions and/or behavioral actions institutions' managers take during times of stress can lead to additional channels of contagion. In particular, institutions may start liquidating assets which are perceived as risky, leading to liquidity spirals (Hoerova et al., 2009, Acharya and Skeie, 2011, Gai et al., 2011). Institutions' responses to stress are inherently uncertain, however, so by assuming no liquidation of assets prior to default we avoid this source of uncertainty. Specifically, as institutions are assumed not to withdraw their deposits from banks, banks cannot default through illiquidity in our model.

Shareholder Contagion

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

Overlapping Portfolio Contagion

As explained, the liquidity-weighted portfolio overlap between two institutions gives the mark-to-market loss that either suffers when the other defaults. Overlapping portfolio contagion refers to these mark-to-market losses that institutions suffer on their tradable securities due to the price impacts of other institutions liquidating the securities upon default. Hence, the overlapping portfolio contagion that an institution suffers in round $n + 1$ is driven by the price impacts due to the liquidation of the tradable securities of the institutions that default in round n.

Counterparty Risk Contagion

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

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

3.3.4.2 Interacting Contagion Equations

Before we introduce the contagion equations, we first show how institutions' assets evolve over the rounds, since the asset losses in round n determine the contagion in round $n+1$. The initial total assets $A_{kj,1}$ of a bank or fund $k \in \mathcal{I}$ are given by the data, and k's total assets in the $(n+1)^{th}$ round after j's default are found by subtracting k's losses in previous rounds:

$$A_{kj,n+1} = A_{kj,n} - L_{kj,n}. \tag{3.18}$$

The contagion equations introduced below guarantee that $L_{kj,n} \leq A_{kj,n}$ so $A_{kj,n+1}$ cannot become negative.

Note that the NAV $N_{kj,n}$ of a fund is proportional to its total assets, so

$$N_{kj,n} = \frac{A_{kj,n}}{A_{kj,1}} N_{kj,1} = \frac{A_{kj,n}}{A_{kj,1}} \text{ZAR}, \tag{3.19}$$

because we have normalized the initial NAV to 1 ZAR, and the NAV evolves according to

$$\Delta N_{kj,n} = N_{kj,n} - N_{kj,n+1}. \tag{3.20}$$

Let us also introduce institutions' *buffers* as these affect all asset values. We refer to the amount of losses that institution k can absorb before it defaults as k's buffer $B_{kj,n}$.[14] A bank is assumed to default when its equity is reduced to zero, because of insolvency. Funds cannot become insolvent as they do not have debt, but are vulnerable to runs that could make them default through illiquidity. We assume that a run on a fund is initiated (causing the fund to default) when the fund loses 45% of its total asset value. This threshold is in line with Cont and Wagalath [2013] and we explore how this assumption affects exposures in the results section. Hence, if k is a bank, its initial buffer is equal to its equity $e_{kj,1}$ (as given by the data), and when k is a fund, its initial buffer is equal to 45% of its total asset A_k:

$$B_{kj,1} = \begin{cases} e_{kj,1}, & \text{if } \ k \in \mathcal{B} \\ 0.45 A_k, & \text{if } \ k \in \mathcal{F} \end{cases}. \tag{3.21}$$

Banks' and funds' buffers are updated according to

$$B_{kj,n+1} = \max\left\{B_{kj,n} - L_{kj,n}, 0\right\}, \tag{3.22}$$

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

$$e_{kj,n} = B_{kj,n}, \tag{3.23}$$

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

$$B_{gj,n} = e_{gj,n} = 1, \quad g \in \mathcal{G} \cup \mathcal{C}. \tag{3.24}$$

[14]Note that the model continues to record losses on an institution's assets after it defaults. This is for two reasons. First, when (in reality) a defaulted institution is resolved, the residual value of its assets determines the recoveries, so the depreciation of a defaulted institution's assets has economic relevance. Second, we vary institutions' buffers to study how this affects stability. When we do not record an institution's losses after its default, the size of the institution's buffer maximizes the amount of losses that an institution can incur. Hence, decreasing institutions buffers would artificially limit the amount of losses that institutions can incur and therefore the destabilizing impact of decreasing buffers would not be properly reflected in the losses. Figure A.4 in the appendix shows that the assumption that institutions continue to incur losses after their default does not significantly affect our main results.

Finally, we use $\mathcal{D}_{j,n}$ to denote the set of institutions who exhaust their buffer during round n (causing default), i.e.

$$\mathcal{D}_{j,n} = \{k | B_{kj,n} = 0 \cap B_{kj,n-1} > 0\} \tag{3.25}$$

Note that as an institution can only default once, it can only be included in $\mathcal{D}_{j,n}$ for one value of n.
We are now ready to discuss how to calculate the terms $L^{\alpha}_{ij,n}$ in equation (3.16) for each investment type $\alpha \in \{l, b, m, e, f\}$. We first introduce the value of an investment in round n and then discuss how the interacting contagion channels drive the losses in the investments' values.

Loans and Deposits

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

$$v^{l}_{ikj,n} = \epsilon^{l}_{kj,n} w^{l}_{ik}. \tag{3.26}$$

The risk adjustment factor $\epsilon^{\rho}_{kj,n}$ reflects the expected loss on the investment and may be based on sophisticated models that estimate losses using k's buffer, profitability, collaterals, and other institution-specific and market-level features (as is common in expected loss modeling for capital and accounting purposes). We derive a simple estimate of the risk-adjusment factor $\epsilon^{\rho}_{kj,n}$ based only k's buffer. Although our framework allows for calcuating higher-order exposures using more advanced models of the risk-adjustment factor, the estimate of the risk-adjusment factor derived here is sufficient for our present purposes.
Assume for simplicity that the initial risk adjustment factor $\epsilon^{\rho}_{kj,1} = 1$ for any investment $\rho \in \{l, b, m, f\}$, i.e. that the initial risk-adjusted value of any investment is equal to the investment's nominal value given by the data. Hence, we assume that all counterparties initially have zero default risk. This is only for simplicity as it does not affect our results. As explained previously, we assume that short-term recoveries are zero, which implies that the risk-adjusted value of an investment is zero once the counterparty's buffer is exhausted (causing default). This, together with the assumption that the risk-adjusted value of an investment is proportional to the

counterparty's buffer (as in Bardoscia et al. [2017] and Wiersema et al. [2019]), implies that the risk-adjustment factor in round n for any investment $\rho \in \{l, b, m, f\}$ in k is equal to

$$\epsilon^{\rho}_{kj,n} \stackrel{\text{def}}{=} \frac{B_{kj,n}}{B_k}. \tag{3.27}$$

We will also use

$$\Delta\epsilon^{\rho}_{kj,n} = \epsilon^{\rho}_{kj,n} - \epsilon^{\rho}_{kj,n+1}. \tag{3.28}$$

The loss that i suffers on its portfolio of loans and deposits in round $n+1$ after j's default is given by the decrease in the loans' and deposits' risk-adjusted value due to counterparty risk contagion:

$$L^{l}_{ij,n+1} = \sum_{k\in\mathcal{B}} v^{l}_{ikj,n} - v^{l}_{ikj,n+1} = \sum_{k\in\mathcal{B}} \Delta\epsilon^{\rho}_{kj,n} w^{l}_{ik}, \tag{3.29}$$

where the sum runs over banks $k \in \mathcal{B}$ only because funds do not take loans or deposits and (the solvencies of) the government and non-financial corporates are not modeled.

Fund Shares

When a fund suffers a loss, the NAV of the shares it has issued falls through shareholder contagion. Still, this does not cover the fund's increased default risk, so the (expected) value of the shares in our model is given by their risk-adjusted NAV. The value $v^{f}_{ikj,n}$ of institution i's shares in fund k in round n is therefore given by

$$v^{f}_{ikj,n} = \epsilon^{f}_{kj,n} \frac{N_{kj,n}}{N_{kj,1}} w^{f}_{ik}, \tag{3.30}$$

where $\epsilon^{f}_{kj,n} N_{kj,n}/N_{kj,1}$ gives the risk-adjusted NAV in round n as a fraction of the initial NAV. We use equation (3.20) to find that

$$v^{f}_{ikj,n} - v^{f}_{ikj,n+1} = \epsilon^{f}_{kj,n} \frac{N_{kj,n}}{N_{kj,1}} w^{f}_{ik} - \epsilon^{f}_{kj,n+1} \frac{N_{kj,n} - \Delta N_{kj,n}}{N_{kj,1}} w^{f}_{ik}, \tag{3.31}$$

and equation (3.28) to find that the loss that i suffers on its fund shares portfolio in round $n+1$ after j's default due to counterparty risk and shareholder contagion losses is given by

$$L^{f}_{ij,n+1} = \sum_{k\in\mathcal{F}} v^{f}_{ikj,n} - v^{f}_{ikj,n+1} = \sum_{k\in\mathcal{F}} \left(\Delta\epsilon^{f}_{kj,n} \frac{N_{kj,n}}{N_{kj,1}} + \epsilon^{f}_{kj,n+1} \frac{\Delta N_{kj,n}}{N_{kj,1}} \right) w^{f}_{ik} \tag{3.32}$$

The first term on the right hand side of equation (3.32) gives the counterparty risk contagion that i suffers due to risk adjustment $\Delta\epsilon^{f}_{kj,n}$, multiplied by the "interaction term" $N_{kj,n}/N_{kj,1}$. The second term gives the shareholder contagion that i suffers due to the relative NAV loss $\Delta N_{kj,n}/N_{kj,1}$, multiplied by the interaction term

$\epsilon^{f}_{kj,n+1}$. The interaction terms ensure that i's cumulative losses over the rounds can never exceed the initial asset values. We apply the counterparty risk contagion channel first and the shareholder channel second, which is why the interaction term $\epsilon^{f}_{kj,n+1}$ already reflects k's loss in round n but $N_{kj,n}/N_{kj,1}$ does not. (The choice of which contagion channel to apply first is arbitrary as it does not affect our results.)

Tradable Securities

The value of a bond or MMI is affected by both overlapping portfolio contagion and counterparty risk contagion. Moreover, the value of an equity share is affected by overlapping portfolio contagion and shareholder contagion. The overlapping portfolio contagion channel reduces the value of a tradable security based on a sudden increase in the security's supply, whereas the counterparty risk and shareholder contagion channels reduce the security's value based on the decrease in the equity (buffer) of the issuer. In informationally efficient markets, the market price of a security reflects both the security's supply and the performance of its issuer (as well as demand and other factors).

The performance of the issuer k of an equity share e may be represented by the share's book value. To reflect the supply of the security, we introduce the liquidity factor $\hat{P}^{\tau}_{kj,n}$ of a tradable security $\tau \in \{b, m, e\}$ issued by institution $k \in \mathcal{H}$. $\hat{P}^{\tau}_{kj,n}$ is a nondimensional factor of order one and the initial liquidity factor is normalized to one;

$$\hat{P}^{\tau}_{kj,1} = 1. \tag{3.33}$$

Absent counterparty risk or shareholder contagion, the liquidity factor $\hat{P}^{\tau}_{kj,n}$ multiplied by 1 ZAR gives the market price $P^{\tau}_{kj,n}$. The liquidity factor evolves according to

$$\hat{P}^{\tau}_{kj,n+1} = \hat{P}^{\tau}_{kj,n} - \Delta\hat{P}^{\tau}_{kj,n}, \tag{3.34}$$

where $\Delta\hat{P}^{\tau}_{kj,n}$ represents the price impact in the n^{th} round after j's default. We use the linear price impact implied by the liquidity weighted portfolio overlap (eq. 3.11), such that the liquidation of the portfolios of all in-default institutions $\mathcal{D}_{j,n}$ causes a price impact

$$\Delta\hat{P}^{\tau}_{kj,n} = \sum_{q \in \mathcal{D}_{j,n}} \frac{w^{\tau}_{qk}}{D^{\tau}_{k}}, \tag{3.35}$$

where the market depth (eq. 3.12) is assumed to be constant and we will only consider market depths large enough that the market price $P^{\tau}_{kj,n}$ is guaranteed to remain positive.

We assume that the market price of an equity share e issued by k is given by

$$P^e_{kj,n} = \frac{e_{kj,n}}{e_{kj,1}} \hat{P}^e_{rj,n}\text{ZAR}, \tag{3.36}$$

where the first term on the right hand side gives the shares' book value in round n relative their initial book value, and the liquidity factor $\hat{P}^e_{rj,n}$ reflects that overlapping portfolio contagion can drive the market value of the shares below their book value. Consequently, the value $v^e_{ikj,n}$ of institution i's equity shares in k in round n is given by

$$v^\tau_{ikj,n} = \frac{P^e_{kj,n}}{P^e_{kj,1}} w^e_{ik} = \epsilon^e_{kj,n} \hat{P}^e_{kj,n} w^e_{ik}, \tag{3.37}$$

where $P^e_{kj,n}/P^e_{kj,1}$ gives the shares market price in round n as a fraction of their initial market price (which is normalized to 1 ZAR), and we have used that $\epsilon^e_{kj,n} = e_{kj,n}/e_{kj,1}$ (see equations 3.23 and 3.27).
Similar to the market price of equity shares, we assume that the market price of bonds or MMIs $\tau \in \{b, m\}$ issued by k is given by

$$P^\tau_{kj,n} = \epsilon^\tau_{kj,n} \hat{P}^\tau_{rj,n}\text{ZAR}, \tag{3.38}$$

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

$$v^\tau_{ikj,n} = \frac{P^e_{kj,n}}{P^e_{kj,1}} w^\tau_{ik} = \epsilon^\tau_{kj,n} \hat{P}^\tau_{kj,n} w^\tau_{ik} \tag{3.39}$$

Comparing equations (3.39) and (3.37) shows that the values of all tradable securities are modeled equivalently, so the corresponding losses are too. Note that this is only the case because of our particular choice of risk-adjustment factor (eq. 3.27) and is not true in general. The loss that i suffers on its tradable securities of type $\tau \in \{b, m, e\}$ in round $n+1$ after j's default is given by the decease in the securities' value due to overlapping portfolio contagion and counterparty risk or shareholder contagion;

$$L^\tau_{ij,n+1} = \sum_{k \in \mathcal{H}} v^\tau_{ikj,n} - v^\tau_{ikj,n+1} = \sum_{k \in \mathcal{H}} \left(\Delta\epsilon^\tau_{kj,n} \hat{P}_{kj,n} + \epsilon^\tau_{kj,n+1} \Delta\hat{P}_{kj,n} \right) w^f_{ik}, \tag{3.40}$$

where we have used equations (3.34), (3.28), and (3.39) (similar to how we found equation 3.31).

The first term on the right hand side of equation (3.40) gives the counterparty risk or shareholder contagion that i suffers due to buffer loss $\Delta\epsilon^{\tau}_{kj,n}$, multiplied by the interaction term $\hat{P}_{kj,n}$. Similarly, the second term gives the overlapping portfolio contagion that i suffers due to price impact $\Delta\hat{P}_{kj,n}$, multiplied by the interaction term $\epsilon^{\tau}_{kj,n+1}$. The interaction terms again ensure that i's cumulative losses over the rounds cannot exceed the initial asset values. We apply the counterparty risk shareholder contagion channel first and the overlapping portfolio contagion channel second, which is why the interaction term $\epsilon^{\tau}_{kj,n+1}$ already reflects k's loss in round n but $\hat{P}_{kj,n}$ does not. This completes the calculation of the losses $L^{\alpha}_{ij,n+1}$ in equation (3.16) and, accordingly, the n^{th}-order exposures in equation (3.15).

3.4 Results

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

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

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

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

These results are derived in the following sections. We discuss system-wide exposures, sectoral exposures, individual institutions' exposures, and stressed exposures, in that order. Each exposure depicted below is the mean exposure calculated over 1000 realizations of the reconstructed interbank network. The number of samples is sufficiently large that the standard errors of the results are negligible.

3.4.1 System-Wide Exposures

Figure 3.7 shows the exposures (as % of the system's total assets) and HSEs up to n^{th} order of the South African financial system to the default of bank j, where j is one of the six large banks. The system's exposure up to n^{th} order is calculated as the sum of the exposures of all banks' and funds' exposures up to n^{th} order. The figure shows that higher-order exposures are substantial, in particular in comparison to the first-order exposures. This is highlighted by the HSEs plotted in Figure 3.7b, which show that exposures to Absa and Capitec may be underestimated by over 50% when ignoring higher-order exposures. (Note that the HSE up to first order is zero by definition.) Of the higher-order exposures, the second-order exposure is the most severe, as can be seen from the substantial jump in exposure from $n = 1$ to $n = 2$, and the higher-order exposures level out as $n \to 5$.

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

Both figures are limited to $n \leq 5$. The reason for this is threefold. First, as noted in Section 3.3.4, any model has finite accuracy, and because inaccuracies compound, the accuracy of the higher-order exposures is smaller for larger n. As shown in Figure A.1a in the appendix, the distribution of exposures (due to the random realizations of the reconstructed interbank network) fans out to quite substantial levels by $n = 5$. As we do not know the true interbank network, the true exposures may lie anywhere within this distribution. Second, exposures or HSEs up to n^{th} order level out for large n. The vast majority of higher-order exposures materialize as second and third order exposure. This also suggests that the vast majority of higher-order exposures accumulate when our confidence in the accuracy of the model is greatest.

Third, when losses are particularly large and percolate through the financial system for an increasingly large number of rounds, it becomes more likely that the central bank will intervene to stabilize the system. Hence, large exposures of very high order will likely not materialize.

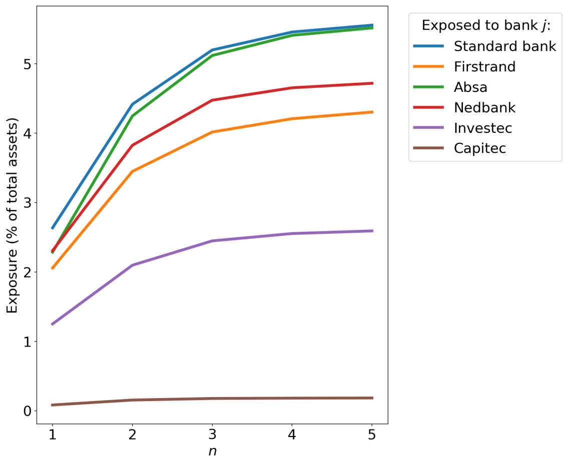

(a) Exposure up to n^{th} order

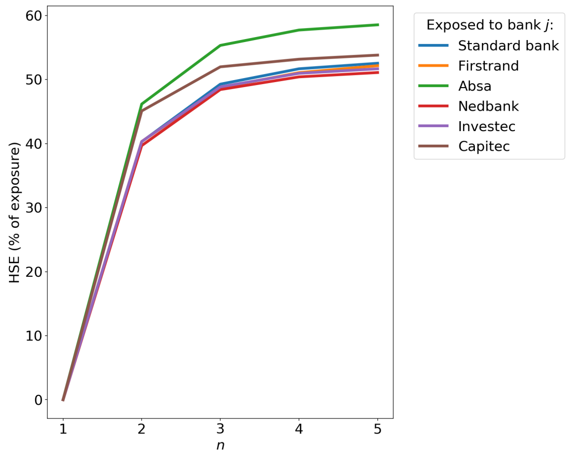

(b) HSE up to n^{th} order

Figure 3.7: **Exposure and higher-order share of exposure (HSE) up to n^{th} order of the South African financial system to the six largest banks.** Plot (a) shows the exposure (as % of the system's total assets) up to n^{th} order of the South African financial system to the default of bank j, where j is one of the six large banks and the system's exposure is calculated as the sum of all banks' and funds' exposures. (Note that the six banks appear in the legend in descending order of total asset size.) The exposure increases substantially from $n = 1$ to $n = 2$, which corresponds to the second-order exposures. Further increases in exposure level out as n approaches 5. Plot (b) shows HSEs up to n^{th} order, which are substantial for all $n > 1$. (Note that the HSE up to order $n = 1$ is zero by definition.) In particular, the HSEs indicate that exposures are underestimated by more than 50% when ignoring higher-order exposures.

3.4.2 Sectoral Exposures

Figures 3.8-3.11 break down the exposures in Figure 3.7 by sector. Figure 3.8 shows the exposures and HSEs of the banking sector to the six largest banks, Figure 3.9 those of the MMF sector, Figure 3.10 those of the OF sector and Figure 3.11 those of the FoF sector.

Figure 3.8 shows that exposures of the banking sector to the six large banks are small, but that the HSE is substantial (above 25 percent) in all cases. This shows that banks in South Africa have significant higher-order exposures to one another. Furthermore, note that Figure 3.8 looks qualitatively different from Figure 3.7, as the vertical ordering of exposures to the six banks is different across figures. This is also reflected in Figures 3.9-3.11. Hence, the sectoral exposures vary qualitatively across sectors and cannot by proxied by the system-wide exposures.

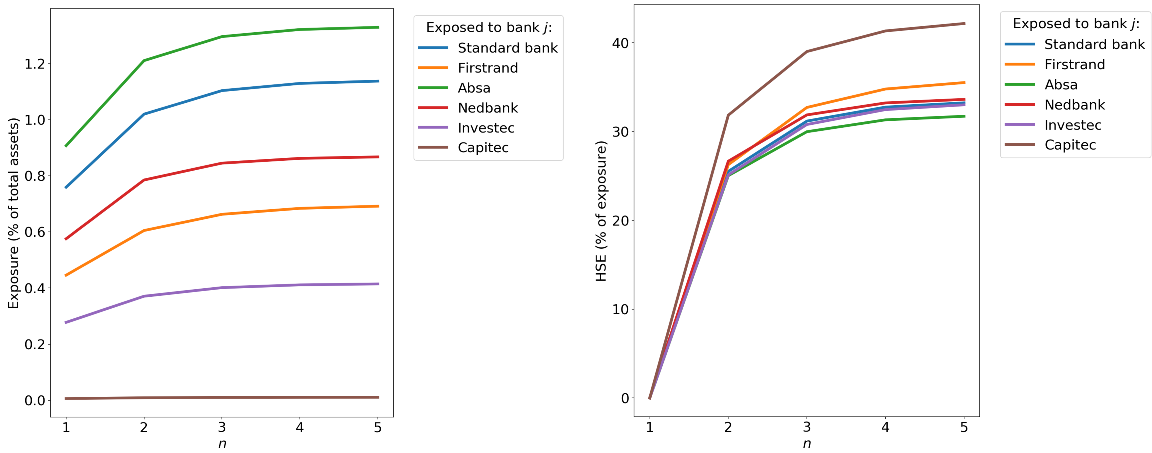

(a) Exposure up to n^{th} order

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

Figure 3.8: **Exposure and HSE up to n^{th} order of the banking sector to the six largest banks.** Plot (a) shows the exposure (as % of the sector's total assets) up to n^{th} order of the South African banking sector to the default of bank j, where j is one of the six large banks and the sector's exposure is the sum of the banks' exposures. Plot (b) shows the corresponding HSE up to n^{th} order. Exposures of the banking sector to the largest six banks are small, yet the HSE of the exposures is substantial.

The MMF sector makes large investments in South Africa's six largest banks, which is reflected in the substantial first-order exposures shown in Figure 3.9. Although the sector's first-order exposures to all six banks remain below the 30% large exposure limit, the higher-order exposures push the exposure to Absa, Nedbank and Standard bank beyond this limit. (Although regulation only requires direct exposures to fall below 30% large exposure limit, higher-order exposures should arguably be recognized too.) Note that the HSE to Capitec is particularly large, but the MMF sector's higher-order exposures to Capitec are negligible in absolute terms.

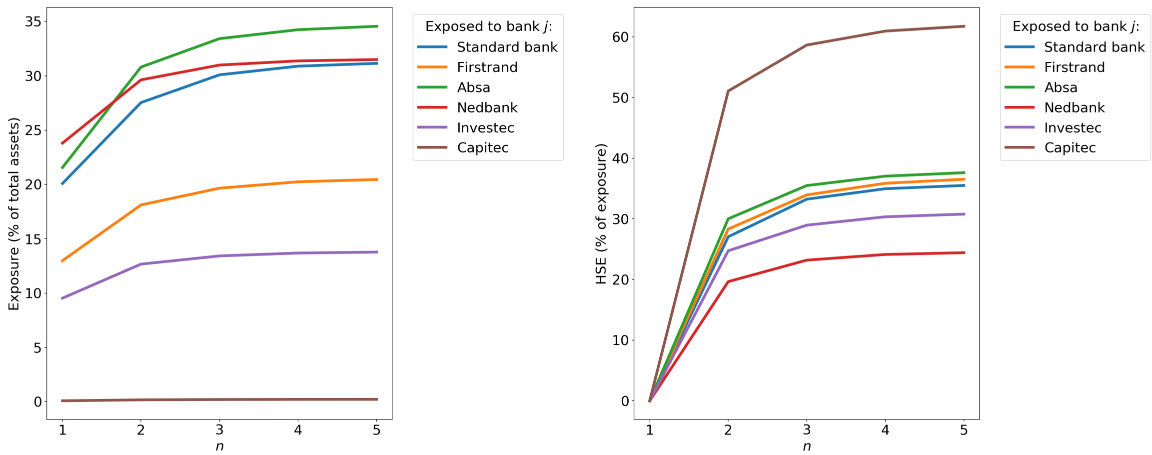

(a) Exposure up to n^{th} order

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

Figure 3.9: **Exposure and HSE up to n^{th} order of the MMF sector to the six largest banks.** Plot (a) shows the exposure (as % of the sector's total assets) up to n^{th} order of the South African MMF sector to the default of bank j, where j is one of the six large banks and the sector's exposure is the sum of the MMFs' exposures. Plot (b) shows the corresponding HSE up to n^{th} order. The MMF sector has substantial first-order exposures to the largest six banks, as banks are the main recipients of MMFs' investments. Moreover, the higher-order exposures push the MMF sector's exposures to Absa, Nedbank and Standard bank beyond the regulatory limit of 30%.

Figure 3.10 shows that exposures of the OF sector to the six large banks are smaller than those of MMF sector. However, HSEs to all banks but Capitec are substantially larger. Ignoring the higher-order exposures to Absa would be particularly problematic and underestimate exposure by more than 60%.

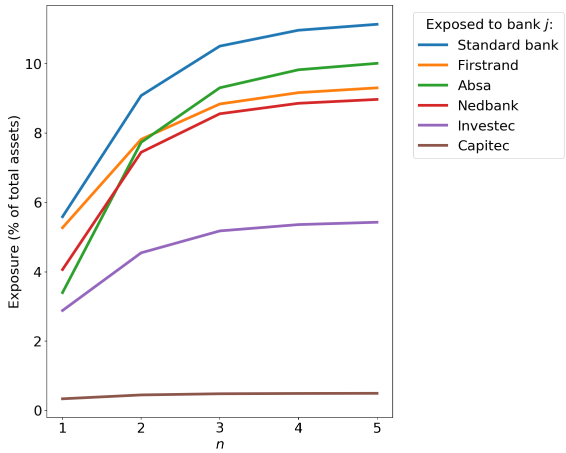

(a) Exposure up to n^{th} order

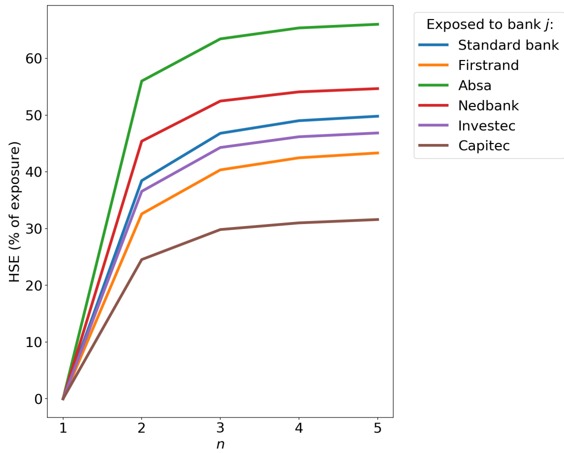

(b) HSE up to n^{th} order

Figure 3.10: **Exposure and HSE up to n^{th} order of the OF sector to the six largest banks.** Plot (a) shows the exposure (as % of the sector's total assets) up to n^{th} order of the OF sector to the default of bank j, where j is one of the six large banks and the sector's exposure is the sum of the OFs' exposures. (b) shows the corresponding HSE up to n^{th} order. Exposures of the OF sector to the six largest banks are smaller than those of the MMFs but HSEs are vast, reaching up to 60% in the case of Absa. Hence, OFs are expected to underestimate their exposure to Absa by more than 60% when only taking first-order exposures into account.

Funds of funds predominantly invest in other funds. Accordingly, the FoFs' first-order exposures to the banks are virtually non-existent and the HSEs are very close to 100%, as shown in Figure 3.11. As a result, conventional exposure metrics (capturing direct and, more recently, also indirect exposure) would find no significant exposures, even though exposures can reach as high as 18% of the total FoF sector's assets in the case of Standard bank. Furthermore, note that although the FoFs' first-order exposures are much smaller than those of the OFs', the FoFs' higher-order exposures are much larger, which shows that the FoFs' strategy of investing in other funds makes them particularly exposed to systemic risk (as compared to the OFs).

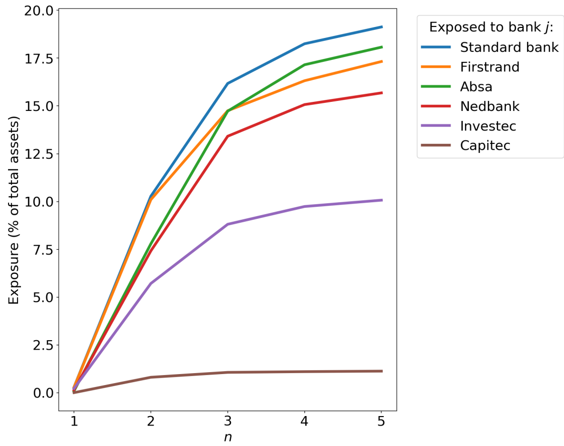

(a) Exposure up to n^{th} order

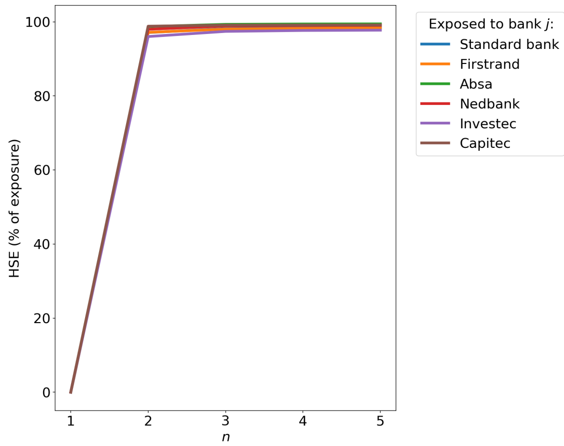

(b) HSE up to n^{th} order

Figure 3.11: **Exposure and HSE up to n^{th} order of the FoF sector to the six largest banks.** Plot (a) shows the exposure (as % of the sector's total assets) up to n^{th} order of the FoF sector to the default of bank j, where j is one of the six large banks and the sector's exposure is the sum of the FoFs' exposures. (b) shows the corresponding HSE up to n^{th} order. FoFs' first-order exposures to the six largest banks are small, as FoFs typically invest in other funds. The FoF sector's higher-order exposures are substantial; exposures to the four largest banks are around 15% of the sector's total assets. Furthermore, because first-order exposures to the largest six banks are close to zero, the FoF sector's HSEs are close to 100%. Hence, FoFs completely overlook their exposure to these banks when not taking higher-order exposures into account.

3.4.3 Individual Exposures

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

Figure 3.12 shows the HSEs and exposures (as % of the fund's total assets) of the ten largest MMFs (by total asset size) and Figure 3.13 the HSEs and exposures of the ten largest OFs. The figures use a color gradient to indicate the magnitude of the exposure of a fund on the vertical axis to one of the six large banks on the horizontal axis. (Note that the color gradients vary across plots, in order to accommodate the heterogeneity of exposure sizes.)

We are primarily interested in variations in HSEs and exposures across the funds (i.e. along the columns). The MMFs' HSEs in Figure 3.12 show little variation; the HSEs to all banks but Capitec are quite modest. (The large HSEs to Capitec are simply due to the very small first-order exposures that MMFs have to Capitec). The only exception to this is MMF6, which has virtually no investment in Investec and a near 100% HSE to Investec accordingly. The homogeneity of exposures across MMFs is due to their highly similar portfolios, as discussed in section 3.3.1.

Of the first-order, second-order and fifth-order exposures, the first-order exposures show the most variation across MMFs, driven by the small variation among MMFs' portolios. In the second-order exposures and, in particular, fifth-order exposures, most of the variation in exposures across the MMFs is damped out and a clear pattern of variation across the banks (i.e. along the rows) emerges. The pattern closely resembles the fifth-order exposures of the MMFs at the sectoral level. (Note that the MMF sector's n^{th}-order exposures are not plotted explicitly, but they can be inferred from the slopes of the exposures up to n^{th} order in Figure 3.9a). This suggests that as the order of the exposure gets higher, sector-level (or system-level) differences between the exposures that the banks generate start to dominate the variations in the MMFs' portfolios. Hence, the individual MMFs' higher-order exposures could potentially be inferred from their higher-order exposures at the sectoral level up to good accuracy. However, as shown below, this is not true for all higher-order exposures, but specific to the MMFs and due to their exceptionally similar portfolios.

Contrary to the MMFs, Figure 3.13 shows that OFs' exposures vary substantially across the OFs. Comparing the first-order and second-order exposures, we see a clear, qualitative difference in the distribution of exposures across the OFs. Hence, the second-order exposures cannot be extrapolated from the first-order exposures. This is also reflected in the variation of the HSEs across the OFs. (Note that the HSEs are predominantly driven by the second-order exposures, as the exposures of order $n \geq 3$ are substantially smaller.) Interestingly, OF8 has the lowest first-order exposures, but highest second-order exposures of all ten OFs, highlighting the danger in attempting to extrapolate higher-order exposures from first-order exposures. Similar to the MMFs, most of the variation across the OFs' exposures is damped out in the fifth-order exposures, which suggests that these exposures may be inferred from their sectoral exposures. However, we already saw this is not the case for the OFs' second-order exposures, which drive the OFs' higher-order exposures as they are the largest. Hence, although the MMFs' higher-order exposures may potentially be inferred from the sectoral exposures due to the MMFs' exceptionally similar portfolios, we conclude that in general the higher-order exposures must be modelled explicitly for the individual institutions and cannot be inferred from first-order or sectoral exposures. This conclusion is supported by Figure A.2 and Figure A.3 in the appendix, which show the exposures of the individual banks and FoFs.

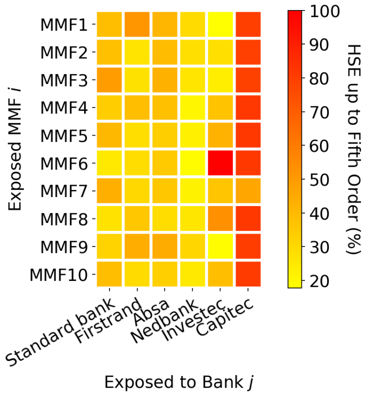

(a) HSE up to fifth-order

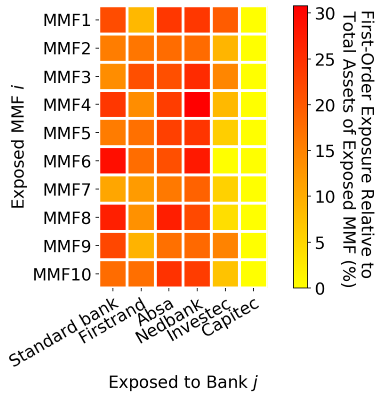

(b) First-order exposure

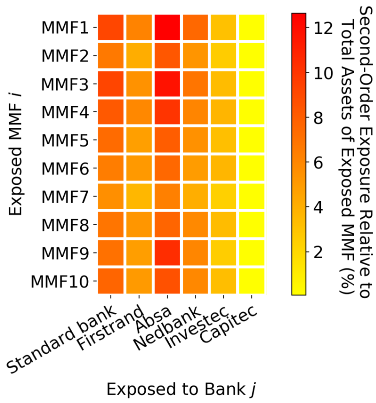

(c) Second-order exposure

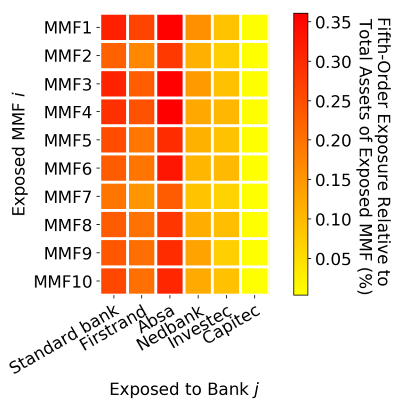

(d) Fifth-order exposure

Figure 3.12: **Individual MMFs' exposures to the six largest banks**. The figure shows the exposures (as % of each MMF's total assets) and HSEs of the ten largest MMFs. The MMFs are ordered by total asset size (descending from top to bottom) and the plots use a color gradient to indicate the HSE or exposure of an MMF on the vertical axis to a bank on the horizontal axis. (Note that the banks are ordered from left to right by descending total asset size.) The plots show little variation in HSEs and exposures across the MMFs, which is due to the MMFs' exceptionally similar portfolios. The first-order exposures (which are predominantly driven by direct exposures) in (b) show that all MMFs' investments in the four largest banks are close to the 30% limit, with Investec receiving the remainder of the investments and Capitec virtually nothing. Plots (c) and, in particular, (d) show that the higher the order of the exposure, the more the variation across the MMFs is damped out and gets dominated by the variation across the banks to which the MMFs are exposed.

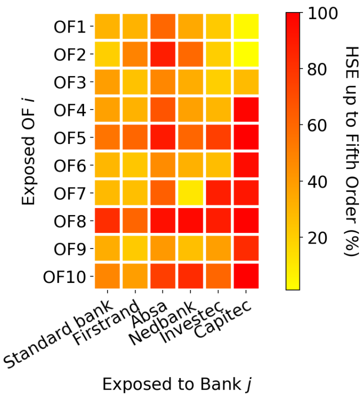

(a) HSE up to fifth-order

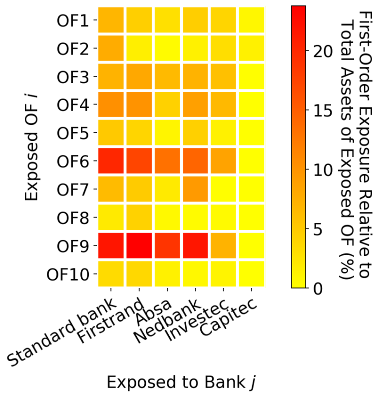

(b) First-order exposure

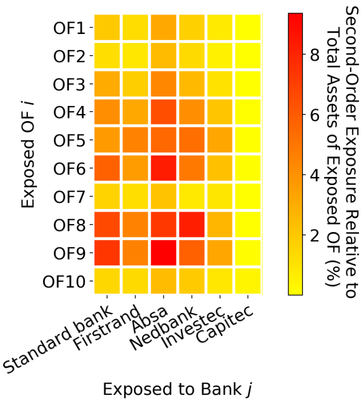

(c) Second-order exposure

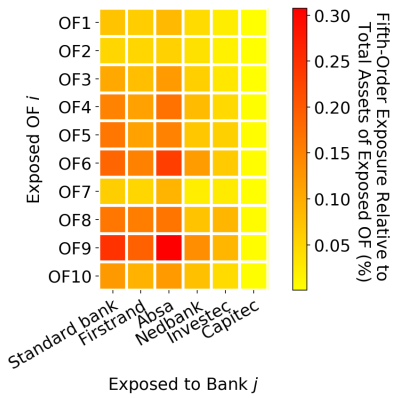

(d) Fifth-order exposure

Figure 3.13: **Individual OFs' exposures to the six largest banks**. The figure shows the exposures (as % of each OF's total assets) and HSEs of the ten largest OFs. The OFs are ordered by total asset size (descending from top to bottom) and the plots use a color gradient to indicate the HSE or exposure of an OF on the vertical axis to a bank on the horizontal axis. Plot (a) shows that HSEs vary strongly across OFs, and comparison of the first-order exposures in (b) to the second-order and fifth-order exposures in (c) and (d) clearly shows that the second-order and fifth-order exposures are qualitatively different from the first-order exposures. Hence, the higher-order exposures cannot be extrapolated from first-order exposures. Furthermore, similar to the MMFs, (d) shows that in the fifth order exposures, most of the variation across the OFs is damped out and gets dominated by the variation across the banks to which the OFs are exposed. However, this is not the case for the second-order exposures, which show substantial variation across the OFs.

3.4.4 Stressed Exposures

Figures 3.14, 3.15 and 3.16 show that higher-order exposures become particularly pronounced during times of financial distress. Put differently, higher-order exposures are at their greatest exactly when they matter most, i.e. in times of crisis when defaults are most likely to occur and, consequently, exposures are most likely to materialize into losses.

Due to adverse macroeconomic conditions (i.e. crisis scenarios), institutions may incur unexpected losses, reducing their buffers. For illustrative purposes, Figure 3.14 shows the exposure up to second-order and up to fifth-order when all banks' and funds' (initial) buffers are reduced by 20%, 40% or 60%. For comparison's sake, the exposures are also shown for the case the institutions' buffers are not changed ("0%"), or increased by 20%[15]. The gray bars show the direct exposures, so all exposures exceeding the gray bars are overlooked when ignoring indirect and higher-order exposures exposures. (Note that while higher-order exposures are affected by buffers, direct and indirect exposures are not.) The figure shows that when buffers are reduced, higher-order exposures increase substantially and start to dominate exposure. Furthermore, (a) shows that exposures up to second order to Standard Bank, Absa and Nedbank become particularly pronounced when buffers are reduced by 60%. (b) shows that exposures up to fifth order to Standard Bank, Absa and Nedbank increase substantially when buffers are reduced by 20% or 40% and that fifth-order exposures to all but Capitec are substantially increased when buffers are reduced by 60%.

[15]Banks may increase their buffers by raising capital. As funds' buffers give the losses that funds can absorb before investors start running, funds increase buffers raising investor confidence, e.g. through increasing funds' liquidity.

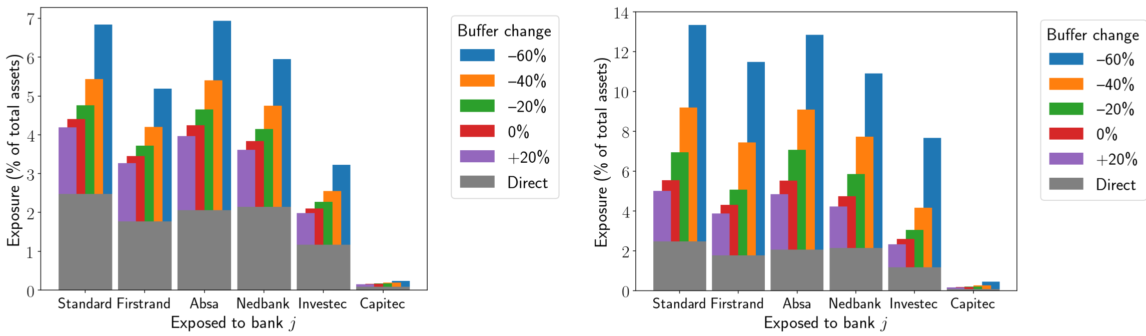

(a) Exposure up to second order

(b) Exposure up to fifth order

Figure 3.14: **Exposure of the South-African financial system to the six largest banks for various values of institutions' initial buffers.** (a) shows the exposure up to second order and (b) the exposure up to fifth order (as a % of total system assets) of the South-African financial system to the default of bank j, where j is one of the six large banks and the system's exposure is the sum of the banks' and funds' exposures. The colors indicate the percentage change applied to all banks' and funds' initial buffers. The gray bars show the direct exposures, which are not affected by the buffers (and neither are indirect exposures). The figures show that higher-order exposures become particularly pronounced when institutions' buffers are reduced, with exposures to Standard bank, Absa and Nedband increasing most strongly. More specifically, (b) shows for exposures up to fifth order, that the exposures to Standard bank, Absa and Nedband are almost doubled even when initial buffers are reduced by just 20%.

Reduced buffers are not the only reason why higher-order exposures increase in times of financial distress. During crises, market liquidity for tradable securities $t \in \{b, m, e\}$ typically falls. In illiquid markets, i.e. when the market depth is reduced, the price impact of the liquidation of a defaulted institution's portfolio increases. Figure 3.15 shows the exposure up to second-order and up to fifth-order for various multiples of the baseline value of the market depth divisor $\mu = .21$. When the market depth divisor $\mu = 0$, market depths are infinite so the overlapping portfolio contagion channel is effectively turned off, whereas when the market depth divisor $\mu = .84$, market depths are reduced by 75% and the overlapping portfolio contagion channel is strongly amplified. The gray bars again show the direct exposures, which are not affected by the market depths (while the indirect and higher-order exposures are). The figure shows that higher-order exposures are exacerbated when market depths fall, but not as much as when buffers fall. Furthermore, note in Figure 3.15 that the increases in exposures are quite proportional to the market depth divisor μ, while exposures increase superlinearly in the buffer decrease in Figure 3.14.

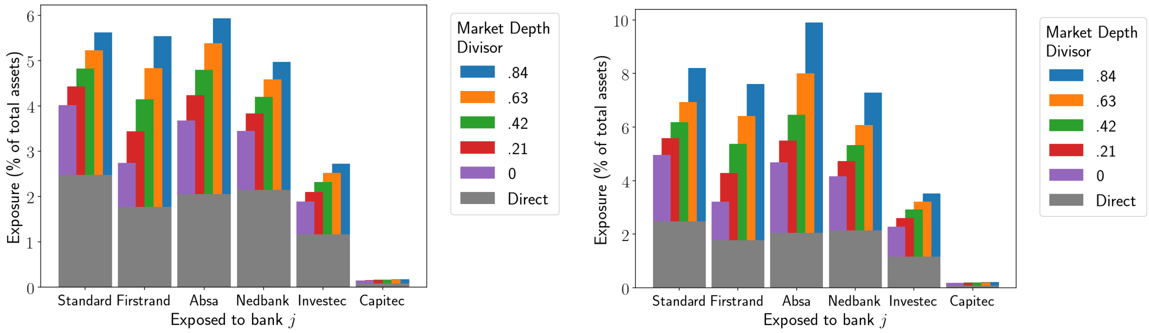

(a) Exposure up to second order

(b) Exposure up to fifth order

Figure 3.15: **Exposure of the South-African financial system to the six largest banks for various market depths.** (a) shows the exposure up to second order and (b) the exposure up to fifth order (as a % of total system assets) of the South-African financial system to the default of bank j, where j is one of the six large banks and the system's exposure is the sum of the banks' and funds' exposures. The colors indicate the market depth divisor μ applied to all securities. The gray bars show the direct exposures, which are not affected by market depths. The figures show that both the exposures up to second and fifth order to all banks but Capitec increase substantially when market depths fall.

Based on the SARB 2016 Financial Stability Review (SARB, 2016a), we formulate a macroeconomic stress scenario, which consists of a 25% reduction in institutions' buffers and 50% reduction in the market depth for all tradable securities. Figure 3.16 shows that when all banks and funds are subjected to the stress scenario, higher-order exposures are substantially larger (compared to Figure 3.7). Furthermore, exposures no longer level out by $n = 5$, so the practice adopted in this paper of only considering exposure up to order $n = 5$ may still underestimate exposure in times of financial distress. This is confirmed by Figure A.1 in the Appendix, which shows that exposures only level out by about $n \approx 8$ when institutions are subjected to the stress scenario.

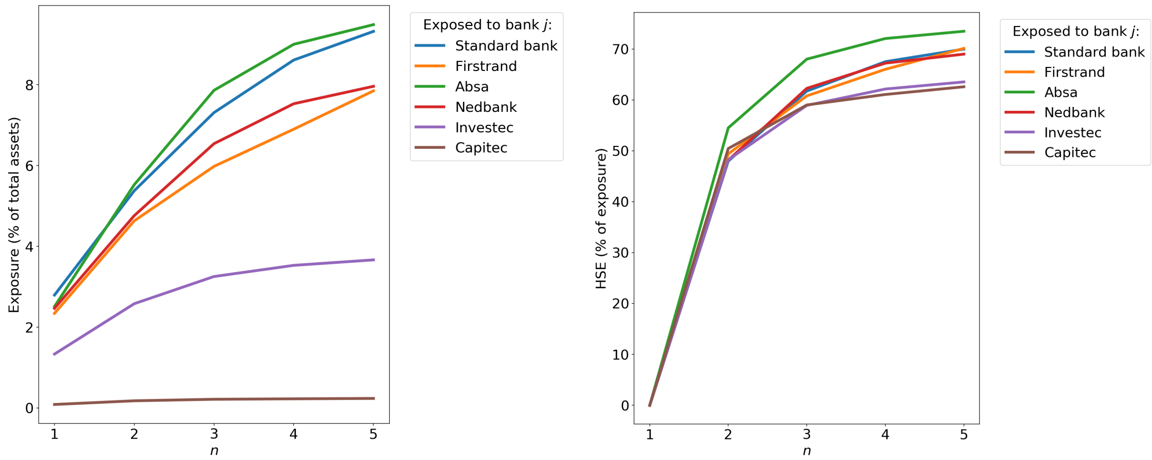

(a) Exposure up to n^{th} order (b) HSE up to n^{th} order

Figure 3.16: **Stressed exposure and HSE up to n^{th} order of the South African financial system to the six largest banks.** All banks and funds are subjected to the macroeconomic stress scenario, which consists of a 25% reduction in the institutions' buffers and 50% reduction in the market depth of all tradable securities (i.e. for the market depth divisor $\mu = .42$). (a) shows the exposure (as % of the system's total assets) up to n^{th} order of the South African financial system to the default of bank j, where j is one of the six large banks and the system's exposure is the sum of the banks' and funds' exposures. (b) shows the corresponding HSE up to order n. Figure (a) shows that exposures no longer level out by $n = 5$. Moreover, figure (b) shows that the stressed HSEs exceed 60% across all six banks. Hence, higher-order exposures become particularly dominant during crises, which is exactly when exposures are most important (as that is when defaults are most likely to cause exposures to materialize into losses).

3.5 Discussion

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

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

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

2. *Capital requirements.* Capital requirements are generally calibrated to direct exposures only, thereby failing to account for the indirect and higher-order exposure. For Global Systemically Important Banks (G-SIBs) this omission is most striking. G-SIBs are required to hold additional capital (referred to as the "G-SIB surcharge") according to their level of systemic importance; G-SIBs that are more systemically important thus face a higher capital surcharge. The current Basel accords use an "indicator-based measurement approach" to measure a G-SIBs systemic importance, which quantifies systemic importance in terms of the impact that the bank's failure could have on the global financial system and the wider economy (BIS, 2018a). While this could be achieved through measuring the direct, indirect and higher-order exposure of the system to its failure, in practice five proxies for its systemic importance are used. One of these is "interconnectedness", which accounts for 20% of the bank's total score. Currently, interconnectedness is assessed through "intra-financial system assets and liabilities" and "securities outstanding", which

strongly correlate to the size of the bank's balance sheet and are not sensitive to how portfolios overlap (Cont and Schaanning, 2019). Our results suggest that this approach may, in some relevant cases, significantly underestimate exposures and would thereby underestimate the systemic importance of the bank, resulting in a surcharge that is too low. We have shown that higher-order exposures are not necessarily correlated with balance sheet size, so extrapolating from the current indicators would not address the problem. Since this methodology is also used to assess the systemic importance of financial institutions more broadly, some systemically important institutions that generate large higher-order exposures may not be identified at all. Incorporating higher-order exposures in calibration exercises of the capital requirements could thus offer insights that are, at the very least, complementary to those obtained using existing methods – expanding the set of systemically important institutions and increasing the capital surcharge for some of them.

3. *Stress test models.* Today's regulatory stress tests are microprudential in nature. Microprudential stress tests are forward-looking exercises that assess the resilience of an institution (as e.g. measured by capital levels) to adverse economic conditions. While these tests measure asset losses resulting from direct exposures, they fail to measure the additional asset losses that could materialize through indirect and higher-order exposures. For stress tests to fulfil their basic function of assessing exposure and institutions' resilience to risk, capturing higher-order exposures is essential. To be able to do that, stress test models should include multiple interacting contagion channels and be designed to study system-wide dynamics (Farmer et al., 2020), because those elements drive higher-order exposures. We have demonstrated that compensating for the lack of explicit system-wide models with direct loss multipliers is inadequate as it gives distorted outcomes. To better assess the resilience of financial institutions, stress tests should thus not only measure the capital impact of asset losses from direct exposures, but also from indirect and higher-order exposures.

4. *Resolution.* In the 2007-2008 financial crisis regulators stood before the terrible choice of either bailing-out a SIB or liquidating it in a potentially disorderly manner. Since then, new resolution regimes have been developed around the world that enable regulators to resolve banks in an orderly manner through a bail-in, thereby aiming to avoid undue disruption to the bank's activities

and contagion effects to the rest of the economy. In a bail-in, the debt of the bank's creditors is written down and in part converted to equity to recapitalize it and revive its short-term viability. To decide whether a bank should be liquidated or bailed-in, authorities must determine that a failing bank cannot go through normal insolvency proceedings without harming public interest and causing financial instability (Kleinnijenhuis et al., 2021). To make this call, regulators must assess the stability implications of a bank failure. This requires a measurement of the losses that the system could suffer if the bank were liquidated. Put differently, it precisely requires the regulator to measure first-order and higher-order exposure of other institutions to the bank's failure. Unfortunately, a measure of higher-order exposure is completely missing in this toolkit.

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

To operationalize the analysis of higher-order exposures, data-gathering mandates should cover more granular data across a wider range of financial institutions. Because higher-order exposures can only be quantified with contagion models that explicitly capture the multi-layered financial network, it is important that regulators and supervisors have the requisite data, i.e.: bilateral, contract-level data on individual institutions' assets and liabilities. The width of the distributions of exposures that we find (resulting from the random realisations of the reconstructed interbank network) highlight the vital importance of having insight into the financial system's network structure. The observation that substantial higher-order exposures may exist between seemingly unconnected parts of the financial system suggests that the scope of data-gathering mandates should be wide, spanning the entire financial system and potentially parts of the real economy (Farmer et al., 2021b, Ullersma and van Lelyveld, 2021). To support this point, we note that without combining our data on the investment fund sector with that on the banking sector, we would not have been able to make the observation that the fund-of-fund sector is highly exposed to the banking sector in South-Africa, even though it has virtually no investments in it.

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

Chapter 4

Scenario-Free Analysis of Financial Stability with Interacting Contagion Channels

Based on the eponymous paper by Garbrand Wiersema, Alissa M. Kleinnijenhuis, Thom Wetzer and J. Doyne Farmer which has been accepted for publication in the Journal of Banking and Finance. This work has been presented at the 2021 Royal Economics Society Conference, the 2021 INET Young Academics Networks Conference at the University of Cambridge, the 2020 ECB-Oxford Workshop on Interconnectedness and Financial Stability at the ECB, the 2019 WEHIA conference at City University (where it was nominated by the Bank of England for best policy-relevant paper), and the 2019 Oxford-ETH Workshop on Mathematical and Computational Finance.

4.1 Summary

Financial stress tests that capture multiple interactions between contagion channels are conditional on specific, subjectively-imposed stress scenarios. Eigenvalue-based approaches, in contrast, provide a scenario-independent measure of systemic stability, but so far only handle a single contagion mechanism. We develop an eigenvalue-based approach that brings the best of both worlds, enabling the analysis of multiple interacting contagion channels without the need to impose a subjective stress scenario. Our model captures the solvency-liquidity nexus, which allows us to demonstrate that the instability due to interacting channels can far exceed that of the sum of the individual channels acting in isolation. The framework we develop is flexible and allows for calibration to the microstructure and contagion channels of

real financial systems. Building on this framework, we derive an analytic stability criterion in the limit of a large number of institutions that gives the instability threshold as a function of the relative size and intensity of contagion channels. This analytical formula requires comparatively little data to elucidate the mechanisms that drive instability in real financial systems and thus complements the insights gained from traditional stress tests.

4.2 Introduction

One of the revelations of the financial crisis was the importance of *systemic risk*.[1] Systemic risk is transmitted between institution to spread and amplify across the financial system through mechanisms referred to as *contagion channels* (Allen and Gale, 2000). Risk control measures that are prudent for a single institution acting on its own may be counterproductive when many institutions act in unison.[2] This problem is complicated by the fact that the financial system is heterogeneous, with different types of actors, such as banks, pension funds, hedge funds, money market funds and insurance companies, and many types of interactions.[3] This makes the financial system a prime example of a complex system and raises the key challenge for policymakers to capture the complex microstructure of the system in models of financial stability (Arinaminpathy et al., 2012, Aymanns et al., 2018).
The microstructural models currently used by policymakers to evaluate the stability of financial systems are premised on stress scenarios consisting of hypothetical exogenous shocks that could potentially threaten the stability of the system.[4] This approach has the obvious drawback that the specification of such scenarios is inherently subjective, giving rise to debates about their realism and relevance to current market conditions (Borio et al., 2014, Aymanns et al., 2018). Moreover, scenarios are by their very nature not comprehensive – the financial system might be stable in one set of scenarios and collapse in other unforeseen or mis-specified scenarios. This challenge can be partially overcome by analyzing *ensembles* of shock scenarios (see e.g. Elsinger et al., 2006, Montagna et al., 2020), but because the space of

[1] See e.g. Cont et al. [2010], Gai and Kapadia [2010], Fouque and Langsam [2013], Glasserman and Young [2016].

[2] See e.g. Adrian and Shin [2010], Thurner et al. [2012], Adrian and Shin [2014], Aymanns and Farmer [2015], Aymanns et al. [2016].

[3] See e.g. Allen and Babus [2009], Gai et al. [2011], Arinaminpathy et al. [2012], Caccioli et al. [2012], Cont et al. [2013], Cont and Schaanning [2017], Aymanns et al. [2018], Farmer et al. [2020].

[4] See e.g. Burrows et al. [2012], IMF [2014], Kok and Montagna [2016], Budnik et al. [2019].

potential scenarios is effectively infinite and the probability of specific scenarios is unknown, the problem is never entirely overcome.

An alternative method that does not suffer from these shortcomings explicitly models the financial network as a dynamical system, so that its stability can be analyzed in terms of its eigenvalues. This approach has been used for studying the effect of contagion channels in isolation (Caccioli et al., 2014, Bardoscia et al., 2017, Cont and Schaanning, 2019). However, financial systems have multiple interacting contagion channels, and studies have shown that the interaction of multiple channels can dramatically amplify instability compared to channels operating in isolation.[5] Up until now there has been no general method for treating interacting channels as a dynamical system. This causes a tenuous state of affairs for policymakers: To take into account interacting contagion channels, they are forced to rely on subjectively imposed stress scenarios.

Our key contribution in this paper is to offer a novel approach that combines the best of both worlds: a systematic method to analyze a financial network with multiple interacting contagion channels as a dynamical system, which significantly complements our ability to understand and monitor the stability of the financial system. This novel approach is realized by expressing contagion channels in terms of shocks to the liquidity and solvency of institutions, which allows us to reduce the multiple layers of the financial system's contagion network to a simple two-layer system and study the nexus of the interactions between liquidity and solvency. Using this method, we compute the linear stability of a financial system exposed to small shocks in a general setting. This makes it possible to estimate the stability of the financial system without having to impose a specific, subjective risk scenario. In contrast to methods such as those used by the EBA and the FED (EBA, 2018, FED, 2018), for example, this has the potential to yield accurate estimates of financial stability that are robust under a wide range of stress scenarios.

The paper proceeds as follows. Section 4.3 introduces our framework, which we refer to as the "shock transmission matrix", explains our modelling of contagion dynamics, and derives our main results. Section 4.4 applies our framework to randomly generated financial systems. The purpose of this exercise is to elucidate the interactions between solvency- and liquidity-mediated contagion channels and to demonstrate that neglecting these interactions may lead to a potentially substantial

[5] See e.g. Caccioli et al. [2013], Poledna et al. [2015], Kok and Montagna [2016], Cont et al. [2020], Detering et al. [2021].

overestimation of stability. Section 4.5 concludes with a discussion of the implications of our findings for financial stability policy and financial stress testing practices.

4.3 Capturing the Solvency-Liquidity Nexus

Financial contagion can take many forms, many of which have been extensively studied (see e.g. Allen and Gale, 2000, Eisenberg and Noe, 2001, Gorton and Metrick, 2012). In this paper, we analyze four principal contagion channels of the financial system, which we call funding contagion, overlapping portfolio contagion, counterparty risk contagion, and deleveraging contagion.

Funding contagion occurs when a borrowing institution depends on short-term loans to provide liquidity and runs the risk that the lender might withdraw its loans (Diamond and Dybvig, 1983, Acharya and Skeie, 2011, Caccioli et al., 2013). *Overlapping portfolio contagion* can materialize when two institutions hold common securities. If either institution sells securities this drives prices down, lowering the securities' value.[6] *Counterparty risk* occurs when a lender runs the risk that a borrower might default.[7] Finally, *deleveraging contagion* takes place when an institution uses borrowed funds to purchase assets.[8] 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.

The culmination of a severe financial crisis is usually the default of one or more institutions (Brunnermeier, 2009, Roukny et al., 2013). A default can be forced by insolvency or illiquidity. *Insolvency* occurs when asset values drop to the point where equity becomes negative – that is, when the value of an institution's liabilities exceeds that of its assets (Amini et al., 2016). Default due to *illiquidity*, on the other hand, occurs when an institution is unable to meet its payment obligations (Cont and Schaanning, 2017). Insolvency and liquidity can be related, but are analytically

[6]See e.g. Adrian and Shin [2010], Caccioli et al. [2013, 2014, 2015], Duarte and Eisenbach [2018], Cont and Schaanning [2017, 2019].

[7]See e.g. Eisenberg and Noe [2001], Furfine [2003], Gai and Kapadia [2010], Battiston et al. [2012], Elliott et al. [2014], Acemoglu et al. [2015], Bardoscia et al. [2015, 2017].

[8]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].

distinct: an institution can default due to a liquidity shock even when it is solvent, and vice versa. During financial crises, liquidity tends to be the more direct threat; an institution may survive temporary insolvency by maintaining liquidity and regaining its solvency at a later date, but for our purposes we neglect this possibility here. In normal economic times, a solvent institution is expected to borrow to avert a liquidity shortage. In times of economic crisis, however, this may not be possible because lending markets malfunction due to uncertainty about asset values, escalating collateral requirements, liquidity hoarding and capital flight, etc. (Gorton and Metrick, 2012).

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 of liquid assets and a valuation shock as a drop in the (expected) value of an institution's assets.[9]

The key insight is that the four contagion channels we distinguish here can be described in terms of the propagation of liquidity and valuation shocks and the conversion of one type of shock into the other:

- *Propagation of liquidity shocks by funding contagion:* If institution i depends on a short-term loan from institution j, if j suddenly withdraws the loan to meet a liquidity shock it receives, then this causes a liquidity shock to i.
- *Propagation of valuation shocks by counterparty risk contagion:* If a valuation shock causes institution i's probability of default to rise, the risk-adjusted value of its debt to institution j falls, causing a valuation shock to j.
- *Conversion of liquidity shocks to valuation shocks by overlapping portfolio contagion:* If institution i suffers a liquidity shock it may be forced to sell securities to raise liquidity. This depresses their price. If institution j also has a position in these securities it experiences a valuation shock.
- *Conversion of valuation shocks to liquidity shocks by deleveraging:* If a valuation shock decreases institution i's equity, its leverage rises. To return to its target leverage, the institution must raise cash to pay off its debt, essentially

[9]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.

triggering a liquidity shock to itself (we do not consider slower mechanisms to raise equity-capital, such as issuing new shares or retaining earnings).

Note that we use the term "contagion channel" to refer to a specific mechanism that propagates or converts a financial shock, and not as a reference to the shock itself.
In the remainder of this section, we show how to describe the collective dynamics of these four interacting contagion channels in a scenario-independent framework. This allows us to characterize the financial system's resilience to a wide range of liquidity and valuation shocks based on the corresponding largest eigenvalue.

4.3.1 The Shock Transmission Matrix

The interactions of the four contagion channels can be captured in a single matrix A which we call the *shock transmission matrix*, as shown in Figure 4.1a. Assume discrete dynamics and let $x^l_{t,i}$ denote the liquidity shock suffered by institution i at time t. The N-dimensional vector $\vec{x}^l_t$ gives the liquidity shocks to all institutions, where N is the number of financial institutions. Similarly, $x^v_{t,i}$ denotes the valuation shock to institution i at time t and $\vec{x}^v_t$ the N-dimensional vector of valuation shocks to all institutions. The combined shock vector $\vec{x}_t$ of length $2N$ is

$$\vec{x}_t = \begin{bmatrix} \vec{x}^l_t \\ \vec{x}^v_t \end{bmatrix}. \tag{4.1}$$

The shock transmission matrix A is the $2N \times 2N$ matrix that acts on the shock vector $\vec{x}_t$ according to

$$\vec{x}_{t+1} = A\vec{x}_t. \tag{4.2}$$

Given the distinction between the top and bottom half of $\vec{x}_t$, we decompose the shock transmission matrix into its four quadrants,

$$A = \begin{bmatrix} A^{ll} & A^{vl} \\ A^{lv} & A^{vv} \end{bmatrix}, \tag{4.3}$$

where each of the components A^{ll}, A^{lv}, A^{vl} and A^{vv} are $N \times N$ matrices, so that Eq. (4.2) can be written in the form

$$\vec{x}_{t+1} = \begin{bmatrix} \vec{x}^l_{t+1} \\ \vec{x}^v_{t+1} \end{bmatrix} = \begin{bmatrix} A^{ll}\vec{x}^l_t + A^{vl}\vec{x}^l_t \\ A^{lv}\vec{x}^l_t + A^{vv}\vec{x}^v_t \end{bmatrix}. \tag{4.4}$$

Eq. (4.4) makes explicit how the diagonal quadrant A^{ll} describes the propagation of liquidity shocks and A^{vv} the propagation of valuation shocks. The off-diagonal quadrant A^{lv} gives the conversion of liquidity to valuation shocks and A^{vl} the conversion of valuation to liquidity shocks. Figure 4.1b shows the corresponding contagion channels.

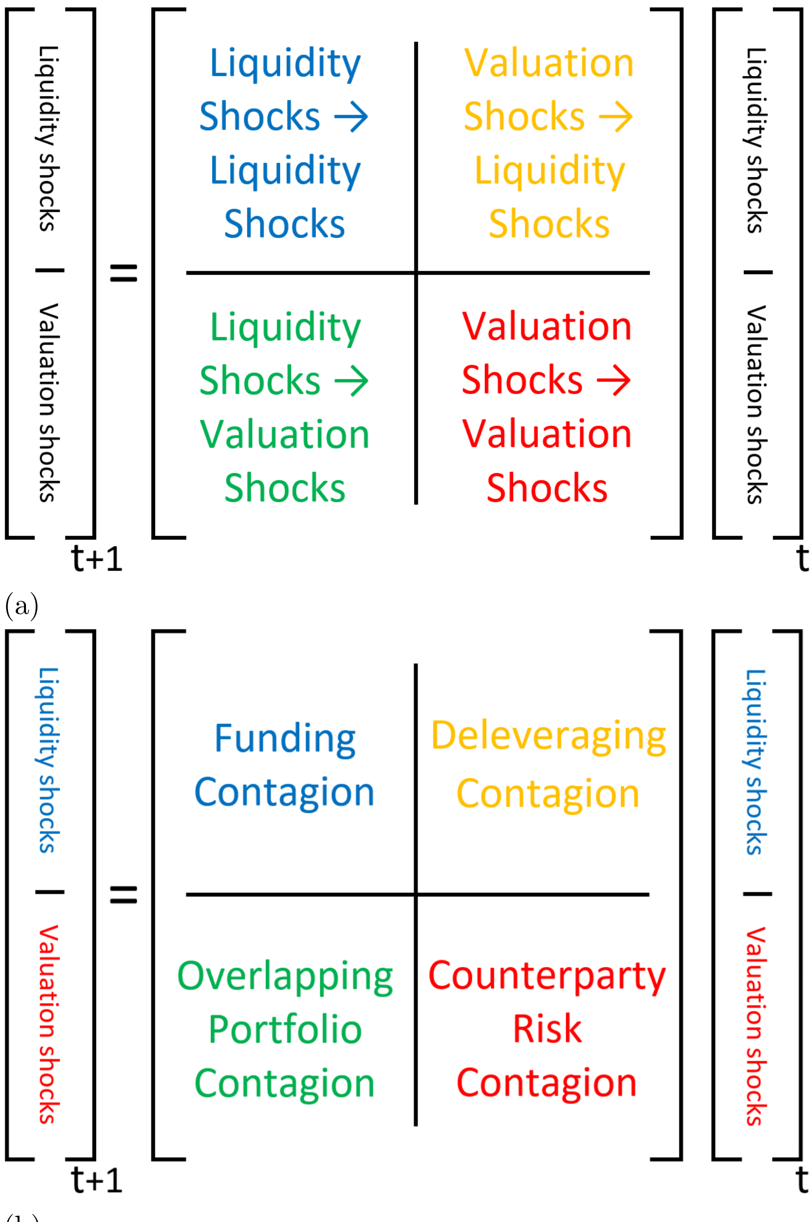

Figure 4.1: **Decomposition of shock dynamics.** The vector of shocks to the financial system can be described as a concatenation of the vector of liquidity shocks and the vector of valuation shocks to each institution. The shock transmission matrix maps the complete vector of shocks in one period to the vector of shocks in the next period. It can be decomposed into its four quadrants as shown in the figure, corresponding to the propagation and conversion of both shock types. Note the correspondence of the quadrants in (a) and (b): Funding contagion propagates liquidity shocks, counterparty risk propagates valuation shocks, overlapping portfolio contagion converts liquidity shocks to valuation shocks and deleveraging converts valuation shocks to liquidity shocks.

The shock transmission matrix A is the adjacency matrix of a weighted, directed, duplex network, where the nodes are institutions and the edges represent the transmission of shocks. Each institution is represented by a node in each layer. As shown in Figure 4.2, the top layer describes the propagation of liquidity shocks by funding contagion and is referred to as the *liquidity shock network*. The bottom layer describes the propagation of valuation shocks by counterparty risk contagion and is referred to as the *valuation shock network*. The edges between the two layers describe liquidity shocks transitioning to valuation shocks and vice versa, according to overlapping portfolio and deleveraging contagion respectively. Because we can express all four contagion mechanisms in this two-layer system, in contrast to earlier methods[10], this method does not require a separate layer for each contagion mechanism.

[10]See e.g. Caccioli et al. [2013], Kok and Montagna [2016], Poledna et al. [2015], Hüser et al. [2018], Bardoscia et al. [2018].

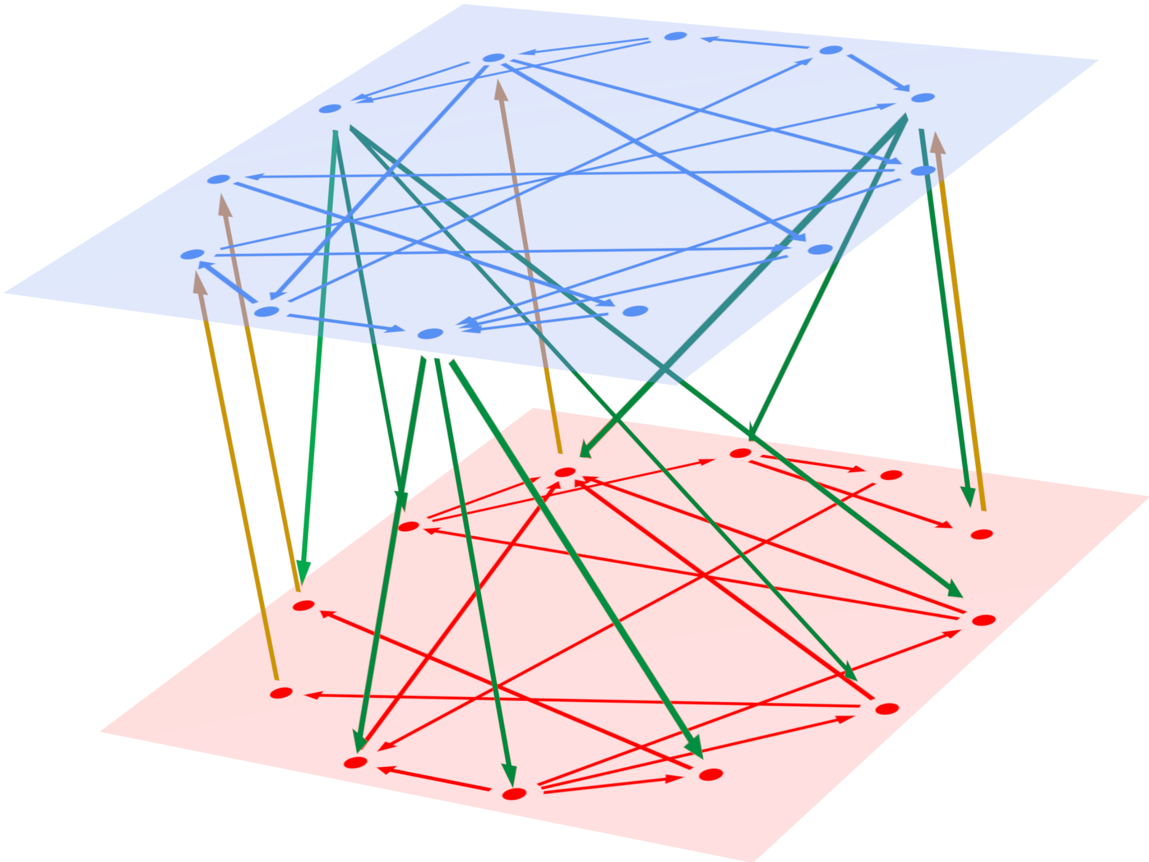

Figure 4.2: **The duplex network underlying the shock transmission matrix A.** The nodes represent institutions and the edges represent the shock transmission between institutions. In this two-layer system, the top layer represents the liquidity shock network (in blue) and the bottom layer the valuation shock network (in red). The green and yellow arrows represent interactions between these two networks; the green arrow represents conversions of shocks from liquidity to valuation, whereas yellow arrows represent the conversion of shocks from valuation to liquidity. The shock transmission matrix is a weighted adjacency matrix that describes both layers and their interactions simultaneously.

The shock transmission matrix can be used to study the system's stability and resilience to shocks. Because all its elements are non-negative, the Perron-Frobenius theorem guarantees that the matrix has a non-negative real eigenvalue greater than or equal to (the absolute values of) the matrix' other eigenvalues. This largest eigenvalue describes the systemic properties of the financial system (Caccioli et al., 2014, Bardoscia et al., 2017, Cont and Schaanning, 2019): If the largest eigenvalue is greater than one, shocks that are not orthogonal to the corresponding eigenvector are amplified without bound and the system is *unstable*; if the eigenvalue is smaller than one, shocks are damped and we refer to it as *stable*. While no system is resilient to arbitrarily large shocks, an unstable system under this definition is not even resilient to small shocks, as it amplifies them without bound over time.

In Eq. (4.2), each contagion mechanism manifests itself in a single time step t. This setup implicitly assumes that all four contagion mechanisms act equally fast, which we know is not necessarily true in reality. However, this simplifying assumption does not affect our results: As we show in the Supplementary Materials B.2, the set of conditions under which the largest eigenvalue is equal to one is independent of this assumption.

4.3.2 Institutions' Responses to Shocks

To study how a system's stability depends on its composition in terms of different types of financial institutions, we classify institutions based on the contagion they transmit. The response to a financial shock generally depends on its magnitude, but here we focus on shocks that are sufficiently small that the dynamics are approximately linear. This can be extended to deal with larger shocks by dynamically updating the shock transmission matrix as the shocks propagate.

When a liquidity shock hits, an institution may have multiple options available to respond. We assume that each institution has a *pecking order* that specifies the sequence in which it uses these options (Kok and Montagna, 2016, Hałaj, 2018). 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). We focus on shocks that are sufficiently small for us to assume that the liquidation option at the top of any institution's pecking order is not exhausted. In reality, the pecking order is institution-specific. Our methodology assumes that every institution has a pecking order, but it is in principle agnostic to what that pecking order looks like. Following the approach of Kok and Montagna [2016] and Hałaj [2018], we assume that institutions adopt the pecking order that minimizes liquidation costs. As a result, any institution that holds sufficient cash on its balance sheet can absorb liquidity shocks without causing any contagion. Such institutions do not transmit any shocks in response to the receipt of a liquidity shock (i.e. its column in the left half of the shock transmission matrix is zero). We refer to these institutions as liquidity sinks. Institutions that can easily access cash, for example by borrowing on the interbank market (Rochet and Tirole, 1996) or accessing central bank credit (Bagehot, 1873), can also act as liquidity sinks (note, however, that borrowing cash is not an option when the liquidity shock arises because the institution needs to pay off its debts to decrease its leverage to return to its leverage target). If an institution holds insufficient cash but has made short-term loans, it can raise cash by not rolling over these loans. Finally, an institution can liquidate securities; we assume that this is done in descending order of liquidity (the most liquid securities are liquidated first, to minimise price impact). In sum: cash sits at the top of the pecking order, followed by withdrawal of short-term loans, and finally by liquidation of securities in descending order of liquidity.

Following a similar logic to that applying to liquidity sinks, we also define valuation sinks. For our purposes, a valuation sink is an institution without leverage. Because it has no creditors to transmit contagion to (it cannot go bankrupt) and cannot deleverage (it has no debt), it *absorbs* valuation shocks (i.e. its column in the right half of the shock transmission matrix is zero). An example of a valuation sink is a defined contribution pension fund which has no debt to the financial system.

We assume that each leveraged institution has a *leverage ceiling*, which reflects the maximum risk an institution is willing or allowed to take. The ceiling may be set by regulation, or it may be implicitly imposed by haircuts on collateralized loans. Because the haircut requires the collateral value to exceed the value of the loan, the borrower must finance the excess collateral with its own funds. This limits the amount of (collateralized) debt the institution can finance given its equity. If an institution operates sufficiently close to its ceiling that a valuation shock would force it to deleverage, then we say that it is *leverage targeting.*[11] In contrast, if the leverage is sufficiently below the ceiling (e.g. due to a leverage buffer, as proposed in recent regulation [Goodhart, 2013, FSB, 2017]) we say that it is *passively leveraged.* The shocks in our model are sufficiently small not to push passively leveraged institutions towards their ceiling to the point where they have to transition towards becoming leverage targeting.

4.3.3 Contagion Equations

We now derive simple representative formulas for each contagion channel.

- *Funding contagion:* Suppose institution i extends a short-term loan of size S_{ij} to institution j, which is part of its short-term loan portfolio of size S_i. On receiving a liquidity shock x_i^l, assume institution i proportionately reduces the size of its short-term loans to each institution j to absorb the entire shock.

[11] See e.g. Adrian and Shin [2010], Duarte and Eisenbach [2018], Greenwood et al. [2015], Cont and Schaanning [2017], Bookstaber [2017].

This means that the liquidity shock that is transmitted to institution j is $A^{ll}_{ji}x^l_i$, where

$$A^{ll}_{ji} = \frac{S_{ij}}{S_i}. \tag{4.5}$$

- *Overlapping portfolio contagion:* Suppose that institution i holds n_{si} shares of security s, which makes up the top of i's pecking order[12], and that i experiences a liquidity shock x^l_i that causes i to sell $\Delta n_{si} = x^l_i/p_s$ shares, where p_s is the price of security s. Assume a price impact function of the form

$$\frac{\Delta p_s}{p_s} = \mu_s \frac{\Delta n_{si}}{n_s}, \tag{4.6}$$

  where n_s is the total number of shares of security s in circulation, and the price impact factor μ_s is a nondimensional constant of order one that is inversely proportional to the liquidity of security s. Setting $\mu_s = 1$ implies that selling n_s shares drives the price to zero. Under the assumption of linearity, $\mu_s = 1$ is an upper bound because the price cannot be negative. The resulting valuation shock to any institution j that holds n_{sj} shares of security s is $\Delta p_s n_{sj} = \mu_s x^l_i n_{sj}/n_s$, which implies

$$A^{lv}_{ji} = \mu_s \frac{n_{sj}}{n_s}. \tag{4.7}$$

  Note that the diagonal component A^{lv}_{ii} is nonzero. We assume for simplicity that institutions do not short securities, so we always have $n_{sj} \geq 0$.[13]

- *Counterparty risk contagion:* Assume passively leveraged institution i has equity E_i and total debt D_i, so that its leverage is $\lambda_i = D_i/E_i$. When institution i experiences a valuation shock, its probability of default rises and the risk-adjusted value of its debt falls (Bardoscia et al., 2017). Institutions with more equity can withstand larger valuation shocks without becoming insolvent. Therefore, we assume that the fractional drop in the value of the debt is

[12] We assume for simplicity that institution i has a single security at the top of its pecking order. Hence, although the institution may hold positions in various securities, i only sells shares in security s (until the position is exhausted) to raise liquidity (but the model can allow for multiple securities in the top pecking order layer by assuming that i liquidates these securities proportionally to its position in each security).

[13] The framework can accommodate short positions by simply allowing n_{sj} to be negative, but some convenient properties of the matrix would no longer be guaranteed by the Perron Frobenius theorem.

proportional to the fractional loss in equity x_i^v/E_i. If institution i owes debt D_{ij} to institution j, then the valuation shock transmitted to institution j is $\delta_i x_i^v/E_i D_{ij}$, so

$$A_{ji}^{vv} = \delta_i \frac{1}{E_i} D_{ij} = \delta_i \lambda_i \frac{D_{ij}}{D_i}, \tag{4.8}$$

where the risk adjustment factor δ_i is a nondimensional constant of order one. Choosing $\delta_i = 1$ implies that a shock of size E_i (which causes bankruptcy) causes the full value of the debt to be lost and passed onto i's creditors as a valuation shock. Under the assumption of linearity, $\delta_i = 1$ is an upper bound as the loss cannot exceed the value of the debt. D includes short-term as well as long-term debt, so in general $D_{ij} \geq S_{ji}$.

- *Deleveraging contagion.* Suppose leverage targeting institution i maintains a leverage target λ_i. If it receives a valuation shock x_i^v it must pay off debt to return to its target. The amount by which it must reduce debt is $\lambda_i x_i^v$, so

$$A_{ii}^{lv} = \lambda_i. \tag{4.9}$$

We assume that institution i's leverage targeting prevents the institution from transmitting counterparty risk contagion to its creditors. This is because the institution averts the risk associated with increased leverage by paying off its debts to keep its leverage constant.

We want to stress that the parameters in the equations vary over time, and hence the shock transmission matrix is defined with respect to a specific time t. For simplicity we omit the time subscripts to the matrix and its entries.
The four contagion equations are summarized in Table 4.1. This set of contagion mechanisms is not exhaustive; for example, information contagion is not included [Aharony and Swary, 1996, Acharya and Yorulmazer, 2008]). Furthermore, in times of crisis, institutions sometimes hoard liquidity in response to liquidity shocks (Acharya and Skeie, 2011, Heider et al., 2009). Liquidity hoarding can be included in the funding (4.5) and overlapping portfolio (4.7) contagion equations by adding a hoarding term that captures the additional liquidity an institution hoards proportionally to the received liquidity shock. We make the simplifying assumption that liquidity hoarding is absent. We have chosen the four forms we study here because they are

all important, but we restrict ourselves to only four contagion channels for simplicity. Our basic methodology applies to any contagion channels and does not depend on the details of the interaction terms.

According to the Perron-Frobrenius theorem, the largest eigenvalue of the shock transmission matrix is bounded by its smallest and largest column sums. The sum of a column's entries gives the size of the aggregate shock the institution transmits relative to a received liquidity or valuation shock (depending on whether the columns is in the left or right half of the matrix). When no column-sum exceeds one, no institution ever transmits an aggregate shock that exceeds the received shock, so there is no shock amplification and the system is stable. Conversely, when all column-sums exceed one, shocks are always amplified and the system is unstable.

The sum of a column corresponding to an institution's transmission of funding contagion (4.5) is equal to $\sum_j S_{ij}/S_i = 1$ and the sum of a column corresponding to overlapping portfolio contagion (4.7) is given by $\sum_j \mu_s n_{sj}/n_s = \mu_s \leq 1$. Hence, the aggregate shock transmitted in response to a liquidity shock is never amplified (but note that the addition of a liquidity hoarding term could change this). Furthermore, the sum of a counterparty risk contagion column is equal to $\sum_j \delta_i \lambda_i D_{ij}/D_i = \delta_i \lambda_i \leq \lambda_i$ and the sum of a deleveraging column is given by its only non-zero element λ_i. Therefore, under the assumption of no liquidity hoarding, *leverages exceeding one are the only source of shock amplification in the system.* This acts through the counterparty risk and deleveraging channels.

Contagion Mechanism	Contagion Equation	Description
Funding Contagion	$A^{ll}_{ji} = \frac{S_{ij}}{S_i}$	Short-term lending withdrawal
Counterparty Risk Contagion	$A^{vv}_{ji} = \delta_i \lambda_i \frac{D_{ij}}{D_i}$	Probability of default increases due to lower valuations
Overlapping Portfolio Contagion	$A^{lv}_{ji} = \mu_s \frac{n_{sj}}{n_s}$	Price-impact of selling securities
Leverage Targeting Contagion	$A^{vl}_{ii} = \lambda_i$	Delevering requires raising liquidity

Table 4.1: Contagion Equations

4.3.4 Stylized Example

We illustrate the approach outlined above using a simple example of a self-contained financial system that includes all four contagion mechanisms as well as both liquidity and valuation sinks. Consider four institutions, as summarized in Figure 4.3.

- Pension fund h has no debt and a cash surplus, making it both a valuation and a liquidity sink. It makes long-term loans L_{hi}, L_{hj} and L_{hk} to institutions i, j and k and has a position n_{sh} in security s.
- Bank i is passively leveraged. It has a position n_{si} in security s, which sits at the top of its pecking order, and debt $D_{ih} = L_{hi}$ and $D_{ij} = S_{ji} + L_{ji}$ to institutions h and j.
- Bank j targets leverage λ_j. It makes short and long-term loans S_{ji}, L_{ji} and S_{jk}, L_{jk} to institutions i and k and has debt $D_{jh} = L_{hj}$ to institution h. The short-term loans sit at the top of its pecking order.
- Bank k has a cash surplus, making it a liquidity sink, and maintains a leverage target λ_k. It has a position n_{sk} in security s, short-term debt $D_{kj} = S_{jk} + L_{jk}$ and long-term debt $D_{kh} = L_{hk}$.

Balance sheet pension fund h	
Assets	Liabilities
Cash: surplus Position in security s Long-term loans to institutions i, j and k	
	Equity

(a)

Balance sheet bank i	
Assets	Liabilities
Cash: required minimum Position in security s	Debt to h and j
	Equity

(b)

Balance sheet bank j	
Assets	Liabilities
Cash: required minimum Short and long-term loans to i and k	Debt to h
	Equity

(c)

Balance sheet bank k	
Assets	Liabilities
Cash: surplus Position in security s	Debt to h and j
	Equity

(d)

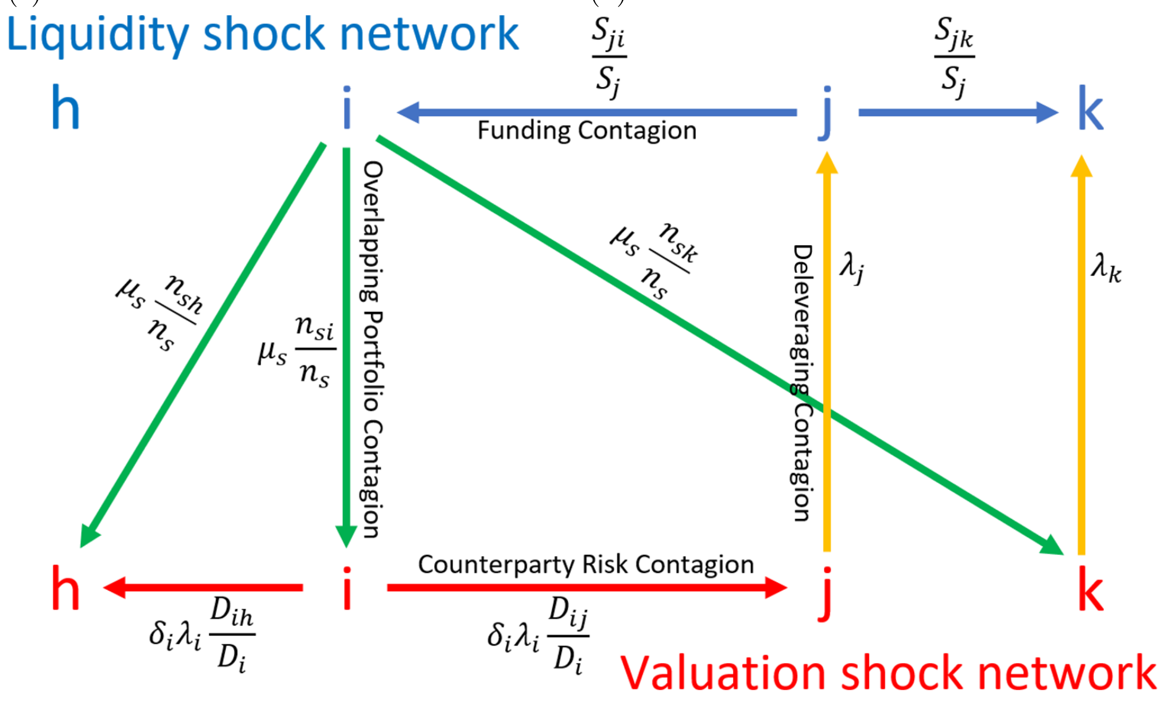

(e)

Figure 4.3: **A stylized example illustrating the interaction of multiple channels of contagion.** Consider four institutions, h, i, j and k, whose balance sheets are given at the top of the figure. Each institution's liquidity is represented by a node in the liquidity shock network, while each institution's solvency is represented by a node in the valuation shocks network. The color of an edge indicates the type of contagion transmitted and the expression next to it represents the size of the interaction. The out-edges of a node in the liquidity shock network give the node's response to liquidity shocks and are dictated by the asset that sits at the top of the institution's pecking order. Similarly, out-edges in the valuation shock network give the response to valuation shocks as given by the institution's leverage strategy: unleveraged (institution h), passively leveraged (institution i), or leverage targeting (institutions j and k).

The shock transmission matrix (and dynamic of the shock vector) of this system is:

$$\begin{array}{c} \\ h^l \\ i^l \\ j^l \\ k^l \\ h^v \\ i^v \\ j^v \\ k^v \end{array}
\begin{array}{c} \begin{array}{cccccccc} h^l & i^l & j^l & k^l & h^v & i^v & j^v & k^v \end{array} \\
\begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & \frac{S_{ji}}{S_j} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & \lambda_j & 0 \\
0 & 0 & \frac{S_{jk}}{S_j} & 0 & 0 & 0 & 0 & \lambda_k \\
0 & \mu_s \frac{n_{sh}}{n_s} & 0 & 0 & 0 & \delta_i \lambda_i \frac{D_{ih}}{D_i} & 0 & 0 \\
0 & \mu_s \frac{n_{si}}{n_s} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & \delta_i \lambda_i \frac{D_{ij}}{D_i} & 0 & 0 \\
0 & \mu_s \frac{n_{sk}}{n_s} & 0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix} \end{array}
\begin{bmatrix} x^l_{t,h} \\ x^l_{t,i} \\ x^l_{t,j} \\ x^l_{t,k} \\ x^v_{t,h} \\ x^v_{t,i} \\ x^v_{t,j} \\ x^v_{t,k} \end{bmatrix}
=
\begin{bmatrix} x^l_{t+1,h} \\ x^l_{t+1,i} \\ x^l_{t+1,j} \\ x^l_{t+1,k} \\ x^v_{t+1,h} \\ x^v_{t+1,i} \\ x^v_{t+1,j} \\ x^v_{t+1,k} \end{bmatrix}, \tag{4.10}$$

where $S_j = S_{ji} + S_{jk}$, $D_i = D_{ih} + D_{ij}$ and $n_s = n_{sh} + n_{si} + n_{sk}$.
To simplify the discussion we set the price-impact and risk adjustment factors to their upper bounds $\mu_s = \delta_i = 1$. For this system, the largest eigenvalue of the shock transmission matrix is equal to

$$\nu = \left(\frac{S_{ji} n_{si} \lambda_i D_{ij} \lambda_j}{S_j n_s D_i} \right)^{1/4}, \tag{4.11}$$

and so the largest eigenvalue in this example is the product of each of the four contagion mechanisms, i.e.

$$\nu^4 = \frac{S_{ji}}{S_j} \times \frac{n_{si}}{n_s} \times \lambda_i \frac{D_{ij}}{D_i} \times \lambda_j.$$

The factors S_{ji}/S_j, n_{si}/n_s and D_{ij}/D_i are all less than or equal to one, and so exert a stabilizing force competing against the potentially destabilizing forces of the leverages λ_i and λ_j. If any of the four channels of contagion is removed the largest eigenvalue becomes zero and the system becomes unconditionally stable. The possibility for instability is caused by the interaction of all four channels, so analyzing each channel in isolation, as often done, gives a misleading result.
Consider some plausible numbers: if $S_{ji}/S_j = D_{ij}/D_i = 1/3$, $n_{si}/n_s = 1/4$ and $\lambda_i = \lambda_j = 6$, then $\nu = 1$ and the system is at its margin of stability. For any system, we define the *critical leverage* $\hat{\lambda}$ as the maximum leverage that all *leveraged* institutions can attain simultaneously without rendering the system unstable (i.e. $\nu \leq 1$, where all unleveraged institutions are assumed to remain so). For the system outlined in the stylized example, the critical leverage is thus given by $\hat{\lambda} = 6$.

The Stabilizing Role of Sinks. Sinks play an essential role in stabilizing the financial system. All else equal, if more shocks are transmitted to sinks, then less shocks are transmitted to potentially destabilizing institutions.[14] The ability of a sink to absorb shocks increases when it holds more short-term lending, more securities and more of other institutions' debt. This can be seen for the example by expanding the denominators in Eq. (4.11),

$$\nu = \left(\frac{S_{ji}}{S_{ji} + S_{jk}} \times \frac{n_{si}}{n_{sh} + n_{si} + n_{sk}} \times \lambda_i \frac{D_{ij}}{D_{ij} + D_{ih}} \times \lambda_j \right)^{1/4}. \tag{4.12}$$

The sinks' short-term debt S_{jk}, securities holdings n_{sh} and n_{sk}, and lending D_{ih} all appear only in the denominator and therefore exert a stabilizing force on the system. Hence, *sinks stabilize the system by absorbing shocks that would otherwise have been transmitted to other institutions.*

4.3.5 Application to Stress Testing

This framework can be calibrated using granular data to accurately represent the microstructure of the financial system and it can be used to study any channels of contagion (he description can be made at any level of granularity, down to individual contracts). Because the shock transmission matrix' largest eigenvalue quantifies stability independently from any specific stress scenario, it provides an objective, robust measure of stability that allows for comparison across time, jurisdictions, policy interventions, and so on.

The eigenvectors associated with the largest eigenvalue also provide useful diagnostic information.[15] The right eigenvector provides a measure of institutions' in-degree centralities that takes the whole system into account and the left eigenvector provides a measure of institutions' out-degree centralities that takes the whole system into account (Newman, 2018). The larger entries of the right eigenvector flag the institutions that are likely to receive the largest shocks and those of the left eigenvector flag the institutions that play the biggest role in transmitting shocks. This can potentially help guide policymakers in identifying systemically important institutions and designing stress tests and interventions.

[14] In any system, the eigenvalues are determined by the shock transmission between strongly connected nodes. Sinks are by definition not strongly connected, so the more shocks transmission to sinks, the less shock transmission between strongly connected nodes and the lower the largest eigenvalue.

[15] The Perron Frobenius theorem guarantees that the right eigenvector is non-negative and, assuming that the network has a single strongly connected component, that it is unique (i.e. the coresponding eigenspace of the largest eigenvalue ν is one-dimensional). The same is true for the left eigenvector.

Scenario-Dependent Stability. Although the framework here assumes infinitesimal shocks, it is also useful for understanding large shocks. In reality institutions do not respond to a shock instantaneously, but rather take a series of actions that are initially close to the dynamics captured by our framework but diverge from these over time, so as the financial system responds to a large shock its stability changes. As the shock plays out, institutions may exhaust the top layers of their pecking orders, market liquidities may fall, and risk adjustment factors may rise. By investigating how this affects (linear) stability, we may gain insight into how large shocks affect stability. The properties of the financial system may also change due to central bank or government intervention in response to the shock, or because of adaptive processes that take place as part of the ongoing evolution of the financial system (Farmer et al., 2021a).

The evolution of the stability of the financial system in response to a large shock embodies the *scenario-dependent* component of the system's shock-dynamics, which is not captured in its linear stability. As the framework can be calibrated to any state of the system, by continuously re-calibrating the shock transmission matrix as the system evolves in response to a shock, we may understand how the instantaneous stability changes, track it over time, and distinguish between the scenario-dependent and independent components of the system's stability. This can be done for both the empirically observed or forecast (by means of some simulation model) evolution of real financial systems.[16] By investigating the sensitivity of a system's stability to specific stress scenarios using the method developed here, policymakers may gain insight into the factors that generate financial instability in any given scenario.

The effect of the liquidity pecking order in particular warrants further investigation, because the stability of financial systems strongly depends on which assets are at the top of institutions' pecking orders. For example, when each institution has a single-layer pecking order (i.e. it liquidates a vertical slice across all its liquid assets), the shock transmission matrix remains constant as long as institutions do not default or change their leverage strategies (under reasonable assumptions about μ_s and δ_i). Liquidation-cost minimizing pecking orders, on the other hand, may be highly sensitive to the choice of stress scenario, because liquidities may change and layers of the pecking order may be exhausted. Therefore, investigating what

[16]Note that the purpose of our framework is to identify instabilities as they emerge, rather than forecast the nature and magnitude of the crisis that may ensue. The framework may provide insight into the simulated evolution of a financial system, but the simulation itself would require another model (or an extension of the one developed here).

pecking orders institutions use during crises may yield valuable insights into the predictability of financial stability.

4.4 Stability Overestimation due to Ignoring Interactions

In the previous sections we derived the shock transmission matrix and explained how it may be used to assess the stability of financial systems with interacting contagion channels. In this section, to underscore the importance of using the shock transmission matrix to complement existing stress tests we demonstrate the dangers of ignoring the interactions between contagion channels. In the stylized example in section 4.3.4, we showed that ignoring any of the four contagion channels may overestimate stability. In this section we demonstrate this in a more general setting. Because most contagion literature studies the counterparty default (risk) channel alone[17], we focus on the overestimation due to only considering counterparty risk contagion and ignoring all other channels. This *stability of pure counterparty risk contagion* is given by the largest eigenvalue of the counterparty risk component of the shock transmission matrix, i.e its bottom-right quadrant. We first discuss two extremes for which the overestimation follows intuitively and then consider intermediate cases.

In the case where all institutions in the financial system are passively leveraged or no institution has tradeable securities at the top of their pecking order, the system's shock transmission matrix is block-triangular. Due to the properties of block-matrix determinants, the largest eigenvalue of this system is given by the largest eigenvalue of the diagonal quadrants, which correspond to counterparty risk and funding contagion. Such a system is stable when pure counterparty risk contagion is stable (as funding contagion is never amplified under the assumption of no liquidity hoarding) so that there is little potential for overestimating stability.

In the opposite situation where all institutions are leverage targeting, the counterparty risk quadrant is zero. Hence, considering pure counterparty risk contagion would lead to the conclusion that the system is unconditionally stable, regardless of the true stability. To understand how ignoring interactions between contagion channels leads to an overestimation of stability in systems with both passively leveraged and leverage targeting institutions, we must understand how network composition

[17]See e.g. Eisenberg and Noe, 2001, Furfine, 2003, Gai and Kapadia, 2010, Battiston et al., 2012, Elliott et al., 2014, Acemoglu et al., 2015, Bardoscia et al., 2015, 2017.

affects stability. As a tractable example, in the sections that follow we apply the framework developed in section 4.3 to randomly generated financial systems.

4.4.1 Application to Randomly Generated Financial Systems

To gain insight into the dynamics of interacting contagion channels we study randomly generated financial systems. To do this we populate the balance sheets of N institutions with randomly generated securities and loans. We do this in such a way that the types of institutions introduced in section 4.3.2 have the following proportions:

1. A fraction ϕ_l of institutions have sufficient cash to absorb shocks. We call ϕ_l the *fraction of liquidity sinks.*
2. A fraction ϕ_v of institutions have no leverage. We call ϕ_v the *fraction of valuation sinks.*
3. A fraction F of institutions provide short-term loans. We call F the *fraction of short-term lenders.*
4. A fraction Λ of *leveraged* institutions are leverage targeting. We call Λ the *fraction of leverage targeters.*

These proportions constrain the random assignment of loans. For each security s out of a possible number N^w of distinct securities, we divide the total number of outstanding shares n_s into N^s blocks of n_s/N^s shares and assign each block to a randomly chosen institution. We do this with uniform probability and with replacement. Similarly, each institution makes N^d loans, each to a randomly chosen *leveraged* institution; for simplicity, all loans an institution *receives* are set equal in size. Any institution that was designated as leveraged but ended up not receiving any loans is allocated a single loan from a randomly chosen institution.
For any institution i, let N_i^s denote the number of blocks of security s received, N_i^d the total number of loans received, and N_{ji}^d the number of loans received from institution j. The leverages of the $N^v = (1 - \phi_v)N$ leveraged institutions are set to the critical leverage $\lambda_i = \hat{\lambda}$, which fixes their debts relative to their equities. Once we choose the N^w distinct securities' market capitalizations $C_s = p_s n_s$, the requirement that any institution's assets (LHS) must equal the sum of its equity and debt

(RHS) provides N constraints that determine the N institutions' equity E_i,

$$\sum_{s=1}^{N^w} C_s \frac{N_i^s}{N^s} + \sum_{j=1}^{N^v} D_j \frac{N_{ij}^d}{N_j^d} = E_i \left(\lambda_i + 1\right). \qquad (4.13)$$

This allows us to generate a random financial system with any desired values of the parameters ϕ_l, ϕ_v, F and Λ.

Financial systems tend to have sparse, heterogenous topologies (Boss et al., 2004, Cont et al., 2013) that can frequently be characterized as core-periphery structures (Craig and Von Peter, 2014, Fricke and Lux, 2015). Similarly, our randomly generated systems here include many sources of heterogeneity and have a core-periphery structure: Since an institution can either be a liquidity sink or not be a liquidity sink, provide short-term loans or not provide short term loans, and be unleveraged, passively leveraged or have a leverage target, this implies that there are $2 \times 2 \times 3 = 12$ different *types* of institutions. The short-term lending network that results from our method of random construction has a core of institutions that both provide and receive short-term loans. There are also three distinct peripheries - one of institutions that only provide short-term loans, one of institutions that only receive short-term loans, and one of institutions that do not partake in the short-term lending network at all. The long-term lending network has a similar topology but with different institutions at its core and peripheries.

Furthermore, because loans and securities are chosen *with replacement*, institutions can receive multiple loans from the same institution and hold multiple shares in the same security. Consequently, the weights of the edges vary across institutions in all networks. Finally, because the institutions' assets are randomly determined, the endogenously-determined balance sheet sizes vary across institutions. By varying the system parameters ϕ_l, ϕ_v, F, Λ, N, N^d, and N^s, we can control the level of heterogeneity present in the system.

4.4.2 Mean-Field Approximation

In the limit where $N, N^d/N, N^s/N \to \infty$, the randomly generated financial systems reduce to a mean-field model. In fact, as we show below, the mean field model remains a reasonable approximation for much smaller, sparser systems as well. Rather than studying the stability of the generated systems explicitly we focus on the mean-field approximation, which is more insightful because it provides an analytic stability condition.

In Supplementary Materials B.3, we explain that the dynamics of this mean-field model are uniquely defined by the transmission of the aggregate liquidity shock $x_t^l = \sum_i x_{t,i}^l$ and aggregate valuation shock $x_t^v = \sum_i x_{t,i}^v$. This allows us to reduce the system's full shock transmission matrix to a 2×2 matrix that describes the dynamics of the aggregate shocks,

$$\begin{bmatrix} x_{t+1}^l \\ x_{t+1}^v \end{bmatrix} = \left[\begin{array}{c||c} (1-\phi_l)F & \lambda(1-\phi_l)\Lambda \\ \hline\hline \mu(1-\phi_v)(1-F) & \delta\lambda(1-\phi_v)(1-\Lambda) \end{array}\right] \begin{bmatrix} x_t^l \\ x_t^v \end{bmatrix}, \tag{4.14}$$

where μ is the price impact factor of the most liquid security (in which all institutions have a position when $N^s/N \to \infty$), and we have set $\lambda_i = \lambda$ for all leveraged institutions and $\delta_i = \delta$ for passively leveraged institutions.

We compute the characteristic equation of the matrix in Eq. (4.14) and solve for its largest eigenvalue $\nu = 1$ to find the *mean-field critical leverage*,

$$\hat{\lambda} = \frac{1 - (1-\phi_l)\,F}{\mu\Lambda\,(1-\phi_l)\,(1-F) + \delta\,(1-\Lambda)\,(1-(1-\phi_l)\,F)}(1-\phi_v)^{-1}. \tag{4.15}$$

Eq. (4.15) demonstrates that financial stability is the result of a balance between the destabilizing force of leverage and the stabilizing force of sinks (note that the mean-field critical leverage (4.15) is an increasing function of both the liquidity sinks fraction ϕ_l and valuation sinks fraction ϕ_v). We first discuss the accuracy of the mean-field model before we use it to understand how ignoring interactions between contagion channels overestimates financial stability.

Accuracy of the Mean-Field Model. The mean-field model was derived in the limit of a dense network with an infinite number of institutions. In Figure 4.4, we compare the mean-field critical leverage predicted by Eq. (4.15) to that of randomly generated systems with varying sizes and densities and show that the mean-field critical leverage is a good approximation not only for large, dense systems, but also for fairly small, sparse systems.

To choose plausible estimates of the system parameters we do a rough calibration to the European financial system. To estimate the parameters ϕ_v, Λ, F and ϕ_l, we calculate the fraction of financial assets in the Eurosystem held by various sectors, as shown at the top of Figure 4.4.[18] We include $N^w = 10$ distinct securities whose

[18] The aggregate sector assets are listed in Table 4.4a (rounded to one decimal) and the parameter calculations are listed in Table 4.4b (rounded to a multiple of $1/4^{th}$ or $1/5^{th}$ so the number of each type is an integer for any multiple of $N = 10$). We assume that pension funds and insurance corporations are unleveraged, monetary financial institutions (which are mostly banks) are leverage targeting, and investment funds and financial vehicle corporations are passively leveraged.

markets caps are calibrated to the ten largest stocks on the Euronext exchange.[19] The Eurosystem's 119 most *significant institutions* as designated by the ECB have leverages ranging between 10 and 20 (ECB, 2019). Given that the Eurosystem appears to be stable, we choose values of $\delta = \mu = 0.1$, which gives a critical leverage of $\hat{\lambda} = 35$. We also think these are reasonable values for other reasons.[20]

Monetary financial institutions are assumed to be the only providers of short-term loans. Finally, we assume favorable market liquidity conditions such that all institutions but investment funds and financial vehicle corporations are liquidity sinks. Although these assumptions imply correlations between institutions' types, the are assigned independently from one another. Note that these are rough parameter estimates. Accurate calibration requires further (empirical) investigation.

[19] Source: `https://www.statista.com/statistics/546298/euronext-market-capitalization-leading-companies/`; accessed January 22^{nd}, 2019.

[20] The fact that a large fraction of the value of a loan can usually be recovered suggests that δ is substantially below one. Similarly, since institutions liquidate assets in order of liquidity, the price impact factor μ_s of any security at the top of institutions' pecking orders is likely to be substantially less than one.

Sector	Assets
Pension Funds (PF)	2.6^a
Insurance Corporations (IC)	8.0^b
Monetary Financial Insts. (MFI)	30.6^c
Investment Funds (IF)	12.4^d
Financial Vehicle Corps. (FVC)	1.9^e

(a)

Fraction	Calculation	Value
ϕ_v	$\frac{PF+IC}{PF+IC+MFI+IF+FVC}$	0.2
F	$\frac{MFI}{PF+IC+MFI+IF+FVC}$	0.5
ϕ_l	$\frac{PF+IC+MFI}{PF+IC+MFI+IF+FVC}$	0.75
Λ	$\frac{MFI}{MFIs+IFs+FVCs}$	0.75

(b)

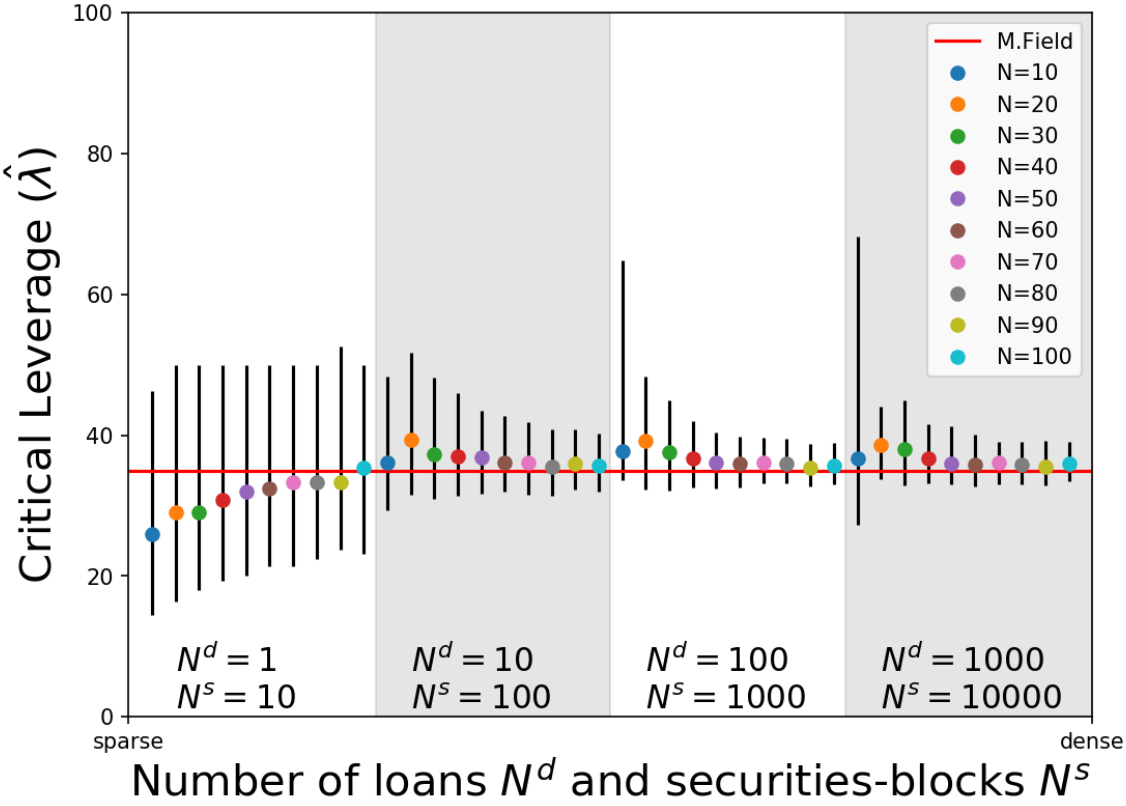

(c)

Figure 4.4: **Comparison of the mean field model to randomly generated financial systems.** The plot compares the critical leverage (red line) predicted by the mean field model, Eq. (4.15), to the critical leverages of randomly generated financial systems. For various combinations of the number of loans N^d, the number of securities-blocks N^s and the number of institutions N, we generate 500 random systems and plot the median (colored dot) and the 15^{th} to 85^{th} percentile interval (black bars) of the distribution of critical leverages. The simulated values converge to the predictions of the mean-field model as N^d, N^s, N increase.

[a] Source: `http://sdw.ecb.europa.eu/reports.do?node=1000005664`; accessed November 3^{rd} 2018

[b] Source: `http://sdw.ecb.europa.eu/reports.do?node=1000005659`; accessed November 3^{rd} 2018

[c] Source: `http://sdw.ecb.europa.eu/reports.do?node=1000005718`; accessed November 3^{rd} 2018

[d] Source: `http://sdw.ecb.europa.eu/reports.do?node=1000003516`; accessed November 3^{rd} 2018

[e] Source: `http://sdw.ecb.europa.eu/reports.do?node=1000003621`; accessed November 3^{rd} 2018

We explore different combinations of the parameters N, N^d and N^s, generating 500 realizations for each set of parameter values. In Figure 4.4 we plot the median critical leverage as well as the 15^{th} to 85^{th} percentile intervals against the prediction of Eq. (4.15). [21] The figure shows that the mean-field model gives a good approximation for systems with at least $N \geq 30$ institutions and at least $N_d \geq 10$ loans made by each institution.

Although this is not a large effect, note that the smallest, sparsest systems in Figure 4.4 are the least stable. This is in contrast to Bardoscia et al. [2017], who find the sparsest systems are most stable. Our model approximately reduces to theirs in the limit where all institutions are passively leveraged ($\Lambda \to 0$). However, Bardoscia et al. [2017] require the sparsest systems to be acyclic, which is the most stable configuration, as the largest eigenvalue is zero by definition. We do not impose this requirement, which is why we get the opposite result.

Our model does not address important features of the financial system, such as hedging with derivatives or short positions. Figure B.2 in the appendix shows that introducing additional sources of heterogeneity increases the variation in critical leverage. Nonetheless, the mean field model does a good job of qualitatively capturing some of the key features of financial stability. We want to stress, however, that when fine-grained data is available, it is far preferable to use the full model developed in section 4.3, which uses weaker assumptions and contains fewer approximations.

4.4.3 The Misclassification Region

We now demonstrate that stability is almost always overestimated, and sometimes dramatically so, when ignoring the interactions between contagion channels. For simplicity, because this is only a qualitative demonstration, we set the price-impact and risk-adjustment factors in the mean-field critical leverage (4.15) equal to their upper bounds $\mu = \delta = 1$. This simplifies Eq. (4.15) to

$$\hat{\lambda} = \left(\frac{1 - (1 - \phi_l)F}{1 - (1 - \phi_l)F - \phi_l \Lambda} \right) (1 - \phi_v)^{-1}. \tag{4.16}$$

Eq. (4.16) is a product of two terms; the first captures the intensity of the feedback loop between solvency and liquidity, and the second captures the way in which valuation sinks counterbalance the destabilizing force of leverage.

[21] Note that the percentile bars visualize the width of the observed distributions. They are not error bars – these are negligible given the large number of samples and are therefore not plotted.

The stability of pure counterparty risk is determined by the counterparty risk quadrant alone. We find the critical leverage of pure counterparty risk contagion in the mean-field model by solving for the leverage for which the counterparty risk quadrant in Eq. (4.14) is equal to one,

$$\hat{\lambda} = \frac{1}{1-\Lambda}(1-\phi_v)^{-1}. \tag{4.17}$$

This is equivalent to setting ϕ_l or F in Eq. (4.16) equal to one (so there is no feedback loop between solvency and liquidity).

The counterparty risk critical leverage (4.17) may be shown to severely overestimate the true mean-field critical leverage when the interaction of other contagion channels is taken into account. To take a simple case, when $\phi_l = 0$, i.e. when there are no liquidity sinks, the true critical leverage equals

$$\hat{\lambda} = (1-\phi_v)^{-1}. \tag{4.18}$$

In Figure 4.5 we plot the counterparty risk critical leverage (4.17) and the true critical leverage (4.18). As the fraction of leveraged institutions increases, the discrepancy between the counterparty risk critical leverage and true critical leverage becomes arbitrily large, and hence so does the overestimation of stability due to ignoring the interactions between contagion channels. To show that this happens for any value of $\phi_l < 1$, in Figure B.1 of the appendix we plot the critical leverage for various values of ϕ_l. For $\phi_l = .75$, $F = .5$ and $\Lambda = .75$, for example, we find that the counterparty risk critical leverage overestimates the true critical leverage by 45%, but this soars to 300% when $\phi_l \to 0$, as may be the case when liquidity dries up during financial crises.

Detering et al. [2021] also show in a scenario-independent setting (which also only depends on aggregate system parameters) that stability is overestimated when ignoring the interaction between contagion channels. However, they only include counterparty default and overlapping portfolio contagion and do not capture liquidity in their model. By capturing the solvency-liquidity nexus (and all its contagion channels) in its entirety, we demonstrate that this overestimation is determined by the intensity of the feedback loop between solvency and liquidity.[22] The manifestation of this overestimation in all but unrealistic scenarios makes clear that taking

[22] Comparison of equations 4.17 and 4.18 shows that the overestimation is caused by the first term on the right hand side of Eq. 4.16, as the second term appears identically in all three equations.

account of interacting contagion channels, as our approach outlined in Section 4.3 proposes, is critical when evaluating financial stability.

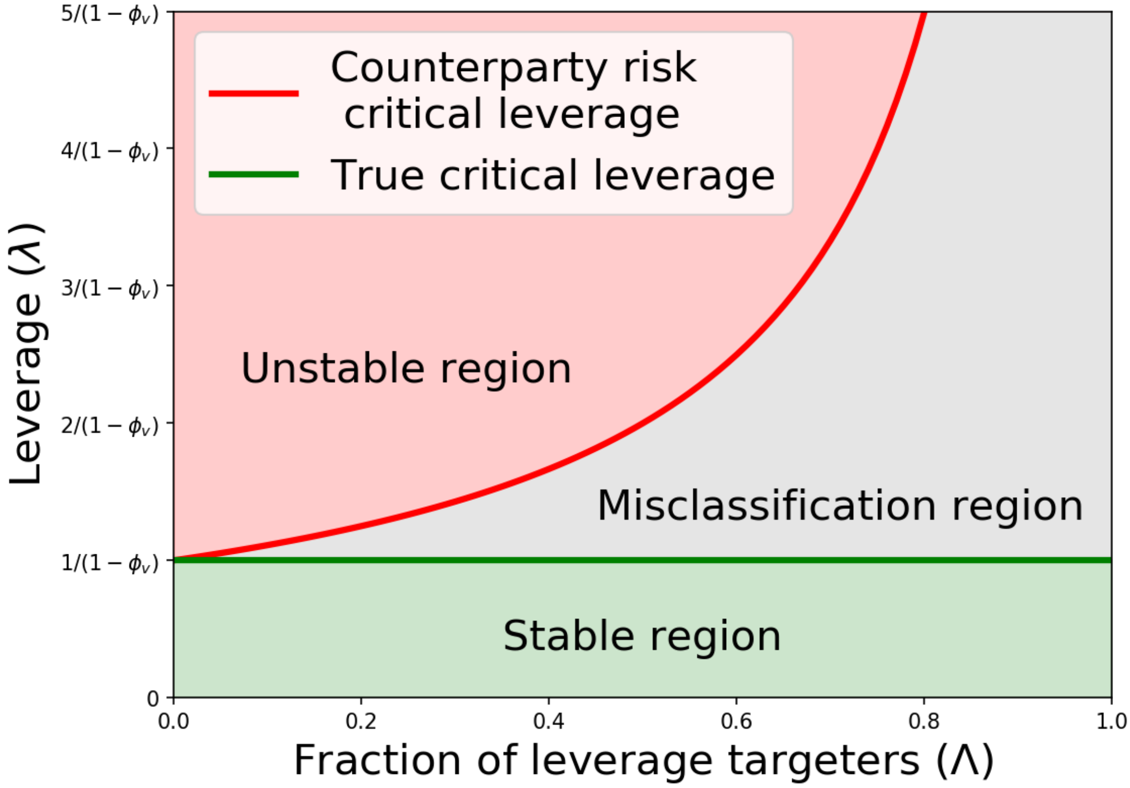

Figure 4.5: **Overestimation of financial stability by omitting the interaction of channels of contagion.** The green line plots the critical leverage of the mean-field model for the fraction of liquidity sinks $\phi_l = 0$, and the red curve plots the counterparty risk critical leverage as a function of the fraction of leverage targeting institutions Λ. The leverages are expressed in units of $(1 - \phi_v)^{-1}$ to reduce the number of free parameters. The region between the true critical leverage and the counterparty risk critical leverage is the *misclassification region* and consists of destabilizing leverages that *seem* stable when considering only pure counterparty risk contagion.

4.5 Concluding Remarks

Financial instability is caused by the endogenous amplification of shocks.[23] We are the first to introduce a scenario-independent measure of the stability of the solvency-liquidity nexus that takes into account the interactions of an arbitrary number of financial contagion channels. By describing the interactions of liquidity and valuation shocks, our method captures the most important contagion mecha-

[23]See e.g. Danielsson and Shin [2003], BIS [2009], Krishnamurthy [2010], Acemoglu et al. [2012], Anderson et al. [2018].

nisms and their interactions in a duplex network consisting of a liquidity and a valuation shock layer. The largest eigenvalue of the system provides a robust measure of the system's stability that is complementary to the insights provided by existing methods because it does not rely on subjectively imposed stress scenarios. Furthermore, the associated eigenvectors provide detailed insights that are valuable for important policy considerations, such as the identification of the most systemically important institutions.

With appropriate microdata this method can be calibrated against real financial systems. To do this it is necessary to estimate which institutions absorb liquidity shocks, to identify their leverage strategies and pecking orders, and measure the price-impact and risk-adjustment factors μ_s and δ_i. While this is not a trivial task, it is feasible with the right data. We hope that our model will help provide an incentive for central banks to collect this data. The analysis presented here relies on the assumption that shocks are small enough that it is only necessary to consider the top level of the pecking order. However, as we have outlined, the method may also be used to monitor the stability of a financial system as it evolves in response to larger shocks. This will be investigated in a follow-up paper.

The framework presented here has the advantage over black-box simulations, such as Kok and Montagna [2016], Cont and Schaanning [2017] or Farmer et al. [2020], that it provides insight into the mechanisms that cause contagion. In fact, it can be used in conjunction with such simulations to provide a deeper understanding of their results. Furthermore, our analysis here complements asymptotic graph techniques (see e.g. Gai and Kapadia, 2010, Amini et al., 2016, Detering et al., 2021), which study the final state to which the system converges in response to a shock. An important advantage of the instantaneous stability quantified by the shock transmission matrix is that it does not suffer from the compounding inaccuracies that any iterative model is inevitably exposed to.

We derive an analytic stability criterion in the limit of a large number of institutions which demonstrates that financial stability results from the balance between stabilizing and destabilizing forces. Although the stability criterion is simple, with only a few parameters, it is powerful enough to derive a wide range of insights about the stability of financial systems. It shows that the shock-amplifying forces of the levels of leverage common in real financial systems must be offset by damping to maintain stability. Understanding the conditions under which institutions absorb financial shocks is crucial to policymakers. Despite this, the absorption of shocks by

sinks and the damping of financial shocks have received little attention in the literature so far. The only previous work that we are aware of that stresses this point is Aymanns et al., 2016), who study the balance between shock-amplification due to deleveraging banks and shock-damping by fundamental value investors. We study the interaction between shock-amplifying and shock-dampening forces in a much more general network setting. The stability criterion that we derive demonstrates the fundamental importance of the balance between stabilizing and destabilizing forces, highlighting that this interplay deserves further investigation.
Building on the work of others who have observed that interactions between contagion channels amplify instabilities in particular settings (Caccioli et al., 2013, Kok and Montagna, 2016, Poledna et al., 2015, Cont et al., 2020,?, Detering et al., 2021), the analytic stability criterion that we develop here elucidates the mechanism responsible for this amplification in general. We show that a feedback loop between liquidity and valuation shocks always exists and that when the interactions between contagion channels are ignored, this feedback loop is overlooked and stability is overestimated, sometimes dramatically so. Because most studies focus on a single type of shocks (see e.g. Eisenberg and Noe, 2001, Caccioli et al., 2014, Bardoscia et al., 2017), financial instabilities may be structurally underestimated. Hence, comprehensive measures of the financial stability implications of interacting contagion channels like the framework developed here are highly complementary to other existing methods.

Chapter 5

Complex Liquidity Spirals

Based on the eponymous paper by Garbrand Wiersema, Esti Kemp and J. Doyne Farmer. This work has been presented at the 2022 Oxford-ETH Workshop on Interconnectedness and Financial Stability at ETH Zurich.

5.1 Summary

We introduce a novel method for studying liquidity spirals and identify these before stock prices plummet and funding markets lock up. We show that liquidity spirals may be underestimated or completely overlooked when interactions between contagion channels are ignored, and 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 show that, depending on the market conditions, a liquidity spirals is sometimes caused by the sectors' collective dynamics, but at other times by one sector's individual impact. We also find that liquidity spirals are exacerbated when the liquidity of institutions worsens, and that central bank-provided liquidity can greatly dampen liquidity spirals. The approach developed here can be used to formulate interventions that specifically target the sector(s) causing the liquidity spiral.

5.2 Introduction

The progressive worsening of market and funding liquidity due to positive feedback loops in the financial system is referred to as a liquidity spiral, and poses a significant risk to financial stability by causing or exacerbate crises such as the Great

Financial Crisis of 2008 (Brunnermeier and Pedersen, 2009). 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 (see e.g. Allen and Gale, 2000, Eisenberg and Noe, 2001, Gorton and Metrick, 2012) and the interactions between different contagion channels have been observed to severely amplify instabilities.[1] 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., 2021). We refer to liquidity spirals that consist of various interacting contagion channels and/or multiple types of institutions as *complex liquidity spirals*.

We identify liquidity spirals before they progressively depress market and funding liquidity using a *shock transmission matrix* (Wiersema et al., 2019). The matrix captures the stability of various interacting contagion channels without relying on any specific, subjective stress scenarios. When the largest eigenvalue of the matrix exceeds one, market and funding liquidity progressively worsen and a liquidity spiral emerges. We find that complex liquidity spirals may be severely underestimated or even completely overlooked when interactions between different types of institutions or contagion channels are ignored, and that the intensity of the spiral greatly depends on which assets institutions choose to liquidate in response to a liquidity shock. 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. Gai and Kapadia [2010] find a similar phenomenon for certain topologies of financial networks. The identification of the robust-yet-fragile tendency of financial systems across multiple dimensions underscores the importance of stability measures that assess a system's resilience to a wide range of shocks, such as the eigenvalue-based approach developed here.

To demonstrate our method, we apply it to a highly granular data set on the South African financial system. The South African financial system consists of a core of five large banks and a periphery of smaller banks and a large number of investment funds. By monitoring both the stability of the individual sectors as well as the combined system, we find that, depending on the market conditions, a liquidity spiral may emerge either as the result of the banking and fund sectors' collective dynamics or due to an individual sector's impact. Interventions that target the banking sector when the investment fund sector is the main cause of the spiral (and

[1] See e.g. Caccioli et al. [2013], Poledna et al. [2015], Kok and Montagna [2016], Wiersema et al. [2019], Cont et al. [2020], Detering et al. [2021].

vice versa) have little effect. In particular, we find that liquidity spirals are exacerbated when the liquidity of institutions falls and that central bank-provided liquidity can greatly dampen the liquidity spiral, but only when provided to the right sector. The approach developed here can be used to determine which sector is the main driver of instability and to formulate interventions that specifically target the sector that is causing the instability.

5.2.1 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] and Hałaj [2018]. 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.[2] Our approach of identifying liquidity spirals based on the contagion dynamics' largest eigenvalue rather than based on a specific, subjectively determined stress scenario further strengthens the comprehensiveness of our analysis (Borio et al., 2014, Wiersema et al., 2019).
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 liquidity spirals 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., 2021). Using a similar data set as we do here, Wiersema et al. [2021] study (counterparty) exposures in the South African financial system and how they are affected by institutions' solvency. The analysis presented in this paper broadens the understanding of the stability of the South African financial system by focusing on liquidity crises.

[2] Caccioli et al. [2013], Kok and Montagna [2016], Poledna et al. [2015], Wiersema et al. [2019], Detering et al. [2021].

5.2.2 Structure

The remainder of this paper is organized as follows: Section 5.3 presents our method for identifying complex liquidity spirals and the insights it offers. In section 5.4, 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 5.5 concludes by discussing the implications and limitations of our results, and provides avenues for further research.

5.3 Identifying Complex Liquidity Spirals

We use the framework developed in Wiersema et al. [2019] to study the conditions under which complex 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 identify a liquidity spiral before a liquidity crisis develops.

5.3.1 The Solvency-Liquidity Nexus

A financial system's contagion dynamics are driven by the *Solvency-Liquidity Nexus*. The culmination of a severe financial crisis is usually the default of one or more institutions (Brunnermeier, 2009, Roukny et al., 2013), where a default can be caused by *insolvency* or *illiquidity*. Insolvency occurs when asset values drop to the point where equity becomes negative – that is, when the value of an institution's liabilities exceeds that of its assets (Amini et al., 2016). Default due to illiquidity, on the other hand, occurs when an institution is unable to meet its payment obligations (Cont and Schaanning, 2017). Insolvency and liquidity can be related, but are analytically distinct: an institution can default due to a liquidity shock even when it is solvent, and vice versa. During financial crises, liquidity tends to be the more direct threat; an institution may survive temporary insolvency by maintaining liquidity and regaining its solvency at a later date. In normal economic times, a solvent institution is expected to borrow to avert a liquidity shortage. In times of economic crisis, however, this may not be possible because lending markets malfunction due to uncertainty about asset values, escalating collateral requirements, liquidity hoarding and capital flight, etc. (Rochet and Vives, 2004, Gorton and Metrick, 2012).

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 of liquid assets and a valuation shock as a drop in the value of an institution's assets.[3] 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 contagion channels may convert valuation shocks to liquidity shocks.

5.3.2 Contagion Channels

We now demonstrate how to capture contagion channels 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. [2019], 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[4]: If institution i suffers a liquidity shock it may be forced to sell securities to raise liquidity. This depresses their price. If institution j also has a position in these securities it experiences a valuation shock. Hence, *overlapping portfolio contagion converts liquidity shocks to valuation shocks.* By increasing the demand for liquidity on trading markets, overlapping portfolio contagion depresses market liquidity.

[3]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.

[4]See e.g. Adrian and Shin [2010], Caccioli et al. [2013, 2014, 2015], Duarte and Eisenbach [2018], Cont and Schaanning [2017, 2019].

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[5]: 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 (Wiersema et al., 2021): 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. Hence, *shareholder contagion propagates valuation shocks.*

Share Redemption Contagion

An institution which issues shares that are redeemable on a short-term basis (typically daily) is at risk of share 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 (Cont and Wagalath, 2013, Wiersema et al., 2021). 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, *share redemption contagion converts valuation shocks to liquidity shocks.*

Deleveraging Contagion

Deleveraging contagion takes place when an institution uses borrowed funds to purchase assets.[6] 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

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

[6] 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].

leverage target, an action that drains the institution's liquidity.[7] 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).

5.3.3 The Shock Transmission Matrix

We show how to characterize the collective stability of the five described contagion channels without relying on any specific, subjective stress scenarios. The interactions of these contagion channels can be captured in a *shock transmission matrix* (Wiersema et al., 2019). Assuming discrete dynamics, the shock transmission matrix A_t is defined with respect to a specific time t as contagion equations may change along with institutions' balance sheets. Furthermore, let $x^l_{t,i}$ denote the liquidity shock suffered by institution i at time t such that the N-dimensional vector $\vec{x}^l_t$ gives the liquidity shocks to all institutions (where N is the number of financial institutions). Similarly, $x^v_{i,t}$ denotes the valuation shock to institution i at time t and $\vec{x}^v_t$ the N-dimensional vector of valuation shocks to all institutions. The combined shock vector $\vec{x}_t$ of length $2N$ is

$$\vec{x}_t = \begin{bmatrix} \vec{x}^l_t \\ \vec{x}^v_t \end{bmatrix}. \tag{5.1}$$

The shock transmission matrix A_t is the $2N \times 2N$ matrix that acts on the shock vector $\vec{x}_t$ according to

$$\vec{x}_{t+1} = A_t \vec{x}_t. \tag{5.2}$$

Given the distinction between the top and bottom half of $\vec{x}_t$, we decompose the shock transmission matrix into its four quadrants,

$$A_t = \begin{bmatrix} A^{ll}_t & A^{vl}_t \\ A^{lv}_t & A^{vv}_t \end{bmatrix}, \tag{5.3}$$

where each of the components $A^{ll}_t, A^{lv}_t, A^{vl}_t$ and A^{vv}_t are $N \times N$ matrices, so that Eq. (5.2) can be written in the form

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

$$\vec{x}_{t+1} = \begin{bmatrix} \vec{x}^l_{t+1} \\ \vec{x}^v_{t+1} \end{bmatrix} = \begin{bmatrix} A^{ll}_t \vec{x}^l_t + A^{vl}_t \vec{x}^l_t \\ A^{lv}_t \vec{x}^l_t + A^{vv}_t \vec{x}^v_t \end{bmatrix}. \tag{5.4}$$

Eq. (5.4) makes explicit how the diagonal quadrant A^{ll}_t describes the propagation of liquidity shocks and A^{vv}_t the propagation of valuation shocks. The off-diagonal quadrant A^{lv}_t gives the conversion of liquidity to valuation shocks and A^{vl}_t the conversion of valuation to liquidity shocks. Hence, each of the five described contagion channels is associated with a specific quadrant of the shock transmission matrix.

The shock transmission matrix can be used to study the system's stability and resilience to shocks. Because all its elements are non-negative, the Perron-Frobenius theorem guarantees that the matrix has a non-negative real eigenvalue greater than or equal to (the magnitude of) the matrix' other eigenvalues. This largest eigenvalue describes the dominant dynamics of the financial system[8]: If the largest eigenvalue is greater than one, shocks that are not orthogonal to the corresponding eigenvector are amplified without bound, causing increasingly more funding and overlapping portfolio contagion. Hence, an eigenvalue greater that one indicates a positive feedback loop that progressively depresses market and funding liquidity, which we refer to as a liquidity spiral. When the largest eigenvalue is smaller than one, contagious propagation of shocks may still worsen market and funding liqiuidities beyond the impact of the initial shock, but liquidity eventually stabilizes as the shock is damped out.

5.3.4 Pecking Orders

What actions institutions choose to take in response to liquidity shocks determines the contagion that institutions transmit and the corresponding entries of the shock transmission matrix. 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 it liquidates its assets (Kok and Montagna, 2016, Hałaj, 2018, Wiersema et al., 2019). 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 [9].

[8] See e.g. Caccioli et al. [2014], Bardoscia et al. [2017], Cont and Schaanning [2019], Wiersema et al. [2019].

[9] See e.g. Allen and Gale [2000], Kok and Montagna [2016], Hałaj [2018], Wiersema et al. [2019].

van den End and Tabbae [2012] observe Dutch banks to use a *uniform pecking order* during crises while employing *liquidity-differentiated pecking orders* in benign times. An institution with a uniform pecking order does not differentiate between liquid assets of various types and uses all simultaneously by liquidating (part of) each asset to respond to shocks. 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 (Kim, 1998, van den End and Tabbae, 2012, Ma et al., 2020). 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 liquidity during crises (see e.g. De Haan and van den End, 2013, De Santis, 2014) or to preemptively divest from illiquid securities to avoid being 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).
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 and the institution is forced to move on to liquidating the asset that is next in line. As long as the asset at the top of the institution's pecking order is not exhausted, that asset exclusively determines the contagion that the institution transmits in response to liquidity shocks (and accordingly determines the corresponding entries of the shock transmission matrix). 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, which changes the dynamics of the system; when the top asset is exhausted, the asset that is next in line becomes the new asset at the top of the pecking order and the contagion transmitted and corresponding entries of the shock transmission matrix change accordingly. As we will see below, pecking orders strongly affect the potential for liquidity spirals to emerge.

5.3.5 Stylized Example

The framework that we have described allows us to study the impact of institutions' pecking orders on the potential for complex liquidity spirals to emerge. As an example, consider a simple banking system where all banks hold only two types of liquid assets; deposits, which can be either at the central bank or at other banks in the system, and (non-redeemable) 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 banks' pecking orders determine whether funding or market liquidity falls first during crises, as the contagion transmitted in response to a liquidity shock is initially determined exclusively by the assets at the top of an institution's pecking order (until liquidity shocks exhaust these assets). 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 are an important determinant of financial stability. When all banks have only deposits at the top of their pecking orders, the lower-left quadrant of the shock transmission matrix (which captures the overlapping portfolio contagion caused by selling securities) is empty so the matrix is upper block-triangular, while for other pecking orders this is not the case. When the matrix is (upper) block-triangular, its largest eigenvalue is determined by its diagonal quadrants and is not affected by the upper-right quadrant, i.e. the deleveraging contagion channel. As banks' leverages are typically on the order of ten in well-developed financial systems, deleveraging strongly amplifies shocks and tends to raise the largest eigenvalue (Wiersema et al., 2019). This makes liquidity spirals more likely when the shock transmission matrix is not upper block-triangular and therefore banks' high leverages push up the largest eigenvalue. Hence, when all banks have the optimistic pecking order, the system may be quite resilient to small shocks. However, when a large shock exhausts the assets at the top of institutions' pecking orders, the shock transmission matrix is no longer upper block-triangular, which raises the largest eigenvalue and may potentially cause a liquidity spiral to emerge.

A financial system which is resilient to small shocks but becomes unstable when hit by a large shock, as may be the case when a sufficient number of institutions have the optimistic pecking order, may be referred to as "robust-yet-fragile". Gai and Kapadia [2010] find a similar phenomenon for certain topologies of financial networks. The identification of this robust-yet-fragile tendency of financial systems across multiple dimensions 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 resilience of a financial system only to specific stress scenarios, which may cause severe instabilities to go unnoticed.

Finally, we discuss why the funding contagion channel, given by the upper-left quadrant of the shock transmission matrix, is the only channel that can cause a liquidity spiral in the absence of other contagion channels in this system. 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 nor overlapping portfolio contagion can individually drive the emergence of a liquidity spiral. Furthermore, when the lower-right quadrant, i.e. 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. 5.4). 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. 5.4). Hence, absent liquidity hoarding, column sums in the funding contagion quadrant cannot exceed one. Note that 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 (so the largest eigenvalue is given by the funding contagion quadrant) and absent liquidity hoarding (so 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.

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

We demonstrate our framework for identifying complex liquidity spiral by applying it 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).

5.4.1 Institutions

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

5.4.1.1 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 5.1a. 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 is well-developed financial systems, and the smaller banks typically have lower leverages.

5.4.1.2 Investment Funds

Investment funds pool investors' money and purchase a portfolio of securities, thereby offering investors the opportunity to obtain exposure to a diverse portfolio of underlying securities, without having to purchase and trade securities directly. From the investor's perspective, investment funds provide investors with an opportunity to earn higher returns than those offered by deposits, in return for taking on greater risk. There are over 1200 open-ended 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 5.1b.
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 anticipate or observe 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., 2021).
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 shares outstanding. 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 share portfolio is identical. Therefore, we assume for simplicity that all fund shares in our model are VNAV-valued.

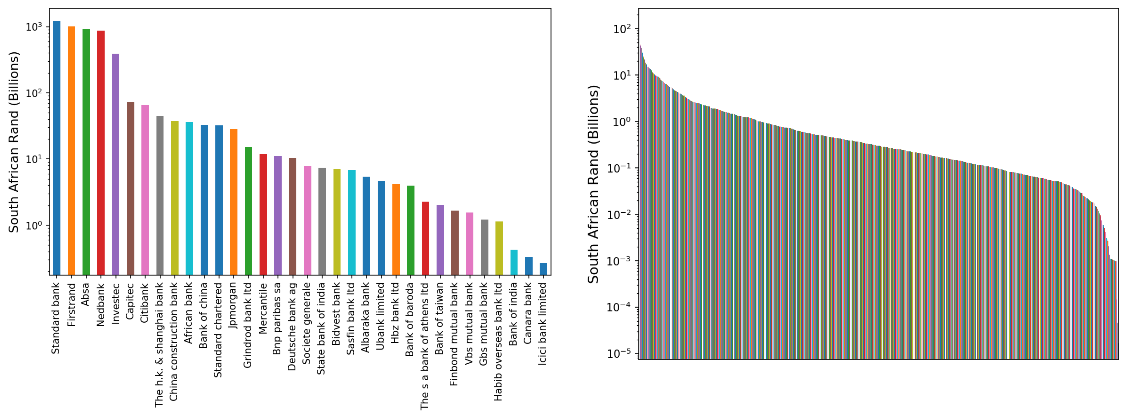

(a) Banks ranked by total assets

(b) Investment Funds ranked by total assets

Figure 5.1: **Distribution of South African financial institutions by asset size.** The institutions are listed on the x-axis in decreasing order of total assets size and their total assets in billions of South African Rand are on the y-axis (log-scale). Note that the funds' names are not listed because they are too numerous. The banking sector consist of a core of five large banks and a periphery of 29 smaller banks. The investment fund sector includes over 1200 funds and also shows a strong concentration in terms of asset size.

5.4.2 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, 2016b). 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, non-financial corporates and the government). 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 investment 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, 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.

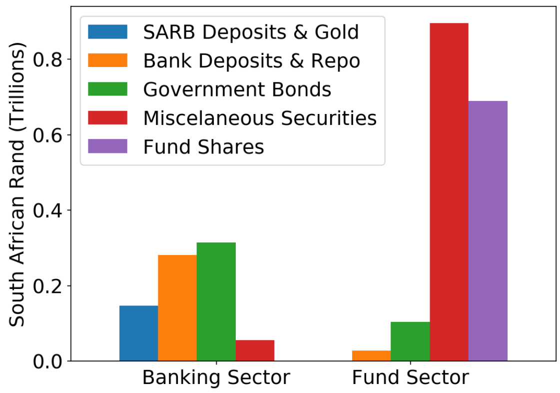

Figure 5.2: **Aggregate asset holdings per category of the banking sector and fund sector.** Asset holdings are aggregated into the categories set out in section 5.4.2.1 for the banking sector and fund sector. Banks do not hold fund shares, while funds do not hold South African Reserve Bank deposits or gold. Miscellaneous securities include non-government bonds, money market instruments, and equity shares. Note that the banks' position in bonds issued by the South African government far exceeds the funds' position, while the funds' position in miscellaneous tradable securities far exceeds that of the banks.

5.4.2.1 Balance Sheet Composition

We include the following assets: Gold, deposits, repurchase agreements, or *repo*, money market instruments, or *MMIs*, bonds, equity shares, and fund shares. We refer to non-government bonds, MMIs, and equity shares as *miscellaneous securities*. Figure 5.1 shows where these assets appear on the stylized balance sheets of banks and investment funds. Note that some investment funds may heavily invest in one type of asset while not investing in other types.

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

(a) Stylised balance sheet of a bank.

Assets		**Liabilities**
Bank Deposits		Fund shares
Tradable securities	MMIs	
	Bonds	
	Equity shares	
Fund shares		
Other assets		

(b) Stylised balance sheet of an investment fund.

Table 5.1: **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 Deposits & Gold

Banks invest in gold and make deposits at the SARB, whereas investment funds do not. We assume that both are perfectly liquid and that neither causes contagion when liquidated. This makes their dynamics in our model identical, so they are grouped together for simplicity.

Bank Deposits & Repurchase Agreements

South African banks receive deposits and issue repo, while investment funds do not (as they do not have debt). Furthermore, while both banks and funds make deposits at (other) South African banks, only the banks buy repo. The banks and funds also make deposits at foreign banks, but these make up a very small part of their portfolios. Because we do not explicitly model collateral and assume that both repo and deposits can be redeemed on a daily basis, their dynamics in our model are identical. As such, we group repo and deposits at (commercial) banks together for simplicity.

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 positions are relatively small. The investment fund data distinguish between bonds issued by different banks, whereas the bank data do not. Contrary to repo and deposits, bonds are tradable.

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. MMIs in our data are exclusively issued by *domestic* banks and are bought by both banks and funds. Like bonds, MMIs are tradable.

Equity shares

Funds invest in listed equity shares issued by the South African banks and corporate sector, while banks invest in listed shares issued by the corporate sector and 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.

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.

5.4.2.2 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). For simplicity, we normalize the initial NAV of each fund share and the initial market price of each tradable security to one South African Rand. This normalization has no effect on our results and is only for simplicity. Furthermore, we assume 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 the issuer's accounting value of that share).

5.4.2.3 Interbank Asset Allocation

The contagion channels that we model require reconstruction of the counterparties of interbank deposits, repo, and unlisted equity shares as the bank data only provide banks' aggregate assets and liabilities. Due to data limitations, we do not distinguish between tradable securities of a specific type issued by different non-financial corporates, nor between tradable securities of a specific type issued by different domestic banks. Therefore, we do not reconstruct counterparties of banks' investments in these securities.

The technique used for the reconstruction of the banks' investments is similar to Kok and Montagna [2016] and Wiersema et al. [2021], and aims to reproduce the high heterogeneity of interconnections observed in financial networks. 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
$\beta \in \{\text{deposits, repo, (unlisted) equity shares}\}$:

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.[10]

[10] 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.

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.

5.4.3 Contagion Equations

We now derive representative formulas for each of the contagion channels that acts on the described asset classes. Note that these forms are chosen for simplicity. 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.

5.4.3.1 Valuation Shock-Induced Contagion

We discuss the contagion caused by valuation shocks first, as liquidity shock-induced contagion requires a more elaborate discussion. We focus on valuation shocks that do not exceed banks' equity as we are interested in how liquidity spirals emerge, while (major) institutions are not expected to default before the liquidity spiral has grown into a systemic crisis.

Deleveraging Contagion

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

$$A^{vl}_{ii,t} = \lambda_i, \tag{5.5}$$

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.

Shareholder Contagion

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

$$A^{vv}_{ji,t} = \frac{s_{ij,t}}{S_{j,t}}, \tag{5.6}$$

when j is a fund.

Let us now consider shareholder contagion for 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. 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 (5.6) also holds for unlisted equity shares issued by banks.

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, an efficient market price reflects the issuer's performance. We 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 (5.6) also holds for listed shares.

Share 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 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 when they themselves suffer a liquidity shock, i.e. through funding contagion). Furthermore, we assume for simplicity that all external holders withdraw liquidity from i proportionally to the loss i suffered at the same *redemption rate* R[11], which implies that

$$A^{vl}_{ii,t} = \epsilon_{i,t} R, \tag{5.7}$$

[11] We show in section C.1 in the appendix 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., 2021).

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.

5.4.3.2 Liquidity Shock-Induced Contagion

The contagion that an institution transmits in response to a liquidity shock is determined by its pecking order. To distinguish between various liquidity-differentiated pecking orders, we group assets in decreasing order of liquidity as follows:

1. Central bank deposits and gold
2. Deposits at commercial banks, repo, and fund shares
3. Government bonds
4. Miscellaneous tradable securities (MMIs, listed equity, bank bonds, and corporate bonds)

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.
Given the limited empirical research into the pecking orders that institutions employ under various circumstances, we consider two more pecking orders for completeness. The *short-term funding pecking order* is the optimistic pecking order but with group 1. moved to the bottom and hence group 2. at the top, and the *government bonds pecking order* is the optimistic pecking order both groups 1. and 2. moved to the bottom and hence group 3 at the top. Hence, institutions with the short-term funding pecking order use deposits, repo and fund shares as their first line of defense against liquidity shocks, which in practice are indeed a common source of liquidity for many institutions, and institutions with the government bonds pecking order use their bonds for this purpose, which is also often observed in practice. The main motivation for including these two additional pecking orders is that each of the four groups of liquid assets we distinguish is now at the top of

one of the four liquidity differentiated pecking orders we consider. Table 5.2 summarizes the various pecking orders that we consider.

	Optimistic pecking order	Short-term funding pecking order	Government bonds pecking order	Conservative pecking order	Uniform pecking order
Top	Central bank deposits and gold	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 and gold	Deposits at commercial banks, repo, fund shares	
Bottom	Miscellaneous tradable securities	Central bank deposits and gold	Deposits at commercial banks, repo, fund shares	Central bank deposits and gold	

Table 5.2: **Pecking orders**

An institution may have multiple assets at the top of its pecking order. 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, repo and fund shares in multiple institutions at the top of its pecking order. When an institution has multiple assets at the top of its pecking order, we assume that the institution liquidates a vertical slice across all these assets, i.e. the institution recovers an amount of liquidity from each asset proportional to that asset's total value. As a consequence, an institution with the uniform pecking order reduces each of its liquid assets by the same proportion in response to a liquidity shock and hence the shock does not change the institution's pecking order, leaving the institution's response to liquidity shocks unchanged. We assume for simplicity that the vertical slice used to respond to the liquidity shock $\vec{x}_t^l$ is based on the asset values at the start of round t (i.e. before the liquidity and valuation shocks $\vec{x}_t^l$ and $\vec{x}_t^v$ are taken into account, which may reduce the value of some securities).

We now derive contagion equations for the liquidation of each type of asset that may be at the top of the pecking order, under the assumption that the liquidity shock does not exceed the top layer. Note that we do not consider liquidity hoarding, so institutions recover liquidity equal to the shock incurred.

Funding Contagion

Suppose institution i has deposits at and/or has bought repo or fund shares issued by institution j with a total value of $d_{ij,t}$. Let $T_{t,i}$ denote the total asset value of the top layer of i's pecking order. 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

$$A^{ll}_{ji,t} = \frac{d_{ij,t}}{T_{i,t}}. \tag{5.8}$$

We assume that the withdrawal of deposits from the SARB or foreign institutions does not cause funding contagion and hence "dilutes" liquidity shock-induced contagion (as less liquidity is required to be recovered from other sources).

Overlapping portfolio contagion

Suppose institution i holds $n_{\sigma i,t}$ shares of security σ with 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 to raise an amount of liquity $x^l_{i,t}n_{\sigma i,t}p_{\sigma,t}/T_{i,t}$ (i.e. a vertical slice). The sale depresses the price of security σ by $\Delta p_{\sigma,t} = p_{\sigma,t} - p_{\sigma,t+1}$, which causes losses to all institutions that have a position in the security. We assume that the price impact $\Delta p_{\sigma,t}$ is linear in the amount of liquidity to be raised by selling shares in security σ[12] (so price impacts are additive when multiple institutions i decide to sell shares in security σ at time t);

$$\Delta p_{\sigma,t} = \sum_{i\in\mathcal{S}_t} \frac{x^l_{i,t}n_{\sigma i,t}p_{\sigma,t}}{T_{i,t}D_\sigma}. \tag{5.9}$$

The market depth D_σ (discussed below) is a non-dimensional constant which gives the liquidity recovered per unit drop in the securities' market price. and $\mathcal{S}_t$ the set of institutions that sell security σ at time t. We also assume for simplicity that all shares sold at time t are sold against the new price $p_{\sigma,t+1}$, such that any institution,

[12]We show in section C.2 in the appendix 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). 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].

including institutions $i \in \mathcal{S}_t$, suffers a valuation shock of $\Delta p_{\sigma,t} n_{\sigma j,t}$[13]. This implies that

$$A^{lv}_{ji} = \sum_{\sigma \in \mathcal{T}_{i,t}} \frac{n_{\sigma i,t} p_{\sigma,t} n_{\sigma j,t}}{T_{i,t} D_\sigma}, \tag{5.10}$$

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 market price of listed equity shares issued by South African banks is depressed by both overlapping portfolio and shareholder contagion. We assume for simplicity that the overlapping portfolio and shareholder contagion channels' impact on the market price is additive such that equation (5.10) also holds for listed shares issued by banks. The combined impact of the overlapping portfolio and shareholder contagion channels on listed shares should not drive their price below zero, as the shares are subject to limited liability. This is guaranteed in our results, as explained in section C.3 in the appendix.

Market depths: For each security class σ, we explore various market depths D_σ by dividing the *baseline estimate* $\hat{D}_\sigma$ by *market depth divisor* δ_σ,

$$D_\sigma = \frac{\hat{D}_\sigma}{\delta_\sigma}. \tag{5.11}$$

The baseline market depth $\hat{D}_\sigma$ is a non-dimensional constant calibrated to the market capitalization of security σ, as explained in appendix C.4, and the market depth divisor δ_σ is a non-dimensional constant which we vary from one to four to explore the sensitivity of our results to the market depth. Hence, when the market depth divisor $\delta_\sigma = 1$, the market depth of σ is equal to the baseline estimate, and when the market depth divisor $\delta_\sigma = 0$, the market depth of σ is infinite and the overlapping portfolio contagion channel is effectively turned off.

Note that our data include tradable securities issued by foreign institutions and gold. Because these securities are predominantly held by foreign institutions, we do not model the overlapping portfolio contagion caused by the sale of these securities, so we effectively assume that these securities have infinite market depth. Consequently, overlapping portfolio contagion only spreads across domestic securities. Hence, foreign securities and gold "dilute" liquidity shock-induced contagion,

[13] In reality, some institutions may decide to sell shares incrementally, while others may under- or overestimate the amount of liquidity they will recover from selling securities, and still others may intentionally liquidate more than the required amount of securities out of conservatism (which is referred to as "liquidity hoarding"). However, for our present purposes, it suffices to simply assume that each institution liquidates the correct number of securities against the new price $p_{\sigma,t+1}$ to recover the required liquidity. Note that the price impact suffered by institutions $i \in \mathcal{S}_t$ implies that the diagonal component of $A^{lv}_{ii,t}$ for $i \in \mathcal{S}_t$ in equation (5.10) is positive.

because the liquidity recovered from foreign securities and gold does not cause contagion but reduces the liquidity required to be recovered from other sources.

5.4.4 Results

Here, we study the contribution of the banking and investment fund sectors to the emergence of complex liquidity spirals for various pecking orders. We do so by comparing the stability of the full system to the stability of each sector individually, where the largest eigenvalue of an individual sector is calculated from the shock transmission matrix that only includes institutions belonging to that sector. The shock transmission matrices are calculated for time $t = 1$, to which securities' market prices, book values, as NAVs are normalized. Note that all results present the means over 1000 random generations of the interbank assets reconstruction. As these interbank assets form a small part of banks' total balance sheets, they do not affect our results significantly and therefore the standard errors of the means are negligible and not visible in the plots.

5.4.5 Uniform Pecking Orders

In Figure 5.3, we plot the largest eigenvalue of the full system, and of the banking sector and fund sector individually, under the assumption that all institutions have the uniform pecking order. In Figure 5.3a we set the market depth divisor to its baseline value of $\delta_\sigma = 1$ for all securities and vary the redemption rate R to explore the impact on stability, and find that a liquidity spiral emerges for about $R = 5$. Note that the banking sector is not affected by the redemption rate R, so increases in the redemption rate affect stability only through the fund sector.
In Figure 5.3b, we set the redemption rate to its baseline value of $R = 1$ and vary the market depth divisor δ_σ for all securities. We find that the stability is greatly affected by the market depth, and that a liquidity spiral emerges as soon as the market depth divisor δ_σ increases beyond one. Hence, a liquidity spiral emerges almost as soon as market depths fall below their baseline values, while only particularly high redemption rates cause a liquidity spiral.
In Figure 5.3c and Figure 5.3d we consider how the South African Reserve Bank, acting as a lender of last resort, may attempt to dampen the liquidity spiral by injecting cash into the banking sector. Because banks use the uniform pecking order and therefore liquidate proportionally to asset values, raising the banks' central bank deposits increases the banks' reliance on their central bank deposits to meet

liquidity shocks and decreases their reliance on other liquid assets. Withdrawing central bank deposits does not cause contagion, while the liquidation of other assets typically does. Hence, increasing the banks' central bank deposits stabilizes the system.

In Figure 5.3c we set the market depth divisor to $\delta_\sigma = 2$ for all securities and the redemption rate to its baseline value of $R = 1$, such that the liquidity spiral is mainly driven by the banking sector. Conversely, in Figure 5.3d we set the market depth divisor to its baseline value of $\delta_\sigma = 1$ for all securities and the redemption rate $R = 5$, such that the liquidity spiral is mainly driven by the fund sector. The increase in banks' central bank deposits as a fraction of their total liquid assets is presented on the x-axis in both figures. Although the liquidity spiral dissipates in both cases, Figure 5.3d shows that the stabilizing effect of the cash injection into the banking sector is minimal when the fund sector is the main driver of the liquidity spiral, while Figure 5.3c shows that the cash injection is more effective when the banking sector is the main driver of the liquidity spiral.

Finally, comparison of Figure 5.3b and Figure 5.3c suggests that, compared to the impact of the market depth, the effect of the cash injection on the largest eigenvalue is modest even when the banking sector is the main driver of the liquidity spiral. However, note that the x-axis Figure 5.3b ranges from an infinitely deep market ($\delta_\sigma = 0$), up to a reduction of the baseline market depth by up to a factor of $\delta_\sigma = 4$, while Figure 5.3c only covers up to a doubling of banks' liquid assets. This explains, at least partially, why the eigenvalue is more strongly impacted in Figure 5.3b than in Figure 5.3c.

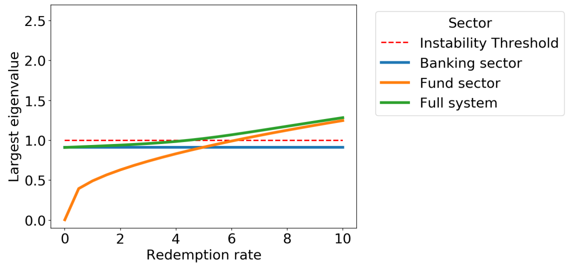

(a) Eigenvalue vs. Redemption Rate R

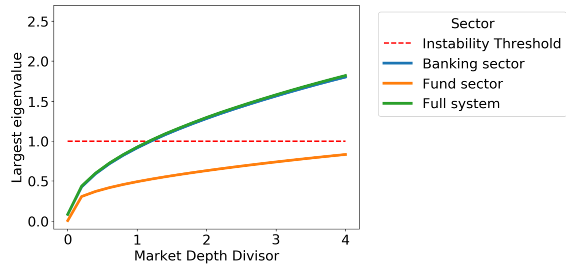

(b) Eigenvalue vs. Market Depth Divisor δ_σ

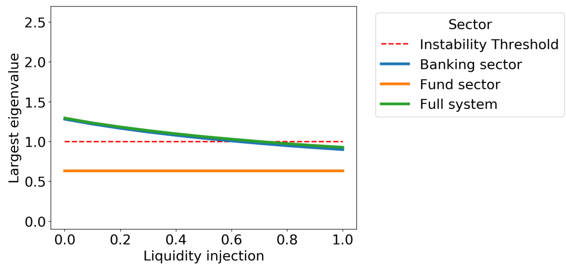

(c) Liquidity Provision for $\delta_\sigma = 2$

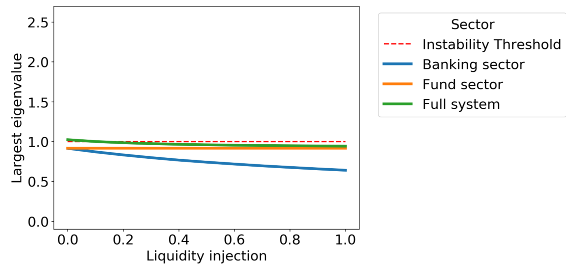

(d) Liquidity Provision for $R = 5$

Figure 5.3: **Eigenvalue Dependency for Uniform Pecking Orders.** All institutions are assumed to have the uniform pecking order. We explore when the largest eigenvalue exceeds one and a liquidity spiral emerges for various redemption rates R in (a) and for various market depth divisors δ_σ in (b). (c) shows the impact of a central bank cash injection into the banking sector when $\delta_\sigma = 2$ and $R = 1$, and (d) shows this impact when $R = 5$ and $\delta_\sigma = 1$.

5.4.6 Liquidity-Differentiated Pecking Orders

We now consider the stability of the four liquidity-differentiated pecking orders, under the assumption that shocks do not exhaust the assets at the top of any institution's pecking order. As such, only the assets at the top of each institution's pecking order need to be considered when calculating the dynamics' largest eigenvalue. As Figure 5.2 gives the aggregate values of all assets that may be at the top of institutions' pecking orders, the Figure gives an indication of how large shocks may be before the assets at the top of institutions' pecking orders are exhausted. For example, Figure 5.2 shows that the assets at the top of funds' government bonds pecking orders are generally exhausted quickly, while for banks this is the case for the assets at the top of conservative pecking orders. The assets at the top of other pecking orders are considerably more substantial. After understanding how the liquidation of each type of asset affects stability, we consider shocks that exceed the assets at the top of institutions' pecking orders (so institutions must liquidate the assets that are next in line) in section 5.4.7.

In Figure 5.4, we plot the largest eigenvalue of the banking sector and fund sector, and of the full system, for the four liquidity-differentiated pecking orders. We set the market depth divisor to its baseline value of $\delta_\sigma = 1$ for all securities and vary the redemption rate R. In Figure 5.4a, all institutions are assumed to have the optimistic pecking order and in Figure 5.4d the conservative pecking order. In Figure 5.4b, all institutions have the short-term funding pecking order, and in Figure 5.4c all institutions have the government bonds pecking order.

The results show that the optimistic and short-term funding pecking orders are very stable, because the liquidation of the assets that are at the top of these pecking orders does not cause high levels of contagion; even for very high redemption rates no liquidity spiral emerges. Conversely, when government bonds are at the top of the pecking order, as shown in 5.4c, the banking sector is highly unstable. The main reason for this is that the majority of the government bonds in our data are held by banks, so the price impact caused by selling the bonds is predominantly suffered by the banks, and that the banks strongly amplify any valuation shocks that they incur through their high leverages (Wiersema et al., 2019). Furthermore, 5.4c shows that for high redemption rates the largest eigenvalue of the full system is substantially higher than the largest eigenvalue of than either of the individual sectors. Hence, the intensity of the liquidity spiral would be underestimated when the interactions between the banking and fund sector are ignored, which highlights the importance of capturing these sectors' combined dynamics.

Figure 5.4d shows that the conservative pecking order is generally more stable than the government bonds pecking order in Figure 5.4c, but is also more strongly affected by the redemption rate R. The main reason for this is that the miscellaneous tradable securities in our data are mainly held by investment funds, so the funds suffer most of the price impact caused by selling the miscellaneous tradable securities, and contrary to banks, funds only amplify valuation shocks when the redemption rate R is high.

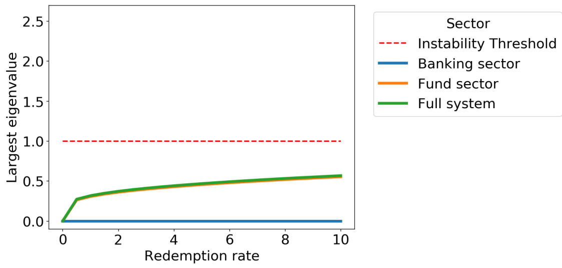

(a) Optimistic Pecking Order

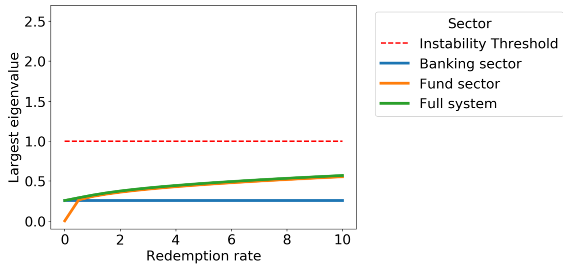

(b) Short-Term Funding Pecking Order

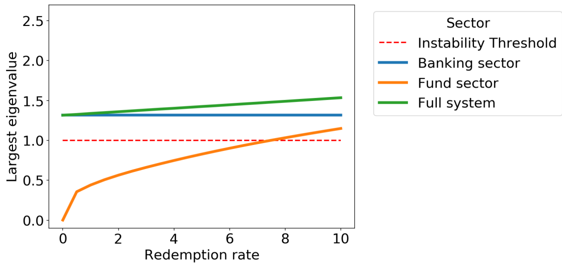

(c) Government Bonds Pecking Order

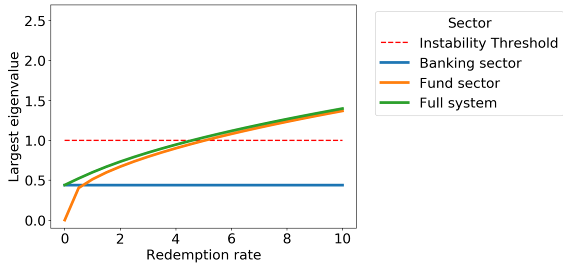

(d) Conservative Pecking Order

Figure 5.4: **Eigenvalue Dependency on Redemption rates for Liquidity-Differentiated Pecking Orders.** We explore when the largest eigenvalue exceeds one and a liquidity spiral emerges for various redemption rates R and various liquidity-differentiated pecking orders. The market depth divisor is set to its baseline value of $\delta_\sigma = 1$. In (a) all institutions are assumed to have the optimistic pecking order, in (b) the short-term funding pecking order, in (c) the government bonds pecking order and in (d) the conservative pecking order.

Figure 5.5 is analogous to Figure 5.4, but here we vary the market depth divisor δ_σ (for all securities) rather than the redemption rate, which is set to its baseline value of $R = 1$. The results again show the standard and short-term funding pecking orders to be very stable as no liquidity spiral emerges in Figure 5.5a and Figure 5.5b even for very illiquid markets (i.e. high δ_σ). When government bonds are at the top of the pecking order, as shown in Figure 5.5c, the banking sector is again highly unstable, and is strongly affected by the market depth. Similar to what we found in Figure 5.4, Figure 5.5d shows that the conservative pecking order is generally more stable than the government bonds pecking order. Furthermore, the eigenvalue of the full system in Figure 5.5d shows that when markets are very illiquid (i.e. high δ_σ), *a liquidity spiral emerges which is completely overlooked when ignoring the interactions between the banking sector and fund sector* (as neither individual sector has an eigenvalue greater than one).

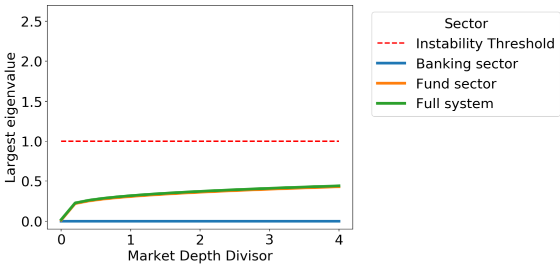

(a) Optimistic Pecking Order

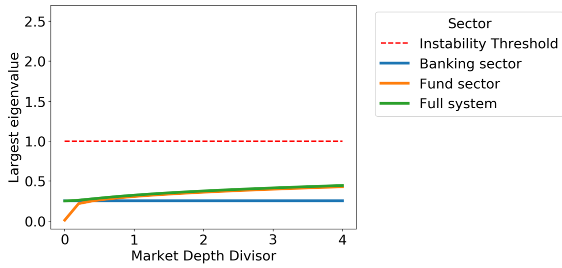

(b) Short-Term Funding Pecking Order

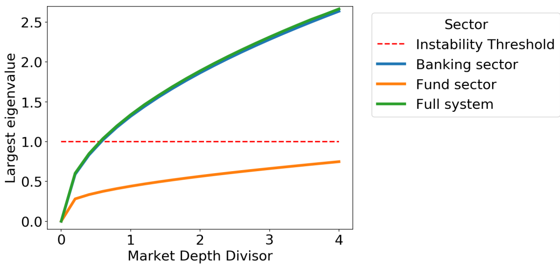

(c) Government Bonds Pecking Order

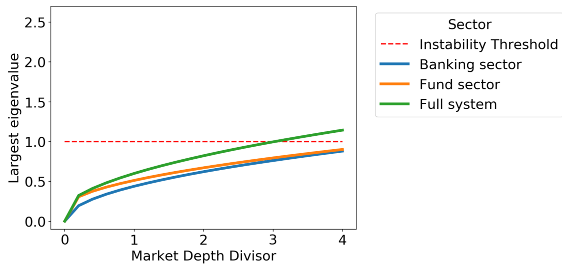

(d) Conservative Pecking Order

Figure 5.5: **Eigenvalue Dependency on Market Depths for Liquidity-Differentiated Pecking Orders.** We explore when the largest eigenvalue exceeds one and a liquidity spiral emerges for various market depth divisors δ_σ and various liquidity-differentiated pecking orders. The redemption rate is set to its baseline value of $R = 1$. In (a) all institutions are assumed to have the optimistic pecking order, in (b) the short-term funding pecking order, in (c) the government bonds pecking order and in (d) the conservative pecking order.

5.4.7 Liquidity Shocks

In Figure 5.6, we consider how the stability of the system is impacted by a large liquidity shock when all institutions have the optimistic pecking order (which empirical evidence suggests to be the most commonly employed pecking order[14]). As we have seen, the system is very resilient to the liquidation of the assets at the top of the optimistic pecking order, so when a liquidity shock exhausts these assets and forces institutions to liquidate assets that are next in-line on their pecking orders, we can expect the system to become less stable. Figures 5.6a and 5.6b show the sensitivity of the largest eigenvalue to the shock size and Figure 5.6c and Figure 5.6d show which asset types are at the top of institutions' pecking orders after the shock has exhausted part of their liquid assets.

We consider the marginal impact on stability of an increase in the shock size, by calculating the largest eigenvalue of the system according each institution's assets

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

lowest on its pecking order that it is forced to liquidate in response to the liquidity shock. For simplicity, we assume that the liquidity shock reduces each institution's total liquid asset holdings by the same proportion. Although we can reasonably expect an institution's total liquid asset holdings to be (at least somewhat) calibrated to the magnitude of the liquidity shocks that the institution expects to incur, liquiditiy shocks come in various distributions and magnitudes in reality. Future research should therefore explore the stability of financial systems across various distributions of liquidity shocks.

In Figure 5.6a, the redemption rate and market depth divisor are set to their baseline values of $R = 1$ and $\delta_\sigma = 1$. The reduction in institutions' liquid assets due to the liquidity shock is presented as a proportion of the institution's total liquid asset holdings on the x-axis. Note that the proportion cannot exceed one as we do not model illiquidity defaults. Figure 5.6a shows that the liquidity shock has the potential to greatly destabilize the system, as a liquidity spiral emerges when about half of institutions' liquid assets are exhausted and institutions are forced to liquidate assets lower on their pecking orders.

Figure 5.6a also shows that the liquidity spiral that emerges when institutions consume their pecking orders is predominantly driven by the banking sector. Comparison with Figure 5.6c allows us to understand how the gradual depletion of banks' pecking orders causes the liquidity spiral to emerge. Figure 5.6c shows that all but one bank initially have central bank deposits and/or gold at the top of their pecking orders (as shown by the blue line), and hence the system is initially very stable, but the number of banks with these assets at the top of their pecking order drops off quickly as the shock size increases. Nevertheless, this does not immediately cause instabilities as most of these banks start liquidating interbank deposits and/or repo instead (as shown by the orange line) which we have seen in previous sections not to cause liquidity spirals either. The liquidity spiral in 5.6a only emerges when the number of banks with interbank deposits and/or repo at the top of their pecking orders starts falling too and more and more banks start liquidating government bonds instead (as shown by the green line in 5.6c). Indeed, we have seen in previous sections that the liquidation of government bonds drives the emergence of liquidity spirals.

Figure 5.6a also shows that the liquidity spiral dissipates when the shock size approaches the point of fully exhausting institutions' pecking orders. Figure 5.6c shows that this dissipation coincides with a strong drop in the number of institutions that liquidate interbank deposits and/or repo and a steep rise in the number of institutions that sell miscellaneous tradable securities (as shown by the red line). Interestingly, we previously showed for $R = 1$ and $\delta_\sigma = 1$ that the short-term funding pecking order (which has bank deposits, repo and fund shares at the top) is more stable than the conservative pecking order (which has miscellaneous tradable securities at the top) in a system where all institutions have the same pecking order. Conversely, Figure 5.6 suggests that in a system where a substantial number of banks have government bonds at the top of their pecking order, withdrawing interbank deposits and repo is more destabilizing than selling miscellaneous tradable securities to meet a liquidity shock. This is a prime example of *emergence*, as the interactions between institutions with different pecking orders yields dynamics the one would not expect based on the individual pecking orders, and highlights the importance of capturing the complex nature of financial systems.

The dissipation of liquidity spirals when institutions' pecking orders are close to exhaustion would be highly advantageous, as it could stabilize the system before institutions default through illiquidity. However, we show in Figure 5.6b that this dissipation may not materialize during adverse market conditions. In Figure 5.6b, the redemption rate is set to $R = 5$ and the market depth divisor is set to $\delta_m = 2$ for miscellaneous tradable securities, but to $\delta_g = 1$ for government bonds (i.e. the market depth of miscellaneous tradable securities is halved, while the market depth of government bonds remains at its baseline value). For these market conditions, the liquidity spiral is predominantly driven by the fund sector, rather than the banking sector.

Figure 5.6d shows that almost all funds initially have bank deposits and/or fund shares at the top of their pecking orders (as shown by the orange line). However, this number falls of sharply as the shock size reaches just 10% of institutions' total liquid asset holdings. This drop coincides with a strong increase in the number of funds that have miscellaneous tradable securities at the top of their pecking orders (as shown by the red line), and with the emergence of a liquidity spiral as show in Figure 5.6b. The number of funds with bank deposits and/or fund shares at the top of their pecking orders continues to fall, and the number of funds that sell miscellaneous tradable securities continues to rise, until the shock size reaches 100%. The liquidity spiral, rather than dissipate, grows steadily in intensity until the maximum shock size is reached.

Finally, let us compare the optimistic pecking order in 5.6a to the uniform pecking order in Figure 5.3. Comparing largest eigenvalues for the same redemption rate

$R = 1$ and market depth modifier $\delta_\sigma = 1$, we find that the optimistic pecking order yields a system more resilient to small shocks than the uniform pecking order. However, the uniform pecking order is not affected by liquidity shocks, as explained in section 5.4.3.2, while a liquidity spiral emerges in response to large shocks when institutions have the optimistic pecking order. Relative to the uniform pecking order, the optimistic pecking order therefore yields a "robust-yet-fragile" system, which is very resilient to small shocks but may be greatly destabilized by a single, large shock. This is similar to the robust-yet-fragile network topologies identified by Gai and Kapadia [2010], but manifest here in terms of pecking order configurations, and highlights the importance of evaluating the resilience of financial systems against a wide range of shocks.

Note that by focusing on time $t = 1$, we have only considered how a liquidity spiral emerges but not how it evolves as shocks continue to propagate. Even when all institutions have the uniform pecking order, which does not change in response to shocks, the dynamics of the system evolve as shocks propagate and tradable securities change hands (which changes how overlapping portfolio contagion is distributed), shares are redeemed (which dampens the shareholder contagion channel), and additional cash flows enter or leave the system (which may affects institutions' pecking orders). However, modeling such changes to the system would require additional assumption and/or (empirical) research beyond the scope of this paper.

Furthermore, Figure 5.6 shows that a large liquidity shock may raise the largest eigenvalue above one and cause a liquidity spiral to emerge. However, a smaller shock that pushes the eigenvalue close to, but not outright above one, may continue to propagate and drain institutions' pecking orders until a spiral eventually emerges. Although Figure 5.6 indicates how the dynamics would evolve as liquidity losses accumulate and pecking orders are progressively depleted by propagating shocks, this does not take into account the changes to the system mentioned in the previous paragraph that are unrelated to the pecking order. To accurately assess how the stability evolves as shocks continue to propagate, more comprehensive modeling approaches are required.

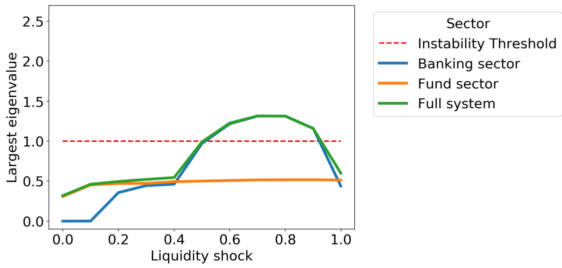

(a) Eigenvalue vs. Liquidity Shock ($R = 1$)

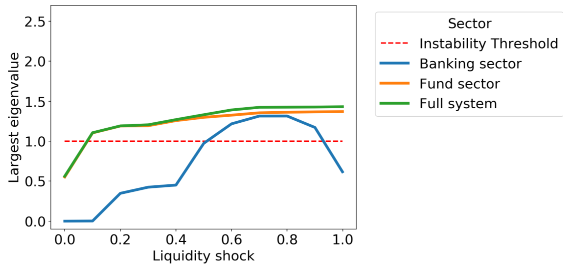

(b) Eigenvalue vs. Liquidity Shock ($R = 5$)

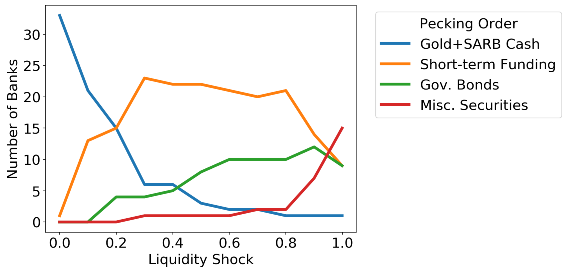

(c) Banks' Top Layer of Pecking Order

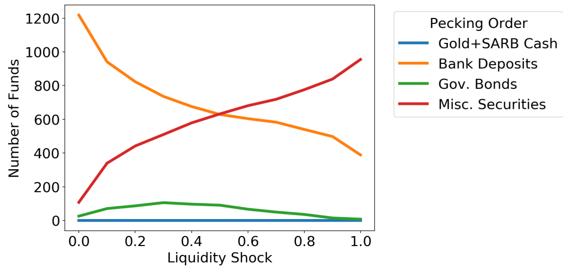

(d) Funds' Top Layer of Pecking Order

Figure 5.6: **Impact of Liquidity Shock on Largest Eigenvalue.** We explore how a large liquidity shock affects stability and may potentially cause the largest eigenvalue to exceed one and a liquidity spiral to emerge. All institution are assumed to have the optimistic pecking order and the reduction in institutions' liquid assets due to the liquidity shock is presented as a proportion of the institution's total liquid assets pool on the x-axis. In (a) the redemption rate is set to its baseline value of $R = 1$ an the market depth divisor to its baseline value of $\delta_\sigma = 1$ for all securities. In (b), the redemption rate is set to $R = 5$. The market depth divisor is set to $\delta_g = 1$ for government bonds g and to $\delta_m = 1$ for all miscellaneous tradable securities m. In (c) and (d) we show the number of banks respectively funds per asset type at the top of their pecking orders depending on the size of liquidity shock incurred. The blue line shows the number of institutions with gold and/or central bank deposits at the top of their pecking order, the orange line the institutions with commercial bank deposits, repo, and/or fund shares at the top of their pecking order, the green line government bonds, and the red line miscellaneous tradable securities.

5.5 Discussion

Liquidity spirals progressively worsen market and funding liquidity (Brunnermeier and Pedersen, 2009). We have introduced the concept of complex liquidity spirals, which consist of various interacting contagion channels and/or multiple types of institutions. To accurately assess these spirals, models are required that can take the interactions between multiple types of contagion and institutions into account. We use the framework developed in Wiersema et al. [2019], which allows us to identify complex liquidity spirals before market and funding liquidity fall progressively. Wiersema et al. [2019]) show in a general setting that financial stability may be greatly overestimated when ignoring the interactions between contagion channels. Here, we demonstrate that complex liquidity spirals are completely overlooked when interactions between different types of institutions or contagion channels are ignored.

The framework allows us to evaluate the impact of institutions' pecking orders on the potential for liquidity spirals to emerge without relying on any specific, subjective stress scenario. We show that institutions' pecking orders strongly affect finnacial stability, with some pecking orders yielding a "robust-yet-fragile" system. The robust-yet-fragile tendency of financial systems has been previously observed by Gai and Kapadia [2010] for certain network topologies. Here, we have seen it manifested for specific pecking order configurations. The identification of robust-yet-fragile tendencies of financial systems across multiple dimensions highlights the dangers of optimizing stability with respect to the small shocks that are incurred on a frequent basis; a financial system that has been optimized to be highly resilient against small shocks may turn out to be highly fragile to large shocks once one eventually materializes. Moreover, it underscores the importance of stability measures that assess a system's resilience to a wide range of shocks, such as the eigenvalue-based approach developed here.

We demonstrate our method by applying it to a highly granular data set on the South African financial system and capture the combined dynamics of the banking and investment fund sector. Wiersema et al. [2021] show that exposures in the South African financial system are underestimated when the interactions between the banking and fund sector are ignored. Here, we identify market conditions for which a liquidity spiral emerges that cannot be identified without taking the interactions between the banking and fund sector into account. These results highlight that comprehensive modeling approaches such as the one presented here are vital for understanding financial stability. We also identify market conditions that yield a liquidity spiral which is predominantly driven by one of the two sectors. This greatly affects the effectiveness of interventions such as liquidity injections into the banking sector. Hence, policy makers may employ the model developed here to decide on what strategies may be most effective at combating liquidity spirals.

We have explored how the system's stability changes in response to a large liquidity shock. We show that when institutions sell their most liquid assets first, the system is very resilient against small liquidity shocks. However, a liquidity spiral may emerge as soon as a substantial part of institutions' pecking orders are exhausted by a sizable liquidity shock. This robust-yet-fragile tendency may appear in any financial system where institutions liquidate assets in order of decreasing liquidity, as contagion typically worsens as institutions are forced to liquidate assets of lesser and lesser liquidity. This highlights the importance of exploring 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. As we have only considered liquidity shocks that are distributed proportionally to institutions' pecking order depths, future research should aim to formulate realistic stress scenarios and investigate how the system evolves over time in response to these. Furthermore, as we have established the important role that pecking orders play in the potential for liquidity spirals emerge, further empirical investigation into the pecking orders that financial institutions employ under various market conditions is warranted.

Chapter 6

Concluding remarks

In this thesis, we have explored measures of systemic risk due to interacting channels of financial contagion that do not rely on any initial stress scenario. These measures serve to complement system-wide stress tests and other tools in the macroprudential policymaker's arsenal. They offer both insights that generalize across financial systems, but also allow for calibration to the specific microstructure of real financial systems to yield accurate and highly detailed observations about the stability of the system in question. These results represent a significant and substantial contribution to the literature as they cannot be derived with the same level of accuracy from traditional measures of systemic risk.

In chapter 2, we identified a lack of models in the current financial stability literature that capture the interactions between the various contagion channels, types of institutions and asset types encountered in financial systems, whilst allowing for accurate calibration to the microstructure of real financial systems and not relying on any specific, subjective stress scenario. We developed such models in chapters 3-5 and in doing so found additional evidence for the importance of models that meet these criteria: In chapter 3 we find substantial exposures that could only be identified by modeling the interactions between at least three different sectors of the financial system simultaneously, and in chapter 4 we demonstrate that ignoring the interactions between contagion channels may arbitrarily overestimate stability. We also find in chapter 4 that using different techniques for reconstructing financial networks can lead to opposite conclusions about the networks' stability, which highlights the relevance of calibrating models to real data of networks' microstructures. Finally, we find evidence for "robust-yet-fragile" financial systems in chapter 5, which are highly resilient to one type of shocks but may be very vulnerable to another. This underscores the importance of models of financial stability that do not rely on a specific (and often subjective) stress scenario, but instead assess the resilience of the financial system to a wide range of shocks.

In chapter 3, we have introduced the concept of higher-order exposures. This concept, supported by our findings of of the South African financial system, has substantial implications for prudential policymakers, regulators, and supervisors with financial stability mandates. We have shown that direct and indirect exposures, which are traditionally used to calculate exposures, only capture part (and, in some cases, only a small part) of the exposures between financial institutions. Higher-order exposures can be significant, heterogeneous, and particularly high in times of crisis – when exposures matter most. Moreover, these exposures cannot easily be extrapolated from traditional measures of exposure and therefore require complementary analysis focusing specifically on higher-order exposures. We have outlined the implication of our results for the design and calibration of various tools in the regulatory arsenal. The overarching takeaway is that, without explicitly capturing higher-order exposures, regulators and supervisors are flying blind. That leaves them ill-equipped to assess the resilience of the financial system they oversee, and ill-prepared to respond to crises once they inevitably materialize.

As exposures only consider the losses that follow from the (idiosyncratic) default of an institution, we developed a more holistic measure of financial stability in chapter 4, which studies the inherent tendency of a financial system to dampen or amplify shocks without reference to any stress scenario. The key insight is that by describing the interactions of liquidity and valuation shocks, our method can capture multiple contagion mechanisms and their interactions in a duplex network, consisting of a liquidity and a valuation shock layer. Although the model is simple, with only a few parameters, it is powerful enough to derive a wide range of insights about the stability of financial systems. Using this approach, we have shown that a feedback loop between liquidity and valuation shocks always exists. When omitting the interaction between liquidity and valuation shocks, this feedback loop is ignored and stability is overestimated, sometimes dramatically so. Since most studies focus on a single type of shocks, financial instabilities may be structurally underestimated.

In chapter 5, we have applied the developed eigenvalue-based approach to the data of the South-African financial system to study liquidity spirals. We have introduced the concept of complex liquidity spirals, which consist of various interacting contagion channels and/or types of institutions. To accurately assess these spirals, models are required that can take the interactions between multiple types of contagion and institutions into account. We have shown that the stability of the system depends greatly on how institutions choose to respond to liquidity shocks, and that to understand stability, both the individual sectors' contribution as well as their collective impact on stability should be assessed. We have explored how stability is affected by liquidity shocks of various sizes and found that, depending on how institutions choose to respond to liquidity shocks, the system may be "robust-yet-fragile"; that is, the system may be highly resilient to small liquidity shocks, but a single, large liquidity shock may cause the system to become highly unstable. This highlights the importance of stability measures that assess a system's resilience to a wide range of shocks, such as the eigenvalue-based approach developed here.

The specific limitations and opportunities for future research of each analysis presented in this thesis were discussed at the end of the corresponding chapters. Here, we conclude with a general discussion of the most important opportunities for how future research could expand and build upon the presented analyses. In chapter 3, we have shown that traditional exposure measures insufficiently capture systemic risk. Although systemic risk has become more widely recognized since the Great Financial Crisis, many areas that do not (sufficiently) consider systemic risk (but should) still exist. A number of these areas were already identified in the previous chapters, but this list is not exhaustive. Furthermore, in chapter 4, we have shown that the levels of leverage common in financial systems must necessarily be offset by damping to maintain stability. As most financial systems are (seemingly) stable, this suggests that damping mechanisms are present in these systems and, in fact, that they play an important role in stabilizing these systems. Although this insight follows straightforwardly from the dynamical systems approach taken in chapter 4, the mechanisms behind the damping of financial shocks in complex systems have received little attention in the literature so far. This thesis has laid (some of) the foundations, but the conditions under which financial shocks are absorbed, and the interplay of the amplification and damping of financial shocks in general, should be further investigated.

Chapters 4 and 5 have both contributed to the understanding of how liquidity and solvency interact and that they should be studied in tandem rather that individually. Academics and policymakers should strive to further integrate liquidity and solvency in all models where financial stability matters. Moreover, chapters 3 and 5 have demonstrated that risks may be underestimated when sectors of the financial system are studied in isolation from one another. This implies that the global financial system, as well as the global economy within which the financial system operates, should be modeled as a single complex system to assess emergent risks accurately. Although the tools and data sets required for such models do not exist today, nor are likely to become available in the near future, this should nevertheless be the end goal; models will never rival Laplace's Demon, yet should unceasingly progress towards approximating it more and more closely.

I conclude by expressing my particular interest in employing the measures of financial stability proposed in this thesis to capture *climate risk*. Climate change will continue to pose a significant risk to the global economy in both the near and distant future. Financial stability models that consider the negative externalities of carbon emissions both serve to prepare for the materialization of these risks, as well as to accelerate industrial decarbonization by quantifying the future costs of emissions more accurately. I hope to pursue these interests upon completing my degree.

Appendix A

Higher-Order Exposures

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

Table A.1: **Notation**

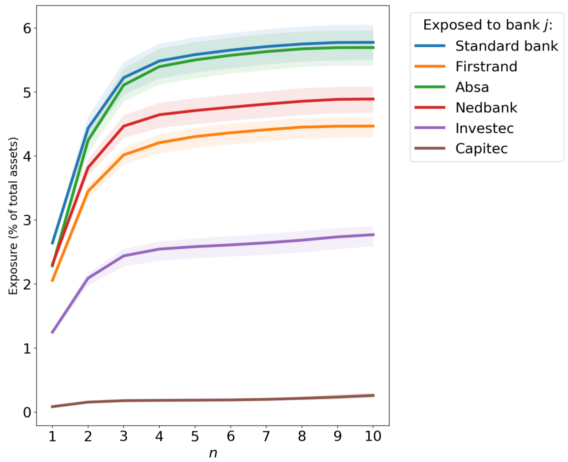

(a) Baseline exposure up to n^{th} order

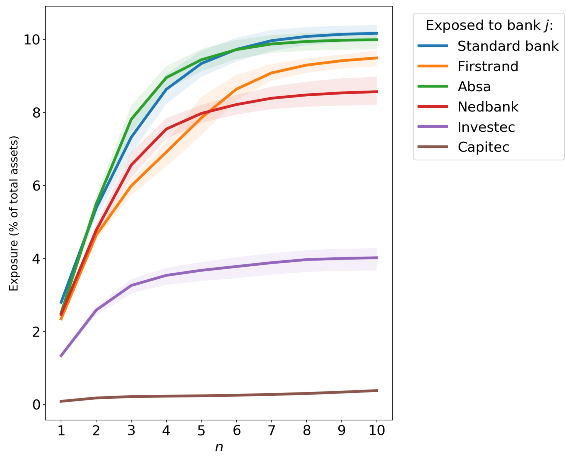

(b) Stressed exposure up to n^{th} order

Figure A.1: **25^{th} and 75^{th} percentiles of exposures of the South African financial system to the six largest banks.** We plot exposure (as % of the system's total assets) up to n^{th} order of the South African financial system to the default of bank j, where $n \leq 10$, j is one of the six large banks and the system's exposure is the sum of the banks' and funds' exposures. (a) shows the baseline exposures and (b) the exposures when institutions are subjected to the stress scenario, which consists of a 25% reduction in all institutions' buffers and a 50% reduction in the liquidity of all tradable assets. We find a distribution of exposures over the 1000 realized samples of the reconstructed interbank network. As we do not know the true interbank network, the true exposures may lie anywhere within this distribution. We plot the mean of the distribution as a solid line and the area between the 25^{th} and 75^{th} percentiles of the distributions as a shaded region in the same color as the mean. Plot (a) shows that baseline exposures level out around $n = 5$ and that the distribution fans out as the order of the exposure increases, which is to be expected because inaccuracies compound. However, (b) shows that stressed exposures to Standard bank, Absa and Nedbank only level out by $n \approx 8$, which suggests that the approach taken in this paper of calculating exposures up to fifth order may substantially underestimate exposures under stressed conditions.

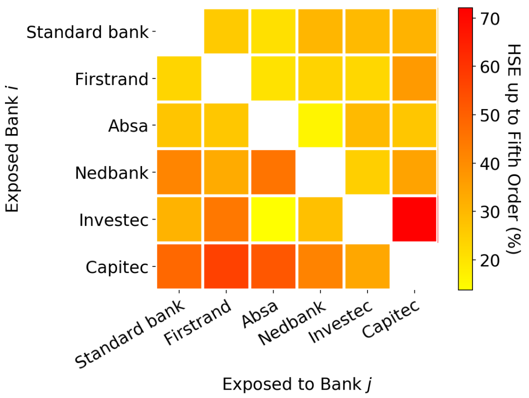

(a) HSE up to fifth-order

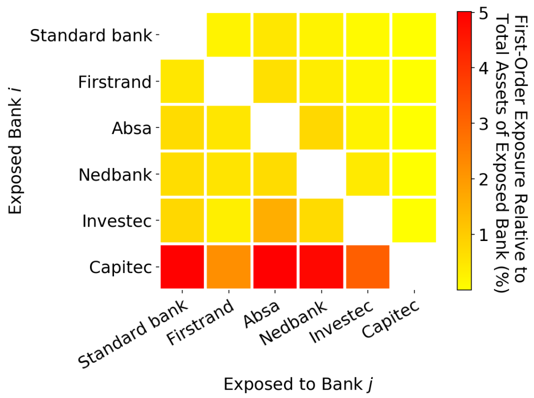

(b) First-order exposure

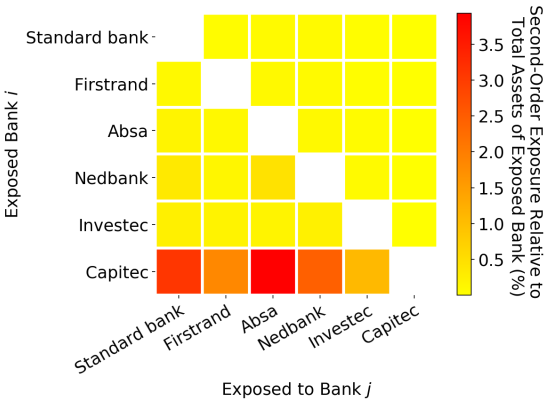

(c) Second-order exposure

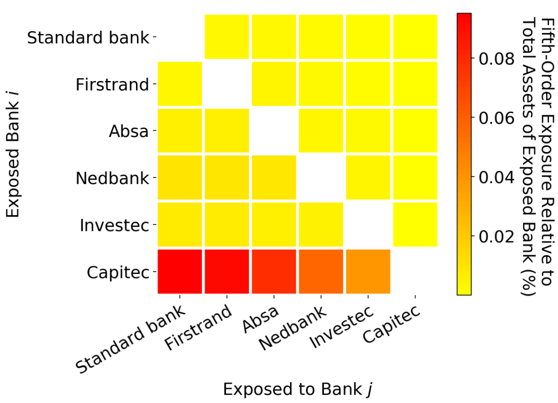

(d) Fifth-order exposure

Figure A.2: **Individual banks' exposures to the six largest banks**. The figure shows the exposures (as % of each bank's total assets) and HSEs between the six large banks. The banks are ordered by total asset size (descending from top to bottom) and the plots use a color gradient to indicate the HSE or exposure of one bank on the vertical axis to another bank on the horizontal axis. (Note that the banks on the horizontal axis are ordered from left to right by descending total asset size.) In general, figures (b)-(d) show that the banks have very modest exposures between them. Other than Capitec, the banks have both low first-order and low higher-order exposures. Yet, figure (a) highlights a few cases where the exposures' HSEs are substantially smaller or larger than the average. Hence, the higher-order exposures cannot be proxied by "scaled" first-order exposures in general.

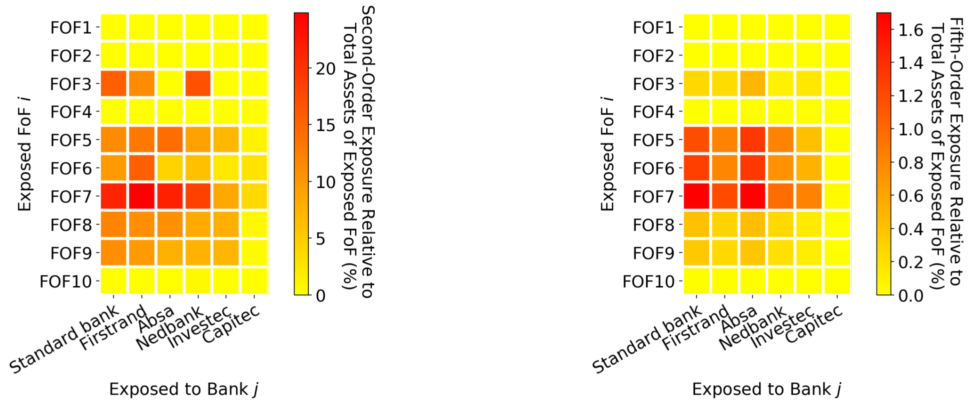

(a) Second-order exposure

(b) Fifth-order exposure

Figure A.3: **Individual FoFs' exposures to the six largest banks**. The figure shows the second-order exposures in (a) and fifth-order exposures in (b) of the ten largest FoFs, as % of each FoF's total assets. The FoFs are ordered by total asset size (descending from top to bottom) and the plots use a color gradient to indicate the exposure of a FoF on the vertical axis to a bank on the horizontal axis. (Note that the banks are ordered from left to right by descending total asset size.) We do not show the FoFs' HSEs or first-order exposures, as the first-order exposures are all zero and, consequently, the HSEs are all equal to one hundred percent. The second and fifth-order exposures show strong variation across the funds. Hence, the FoFs' exposures cannot be modelled at the sectoral level but must be modelled explicitly for individual FoFs .

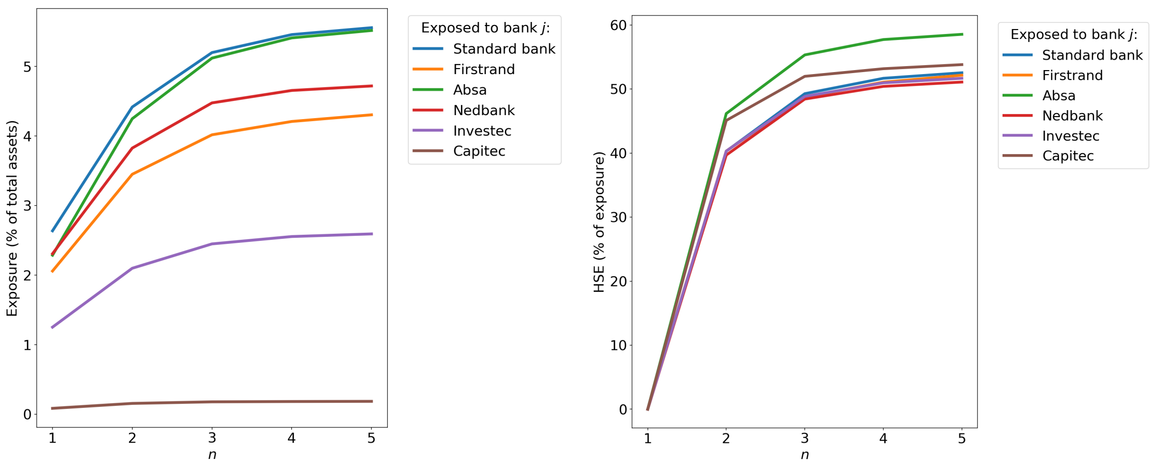

(a) Exposure up to n^{th} order (b) HSE up to n^{th} order

Figure A.4: **Exposure and higher-order share of exposure (HSE) up to n^{th} order of the South African financial system to the six largest banks when ignoring defaulted institutions' losses.** We reproduce the results in figure 3.7 under the assumption that institutions do not suffer losses after they default, and that the securities sold upon default do not cause any subsequent contagion after the sale. The figure shows that this assumption does not impact our results significantly.

Appendix B

Scenario-Free Analysis of Financial Stability with Interacting Contagion Channels

Notation	Description
E	Equity
D	Debt
λ	Leverage
$\hat{\lambda}$	Critical leverage
ν	Largest eigenvalue
A	Shock Transmission Matrix
$\vec{x}$	Shock vector
$\vec{x}^l$	Liquidity shocks vector
$\vec{x}^v$	Valuation shocks vector
S	Short-term loan
L	Long-term loan
n_s	Total number of shares in security s in circulation
μ_s	Price-impact factor for security s
p_s	Price of security s
δ_i	Risk-adjustment factor for institution i
ϕ_l	Fraction of (institutions that are) liquidity sinks
ϕ_v	Fraction of (institutions that are) valuation sinks
F	Fraction of short-term lenders (institutions that provide short-term loans)
Λ	Fraction of leverage targeters (institutions that are leverage targeting)
C_s	Market capitalization of security s
N	Number of institutions
N^v	Number of leveraged institutions
N^w	Number of distinct securities
N^s	Number of blocks of shares of security s
N^d	Number of debts (loans)
N^s_i	Number of blocks of security s received by institution i
N^d_i	Number of loans received by institution i
N^d_{ij}	Number of loans from institution i to institution j

Table B.1: Notation

B.1 Validation Tests

In Figure B.1, we compare the overestimation of the critical leverage of the mean-field model to the overestimation in randomly generated financial systems. The overestimation is calculated as the percentage increase from the true critical leverage to the counterparty risk critical leverage (i.e. the critical leverage found when considering only pure counterparty risk contagion). The financial systems are generated using the algorithm outlined in section 4.4.1 and system parameters derived in section 4.4.2. The systems include $N = 100$ institutions, such that the fraction of leverage targeting institutions Λ can be increased in 100 increments. The figure shows that the overestimation becomes arbitrarily large as $\Lambda \to 1$ for any $\phi_l < 1$, and that the overestimation in the mean-field model closely approximates the overestimation in the randomly generated financial systems.

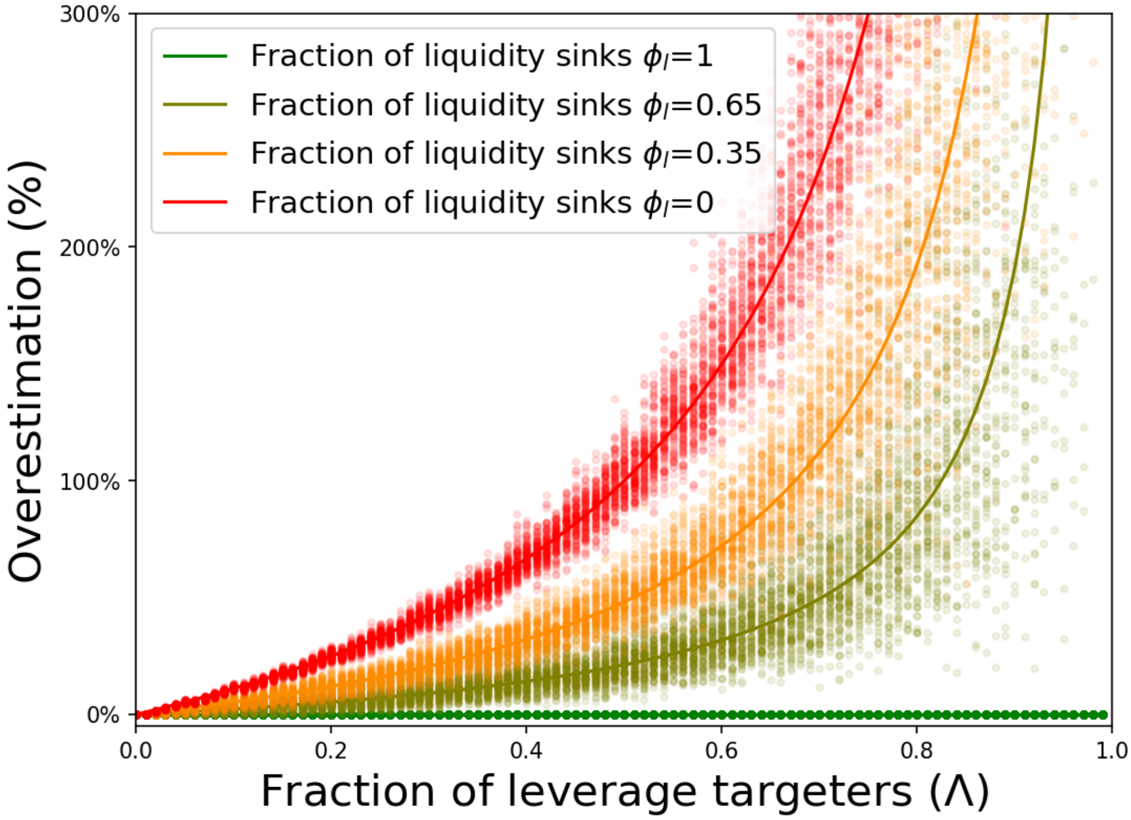

Figure B.1: **Overestimation of the Critical Leverage.** Comparison of the percentage increase from the true critical leverage to the counterparty risk critical leverage of randomly generated systems (dots) and the mean-field model (solid lines). Fixed parameters: $F = 0.5$, $\phi_v = 0.2$, $\mu_s = .1$, $\delta_i = .1$, $N = 100$, $N^d = 10$, $N^s = 100$, and $N^w = 10$.

Figure B.2 compares the mean-field critical leverage to the critical leverages of randomly generated financial systems. The financial systems are generated using the algorithm outlined in section 4.4.1, but with various modifications to the algorithm that generate additional heterogeneity. We use the system parameters derived in section 4.4.2 and set the number of loans and blocks of securities $N^d = N^s = 10$ to generate sparse financial systems. Each labeled column of the figure presents the critical leverages of systems that were generated with a single modification to the algorithm:

- For systems in the column "Pareto Weights", the weight of each generated edge is drawn from a Pareto distribution with shape equal to two and scale equal to one, to create additional heterogeneity in the distribution of edge weights.

- For systems in the column "Pareto In-Degree", the number of loans received by each institution is drawn from a Pareto distribution with shape equal to two and scale equal to one (rounded down and limited to 100 for computational efficiency), to create additional heterogeneity in the distribution of institutions' in-degrees.

- For systems in the column "Pareto Out-Degree", the number of loans made by each institution is drawn from a Pareto distribution with shape equal to two and scale equal to one (rounded down and limited to 100 for computational efficiency), to create additional heterogeneity in the distribution of institutions' out-degrees.

- For systems in the column "In Core-Periphery", the set of non-sink institutions is divided into two halves; one half is excluded from receiving loans when each institution makes its N^d loans, to create a core of institutions that receive the majority of loans. Note that institutions designated as leveraged that do not receive any loans, are allocated a single random loan at the end of the algorithm (see section 4.4.1).

- For systems in the column "Out Core-Periphery", the set of non-sink institutions is divided into two halves; one half is excluded from making loans when each institution makes its N^d loans, to create a core of institutions that make the majority of loans. Institutions designated as short-term lenders that are excluded from making loans are allocated a single loan to a randomly chosen institution.

Comparison with Figure 4.4 shows that the additional sources of heterogeneity increase the variation in critical leverages in Figure B.2, in particular when institutions' in- or out degrees are drawn from a Pareto distribution.

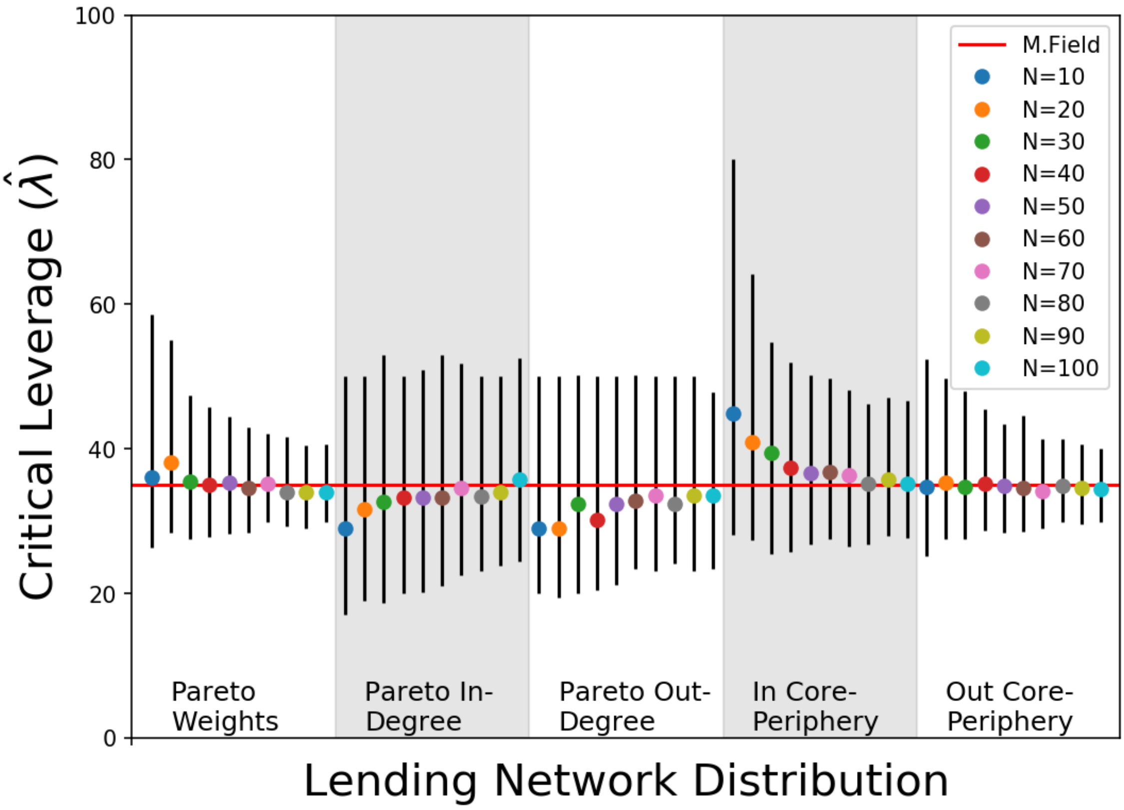

Figure B.2: **Generating Financial Systems with Additional Sources of Heterogeneity.** Comparison of the mean-field critical leverage (red line) to the critical leverages of randomly generated financial systems. For various modifications to the generation algorithm (which introduce additional heterogeneity), we generate 500 random systems and plot the median (colored dot) and the 15^{th} to 85^{th} percentile interval (black bars) of the distribution of critical leverages. Fixed parameters: $F = 0.5$, $\phi_v = 0.2$, $\mu_s = .1$, $\delta_i = .1$, $N = 100$, $N^d = 10$, $N^s = 10$, and $N^w = 10$.

B.2 Eigenvalue Time Dependence

We show that the shock transmission matrix' largest eigenvalue, when it is equal to one, is independent of the relative speeds at which the various channels act. When the largest eigenvalue $\nu = 1$, the corresponding (right) eigenvector $\vec{v}$ is invariant under multiplication by the shock transmission matrix,

$$\vec{v}_{t+1} = A\vec{v}_t = \nu\vec{v}_t = \vec{v}_t. \tag{B.1}$$

Hence, for any element v_k (corresponding to the network's k^{th} node) of the eigenvector $\vec{v}$, we have that

$$v_{k,t+1} = v_{k,t}. \tag{B.2}$$

From the matrix-vector product, we know that

$$v_{k,t+1} = \sum_{y \in \mathcal{A}} w_{yk} v_{y,t}, \tag{B.3}$$

where $\mathcal{A}$ denotes the network's set of nodes and $w_{yk} = A_{ky}$ denotes the weight of the edge from node y to node k. Hence,

$$v_{k,t} = \sum_{y \in \mathcal{A}} w_{yk} v_{y,t}. \tag{B.4}$$

Let us now take a specific pair of nodes i, k. We add a "dummy node" *between* nodes i and k, by replacing the edge from i to k by an edge with identical weight from j to k and adding an edge with weight equal to one from node i to node j. Using $\hat{w}_{xy}$ to denote edges in the new network, we have that $\hat{w}_{ij} = 1$, $\hat{w}_{ik} = 0$, $\hat{w}_{jk} = w_{ik}$, and $\hat{w}_{yx} = w_{yx}$ for any pair of nodes yx other than the pairs ij, ik and jk. Hence, any shock that was previously transmitted directly from i to k is now delayed by one iteration before arriving at node k, while the shock's magnitude is unaffected.

Compared to $\vec{v}$, the new network's eigenvector $\hat{\vec{v}}$ has an additional entry, $\hat{v}_j$. For all $y \neq j$ we set $\hat{v}_{y,t} = v_{y,t}$, and since $\hat{v}_{j,t+1} = \hat{w}_{ij}\hat{v}_{i,t} = \hat{v}_{i,t}$, we set $\hat{v}_{j,t} = v_{i,t}$. Using $\hat{\mathcal{A}}$ to denote the new network's set of nodes, we immediately see that

$$\hat{v}_{k,t+1} = \sum_{y \in \hat{\mathcal{A}}} \hat{w}_{yk} \hat{v}_{y,t} = \sum_{y \in \mathcal{A}} w_{yk} v_{y,t} = v_{k,t} = \hat{v}_{k,t}, \tag{B.5}$$

so the invariance in equation (B.2) is conserved (as well as the invariance of the entire eigenvector (B.1), as the rest of the network remains unchanged). Hence, when the largest eigenvalue is equal to one, we can add a dummy node that slows down shock transmission without affecting the invariance of the corresponding eigenvector under multiplication by the shock transmission matrix.

Note from Figure 4.2 that each institution is represented by both a liquidity and valuation node, such that liquidity and valuation shocks never travel over the same edges. In principle, we can add any number of dummy nodes to "tune" the relative speeds at which contagion channels operate with any desired granularity, without affecting the network's stability. Thus, when the largest eigenvalue is equal to one, the network's stability is independent of the relative speeds at which contagion channels operate.

B.3 Derivation of the Mean-Field Model

Here, we show that the shock transmission matrix reduces to a 2×2 matrix when passively leveraged institutions have the same risk-adjustment factor $\delta_i = \delta$, all leveraged institutions have the same leverage $\lambda_i = \lambda$, and $N^s/N, N^d/N, N \to \infty$. We denote the number of leveraged institutions as $N^v = (1 - \phi_v)N$ and the number of non-sink institutions as $N^l = (1 - \phi_l)N^v$.
We simplify the notation of the shock transmission matrix in three steps:

1. We reorder the nodes to move all empty columns (corresponding to liquidity sinks' absorption of liquidity shocks and valuation sinks' absorption of valuation shocks) to the right of the matrix. The resulting matrix is lower block triangular and its eigenvalues are given by those of the only non-zero diagonal block, i.e. the upper-left diagonal block. Hence, we can obtain a reduced matrix by removing liquidity sinks' liquidity nodes and valuation sinks' valuation nodes from the matrix without affecting the largest eigenvalue.

2. We reorder the nodes in the reduced matrix obtained in step 1. to move the empty rows that correspond to the transmission of liquidity shocks to valuation sinks to the bottom of the matrix. (Valuation sinks have no leverage target and no short-term debt so do not receive liquidity shocks.) The resulting matrix is upper block triangular and its eigenvalues are given by those of the only non-zero diagonal block, i.e. the upper-left diagonal block. Hence, we can completely remove valuation sinks from the matrix without affecting the largest eigenvalue.

3. For simplicity, we reorder the nodes to move the columns corresponding to liquidity sinks' shock transmission in response to valuation shocks to the right end of the matrix (i.e. liquidity sinks' valuation nodes are moved to the bottom of the shock vector).

We refer to resulting matrix as the *simplified* shock transmission matrix. The liquidity shock vector $\vec{x}^l$ is of length N^l and the valuation shock vector $\vec{x}^v$ is of length N^v. Hence, the dimensions of the simplified shock transmission matrix' funding contagion quadrant are $N^l \times N^l$, of the counterparty risk contagion quadrant $N^v \times N^v$, of the overlapping portfolio contagion quadrant $N^v \times N^l$, and of the leverage targeting contagion quadrant $N^l \times N^v$.

B.3.1 Simplified Shock Transmission Matrix

We first derive the simplified shock transmission matrix to which systems converge as $N^s/N, N^d/N, N \to \infty$, after which we discuss how to reduce the resulting shock transmission matrix to the 2×2 matrix.

- When $N^s/N \to \infty$ for all securities s, then for any security s, all institutions hold an equal fraction of the security's market capitalization,

$$\lim_{N^s/N \to \infty} \frac{N_i^s}{N^s} = \mathbb{E}\left(\frac{N_i^s}{N^s}\right) = \frac{1}{N}, \tag{B.6}$$

  where $\mathbb{E}(\dots)$ denotes the expectation. All non-sink institutions that do not provide short-term lending have security $\hat{s}$ at the top of their pecking order, which is the most liquid security among the N^w distinct securities. The price-impact factor of security $\hat{s}$ is denoted as μ (i.e. without any subscript).

- When $N^d/N, N \to \infty$, each leveraged institution's debt is distributed equally over all $N-1$ other institutions,

$$\lim_{N^d/N, N \to \infty} \frac{N_{ji}^d}{N_i^d} = \lim_{N \to \infty} \mathbb{E}\left(\frac{N_{ji}^d}{N_i^d}\right) = \lim_{N \to \infty} \frac{1}{N-1} = \frac{1}{N}. \tag{B.7}$$

The distribution of equities E_i that solves the balance sheet identity (4.13) and equations (B.6) and (B.7) is $E_i\,(1+\lambda_i) \;=\; E_j\,(1+\lambda_j)$ for any institutions i and j, where $\lambda_i = 0$ if institution i is a valuation sink and $\lambda_i = \lambda$ otherwise. That is, all leveraged institutions have the same equity and debt, and all valuation sinks have the same equity, which is equal to the sum of the equity and debt of any leveraged institution.

As equations (B.6) and (B.7) tell us how securities and total debt are distributed (and hence how overlapping portfolio contagion and counterparty risk contagion are distributed), let us now consider how short-term lending is distributed.

When $N^d/N, N \to \infty$, the fraction of short-term lender i's total short-term lending S_i provided to any leveraged institution j is equal to

$$\begin{aligned}
\lim_{N^d/N, N \to \infty} \frac{S_{ij}}{S_i} &= \lim_{N^d/N, N \to \infty} \frac{\frac{N_{ij}^d}{N_j^d} D_j}{\sum_{k=1}^{N^v} \frac{N_{ik}^d}{N_k^d} D_k} = \lim_{N^d/N, N \to \infty} \frac{\frac{N_{ij}^d}{N_j^d}}{\sum_{k=1}^{N^v} \frac{N_{ik}^d}{N_k^d}} \\
&= \lim_{N^d/N, N \to \infty} \frac{N_{ij}^d}{\sum_{k=1}^{N^v} N_{ik}^d} = \lim_{N^d/N, N \to \infty} \frac{\frac{N^d}{N^v-1}}{N^d} = \lim_{N \to \infty} \frac{1}{N^v - 1} = \frac{1}{N^v},
\end{aligned} \tag{B.8}$$

where k runs over all leveraged institutions, and we have used that:

- When $N^s/N, N^d/N, N \to \infty$, all N^v leveraged institutions have the same debt D (as discussed above).
- $\lim_{N^d/N\to\infty} N_k^d = \mathbb{E}\left(N_k^d\right) = \frac{N^d N}{N^v}$ is the same for all leveraged institutions k (including institution j).
- The number of loans any leveraged institution i provides to another leveraged institution j is $\lim_{N^d/N\to\infty} N_{ij}^d = \mathbb{E}\left(N_{ij}^d\right) = \frac{N^d}{N^v-1}$, as institution i cannot lend to itself.

From equation (B.8) follows that the funding contagion transmission of a (non-sink) short-term lender, as given by the corresponding column in the simplified shock transmission matrix, is equal to

$$\left[\frac{1}{N^v}, \ldots, \frac{1}{N^v}, 0, \frac{1}{N^v}, \ldots, \frac{1}{N^v}, 0, \ldots, 0\right]^T, \tag{B.9}$$

where for institution i, the i^{th} entry is zero (no lending to itself) and the last N^v terms are zero, which corresponds to the institution's non-transmission of overlapping portfolio contagion.
When $N \to \infty$, the institutions become a continuum and shock transmission to individual institutions vanishes. Hence, for $N \to \infty$, the funding contagion transmission vector reduces to[1]

$$\left[\frac{1}{N^v}, \ldots, \frac{1}{N^v}, 0, \ldots, 0\right]^T, \tag{B.10}$$

such that each (non-sink) short-term lender's funding contagion is distributed homogeneously over the continuum of leveraged institutions.
From equation (B.7), it follows that for each passively leveraged institution, the counterparty risk contagion transmission as given by the corresponding column of the simplified shock transmission matrix is equal to

$$\left[0, \ldots, 0, \frac{\delta\lambda}{N}, \ldots, \frac{\delta\lambda}{N}, 0, \frac{\delta\lambda}{N}, \ldots, \frac{\delta\lambda}{N}\right]^T, \tag{B.11}$$

[1] Formally, the difference between B.9 and B.10 vanishes in the limit $N \to \infty$; $\lim_{N\to\infty} \left\| \left[\frac{1}{N^v}, \ldots, \frac{1}{N^v}, 0, \ldots, 0\right]^T - \left[\frac{1}{N^v-1}, \ldots, \frac{1}{N^v-1}, 0, \frac{1}{N^v-1}, \ldots, \frac{1}{N^v-1}, 0, \ldots, 0\right]^T \right\| = 0.$

where for institution i, the i^{th} entry is zero (no debt to itself), and the first N^l entries of the vector are zero, which corresponds to the institution's non-transmission of leverage targeting contagion.

Similar to funding contagion, when $N \to \infty$, the counterparty risk contagion transmission vector reduces to

$$\left[0, \ldots, 0, \frac{\delta\lambda}{N}, \ldots, \frac{\delta\lambda}{N}\right]^T, \tag{B.12}$$

such that each passively leveraged institution's counterparty risk contagion is distributed homogeneously over the continuum of institutions.

Lastly, from equation (B.6), the overlapping portfolio contagion shock transmission vector for any non-sink institution that does not provide short-term lending is equal to

$$\left[0, \ldots, 0, \frac{\mu}{N}, \ldots, \frac{\mu}{N}\right]^T, \tag{B.13}$$

where the first N^l terms are zero, which corresponds to the institution's non-transmission of funding contagion. Hence, the overlapping portfolio contagion transmitted by any non-sink institution that does not provide short-term lending is distributed homogeneously over the continuum of institutions.

From equations (B.10), (B.12) and (B.13), we find that for $N^s/N, N^d/N, N \to \infty$, the simplified shock transmission matrix is given by

$$\left[\begin{array}{ccc||ccccccc}
I_1^f \frac{1}{N^v} & \cdots & I_{N^l}^f \frac{1}{N^v} & I_1^\lambda \lambda & 0 & \cdots & 0 & 0 & \cdots & 0 \\
\vdots & \ddots & \vdots & 0 & I_2^\lambda \lambda & \ddots & \vdots & \vdots & \ddots & \vdots \\
\vdots & \ddots & \vdots & \vdots & \ddots & \ddots & 0 & \vdots & \ddots & \vdots \\
I_1^f \frac{1}{N^v} & \cdots & I_{N^l}^f \frac{1}{N^v} & 0 & \cdots & 0 & I_{N^l}^\lambda \lambda & 0 & \cdots & 0 \\
\hline\hline
\left(1 - I_1^f\right) \frac{\mu}{N} & \cdots & \left(1 - I_{N^l}^f\right) \frac{\mu}{N} & (1 - I_1^\lambda) \frac{\delta\lambda}{N} & \cdots & \cdots & \left(1 - I_{N^l}^\lambda\right) \frac{\delta\lambda}{N} & \cdots & \cdots & (1 - I_{N^v}^\lambda) \frac{\delta\lambda}{N} \\
\vdots & \ddots & \vdots & \vdots & \ddots & \ddots & \vdots & \ddots & \ddots & \vdots \\
\left(1 - I_1^f\right) \frac{\mu}{N} & \cdots & \left(1 - I_{N^l}^f\right) \frac{\mu}{N} & (1 - I_1^\lambda) \frac{\delta\lambda}{N} & \cdots & \cdots & \left(1 - I_{N^l}^\lambda\right) \frac{\delta\lambda}{N} & \cdots & \cdots & (1 - I_{N^v}^\lambda) \frac{\delta\lambda}{N}
\end{array}\right]. \tag{B.14}$$

where $I_i^f = 1$ if institution i is a short-term lender and $I_i^f = 0$ otherwise, and $I_i^\lambda = 1$ if institution i has a leverage target and $I_i^\lambda = 0$ otherwise.

B.3.2 Reduction to 2×2 Matrix

We now show that the simplified shock transmission matrix in equation (B.14) can be reduced to a 2×2 matrix with largest eigenvalue identical to that of the shock transmission matrix in equation (B.14). We do so by showing that system is uniquely determined by the dynamics of the aggregate liquidity and valuation shocks x_t^l and x_t^v.

Let $x_{t,i}^l$ denote the liquidity shock received by (non-sink) institution i at time t, such that $\vec{x}_t^l = \left[x_{t,1}^l, \ldots, x_{t,N^l}^l\right]^T$ and let $x_{t,j}^v$ be the valuation shock received by (leveraged) institution j at time t, such that $\vec{x}_t^v = \left[x_{t,1}^v, \ldots, x_{t,N^v}^v\right]^T$. Furthermore, let $x_t^l = \sum_{i=1}^{N^l} x_{t,i}^l$ be the aggregate liquidity shock received by all non-sink institutions at time t and $x_t^v = \sum_{i=1}^{N^v} x_{t,i}^v$ be the aggregate valuation shock received by all leveraged institutions at time t. Lastly, at time t, let the fraction of the aggregate liquidity shock x_t^l received by non-sink institutions with short-term lending be denoted as $\hat{F}_t = \sum_{i=1}^{N^l} I_i^f x_{t,i}^l / x_t^l$, such that $(1 - \hat{F}_t) = \sum_{i=1}^{N^l} \left(1 - I_i^f\right) x_{t,i}^l / x_t^l$, and let the fraction of the aggregate valuation shock x_t^v received by leverage targeting institutions be denoted as $\hat{\Lambda}_t = \sum_{i=1}^{N^v} I_i^\lambda x_{t,i}^v / x_t^v$, such that $(1 - \hat{\Lambda}_t) = \sum_{i=1}^{N^v} \left(1 - I_i^\lambda\right) x_{t,i}^v / x_t^v$.
We use the following properties throughout the derivation:

$$\frac{N^l}{N^v} = 1 - \phi_l, \tag{B.15}$$

$$\frac{\sum_{i=1}^{N^v} I_i^\lambda}{N^v} = \Lambda, \tag{B.16}$$

$$\lim_{N\to\infty} \frac{\sum_{i=1}^{N^l} I_i^f}{N^l} = F, \tag{B.17}$$

$$\lim_{N\to\infty} \frac{\sum_{i=1}^{N^l} I_i^f}{N^v} = \lim_{N\to\infty} \frac{(1-\phi_l)\sum_{i=1}^{N^l} I_i^f}{N^l} = (1-\phi_l)F, \tag{B.18}$$

$$\lim_{N\to\infty} \frac{\sum_{i=1}^{N^l} I_i^\lambda}{N^v} = \lim_{N\to\infty} \frac{(1-\phi_l)\sum_{i=1}^{N^l} I_i^\lambda}{N^l} = (1-\phi_l)\Lambda, \tag{B.19}$$

$$\lim_{N\to\infty} \sum_{i=1}^{N^l} I_i^f I_i^\lambda = \sum_i^{N^l} \mathbb{E}\left(I_i^f I_i^\lambda\right) = \sum_i^{N^l} \mathbb{E}\left(I_i^f\right)\mathbb{E}\left(I_i^\lambda\right) = F\Lambda \sum_i^{N^l} 1 = F\Lambda N^l. \tag{B.20}$$

The first identity is simply a restatement of the fact that we designate a fraction ϕ_l of leveraged institution as liquidity sinks. The second is a restatement of the definition that the fraction of leverage targeters Λ is equal to the fraction of leveraged institutions that have a leverage target. The third and fourth identities use that as $N \to \infty$, the fraction of institutions that provide short-term lending is equal to the fraction of non-sink institutions that provide short-term lending (because sinks and short-term lenders are designated independently). Similarly, the fifth identity uses that as $N \to \infty$, the fraction of non-sink institutions that have a leverage target is equal to Λ (because liquidity sinks and leverage strategies are designated independently). Lastly, the sixth identity gives the number of non-sink short-term lenders

that are leverage targeting and follows from the fact that short-term lenders and leverage targeting institutions are designated independently.
Plugging the simplified shock transmission matrix in equation (B.14) into equation (4.2) yields

$$\vec{x}_{t+1} = A\vec{x}_t = A\begin{bmatrix}\vec{x}_t^l\\ \vec{x}_t^v\end{bmatrix} = \begin{bmatrix} \frac{1}{N^v}\sum_{i=1}^{N^l} I_i^f x_{t,i}^l + \lambda I_1^\lambda x_{t,1}^v \\ \vdots \\ \frac{1}{N^v}\sum_{i=1}^{N^l} I_i^f x_{t,i}^l + \lambda I_{N^l}^\lambda x_{t,N^l}^v \\ \frac{\mu}{N}\sum_{i=1}^{N^l}\left(1-I_i^f\right)x_{t,i}^l + \frac{\delta\lambda}{N}\sum_{j=1}^{N^v}\left(1-I_j^\lambda\right)x_{t,j}^v \\ \vdots \\ \frac{\mu}{N}\sum_{i=1}^{N^l}\left(1-I_i^f\right)x_{t,i}^l + \frac{\delta\lambda}{N}\sum_{j=1}^{N^v}\left(1-I_j^\lambda\right)x_{t,j}^v \end{bmatrix} \tag{B.21}$$

$$= \begin{bmatrix} \frac{1}{N^v}\hat{F}_t x_t^l + \lambda I_1^\lambda x_{t,1}^v \\ \vdots \\ \frac{1}{N^v}\hat{F}_t x^l + \lambda I_{N^l}^\lambda x_{t,N^l}^v \\ \frac{\mu}{N}(1-\hat{F}_t)x_t^l + \frac{\delta\lambda}{N}(1-\hat{\Lambda}_t)x_t^v \\ \vdots \\ \frac{\mu}{N}(1-\hat{F}_t)x_t^l + \frac{\delta\lambda}{N}(1-\hat{\Lambda}_t)x_t^v \end{bmatrix}. \tag{B.22}$$

Hence, at time $t+1$, we have for any non-sink institution i that

$$x_{t+1,i}^l = \frac{\hat{F}_t}{N^v}x_t^l + \lambda I_i^\lambda x_{t,i}^v, \tag{B.23}$$

and for any leveraged institution i that

$$x_{t+1,i}^v = \frac{(1-\hat{F}_t)\mu}{N}x_t^l + \frac{(1-\hat{\Lambda}_t)\delta\lambda}{N}x_t^v. \tag{B.24}$$

Equation (B.24) shows that $x_{t+1,i}^v$ is the same for any (leveraged) institution i, because the right-hand side of equation (B.24) does not depend on i. The aggregate valuation shock at time $t+1$ is given by

$$x_{t+1}^v = \sum_{i=1}^{N^v} x_{t+1,i}^v = (1-\hat{F}_t)(1-\phi_v)\mu x_t^l + (1-\hat{\Lambda}_t)(1-\phi_v)\delta\lambda x_t^v, \tag{B.25}$$

and hence we have for any leveraged institution i that

$$x_{t+1,i}^v = \frac{x_{t+1}^v}{N^v}. \tag{B.26}$$

From $x_{t+1,i}^v = x_{t+1}^v/N^v$ follows that

$$\hat{\Lambda}_{t+1} = \frac{\sum_{i=1}^{N^v} I_i^\lambda x_{t+1,i}^v}{x_{t+1}^v} = \frac{\sum_{i=1}^{N^v} I_i^\lambda \frac{x_{t+1}^v}{N^v}}{x_{t+1}^v} = \frac{\sum_{i=1}^{N^v} I_i^\lambda}{N^v} = \Lambda. \tag{B.27}$$

Furthermore, from $x_{t+1,i}^v = x_{t+1}^v / N^v$ also follows that, at time $t+2$,

$$x_{t+2,i}^l = \frac{\hat{F}_{t+1}}{N^v} x_{t+1}^l + \frac{\lambda I_i^\lambda}{N^v} x_{t+1}^v, \tag{B.28}$$

$$x_{t+2}^l = \hat{F}_{t+1}(1-\phi_l) x_{t+1}^l + \lambda x_{t+1}^v \sum_{i=1}^{N^l} \frac{I_i^\lambda}{N^v} = \hat{F}_{t+1}(1-\phi_l) x_{t+1}^l + \Lambda(1-\phi_l)\lambda x_{t+1}^v. \tag{B.29}$$

Using equations (B.28) and (B.29), we find that

$$\hat{F}_{t+2} = \frac{\sum_{i=1}^{N^l} I_i^f x_{t+2,i}^l}{x_{t+2}^l} = \frac{\sum_{i=1}^{N^l} I_i^f \left(\frac{\hat{F}_{t+1}}{N^v} x_{t+1}^l + \frac{\lambda I_i^\lambda}{N^v} x_{t+1}^v \right)}{\hat{F}_{t+1}(1-\phi_l) x_{t+1}^l + \Lambda(1-\phi_l)\lambda x_{t+1}^v} \tag{B.30}$$

$$= \frac{\hat{F}_{t+1} x_{t+1}^l \frac{\sum_{i=1}^{N^l} I_i^f}{N^v} + \lambda x_{t+1}^v \frac{\sum_{i=1}^{N^l} I_i^f I_i^\lambda}{N^v}}{\hat{F}_{t+1}(1-\phi_l) x_{t+1}^l + \Lambda(1-\phi_l)\lambda x_{t+1}^v} \tag{B.31}$$

$$= \frac{F\left(\hat{F}_{t+1}(1-\phi_l) x_{t+1}^l + \Lambda(1-\phi_l)\lambda x_{t+1}^v \right)}{\hat{F}_{t+1}(1-\phi_l) x_{t+1}^l + \Lambda(1-\phi_l)\lambda x_{t+1}^v} = F. \tag{B.32}$$

Equations (B.24) and (B.28) do not depend on individual shocks but only on the aggregate liquidity and valuation shocks. Therefore, we find that for $t > 1$ (where the initial exogenous shock occurs at $t = 0$)[2], the system's shock propagation is uniquely determined by the dynamics of x_t^l and x_t^v, which we find by plugging $\hat{F}_t = F$ and $\hat{\Lambda}_t = \Lambda$ into equations (B.25) and (B.29):

$$\begin{aligned} x_{t+1}^l &= F(1-\phi_l) x_t^l + \Lambda(1-\phi_l)\lambda x_t^v, \\ x_{t+1}^v &= (1-F)(1-\phi_v)\mu x_t^l + (1-\Lambda)(1-\phi_v)\delta\lambda. \end{aligned} \tag{B.33}$$

Equation (B.33) can be written in matrix form as

[2]This limitation is not an artifact of the derivation but an actual constraint: When $\vec{x}_{t=0}$ consists of a single valuation shock to a leverage targeting institution i, $\vec{x}_{t=1}$ consists of a single liquidity shock to the same institution i, so equation (B.34) does not hold for $t = 0$. Depending on whether or not institution i has made short-term loans, institution i transmits either a pure funding or pure overlapping portfolio contagion shock (so either $x_{t=2}^l = 0$ or $x_{t=2}^v = 0$) so equation (B.34) also does not hold for $t = 1$. Because the funding or overlapping portfolio contagion shock is distributed homogeneously over all institutions, equation (B.34) holds from $t = 2$ onward. However, this $\vec{x}_{t=0}$ is not an eigenvector of A, because $\vec{x}_{t=1}$ is orthogonal to $\vec{x}_{t=0}$. When $\vec{x}_{t=0}$ *is* an eigenvector of A, (B.34) holds for $t = 1$.

$$\begin{bmatrix} x^l_{t+1} \\ x^v_{t+1} \end{bmatrix} = \left[\begin{array}{c||c} F(1-\phi_l) & \Lambda(1-\phi_l)\lambda \\ \hline\hline (1-F)(1-\phi_v)\mu & (1-\Lambda)(1-\phi_v)\delta\lambda \end{array} \right] \begin{bmatrix} x^l_t \\ x^v_t \end{bmatrix}, \tag{B.34}$$

which gives the 2×2 matrix.

Let ν be the largest eigenvalue of the 2×2 matrix and $\vec{v}$ the corresponding eigenvector, such that

$$\hat{A}\vec{v} = \hat{A} \begin{bmatrix} \vec{v}^l \\ \vec{v}^v \end{bmatrix} = \nu \begin{bmatrix} \vec{v}^l \\ \vec{v}^v \end{bmatrix}. \tag{B.35}$$

When we rewrite equations (B.24) and (B.28) as a matrix-vector product and use that $\hat{F}_t = F$ and $\hat{\Lambda}_t = \Lambda$ for $t > 1$, we find that

$$\vec{x}_{t+1} = \begin{bmatrix} \frac{F(1+h)}{N^v} & \frac{\lambda I^\lambda_i}{N^v} \\ \vdots & \vdots \\ \frac{F(1+h)}{N^v} & \frac{\lambda I^\lambda_i}{N^v} \\ \frac{(1-F)\mu}{N} & \frac{(1-\Lambda)\delta\lambda}{N} \\ \vdots & \vdots \\ \frac{(1-F)\mu}{N} & \frac{(1-\Lambda)\delta\lambda}{N} \end{bmatrix} \begin{bmatrix} x^l_t \\ x^v_t \end{bmatrix}, \tag{B.36}$$

and because the matrix is constant, we see that $\vec{x}_t$ grows by ν when $\left[x^l_t, x^v_t\right]^T$ grows by ν. Therefore, ν is also the largest eigenvalue of the (full) shock transmission matrix A.

Appendix C

Complex Liquidity Spirals

C.1 Withdrawal of Fund Shares

When we discussed the share redemption contagion channel in section 5.4.3, 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

$$\Delta NAV_{i,t} = x^v_{i,t}/S_{i,t}, \tag{C.1}$$

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

$$NAV_{i,t+1} = NAV_{i,t} - \Delta NAV_{i,t}, \tag{C.2}$$

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;

$$(NAV_{i,t} - \Delta NAV_{i,t})\,\Delta S_{i,t} = \epsilon_{i,t} R x^v_{i,t}, \tag{C.3}$$

where we have used the factor $\epsilon_{i,t}R$ from equation (5.7) to express the amount of liquidity withdrawn in terms of the loss $x^v_{i,t}$. Using (C.1), we find that

$$\Delta S_{i,t} = \min\left\{\frac{\epsilon_{i,t} R S_{i,t} \Delta NAV_{i,t}}{NAV_{i,t} - \Delta NAV_{i,t}}, \epsilon_{i,t} S_{i,t}\right\}, \tag{C.4}$$

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).

C.2 Price Impact of Number of Shares Sold

When we discussed the overlapping portfolio contagion channel in section 5.4.3, 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 (5.9) reduces to

$$\Delta p_{\sigma,t} = \frac{x^l_{i,t}}{D_\sigma} = \frac{(p_{\sigma,t} - \Delta p_{\sigma,t})\,\Delta n_{\sigma i,t}}{D_\sigma}, \tag{C.5}$$

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 (C.5), we find the price impact as a function of the number of shares sold:

$$\Delta p_{\sigma,t} = \frac{p_{\sigma,t}\Delta n_{\sigma i,t}}{D_\sigma + \Delta n_{\sigma i,t}}, \tag{C.6}$$

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 (C.6) that $\Delta p_{\sigma,t} < p_{\sigma,t}$ so the price cannot become negative.

C.3 Market Price of Listed Equity Shares

The overlapping portfolio and shareholder contagion mechanisms derived in section 5.4.3 should not drive the market price of a listed equity share below zero, as the shares are subject to limited liability. This is guaranteed when the contagion mechanisms act in isolation. Here, we show that the combined impact of the two contagion channels also cannot the market price of listed equity shares issued by South African banks below zero.

Remember that we have assumed that the overlapping portfolio contagion channel and shareholder contagion channels are additive. Therefore, the market price of a tradable security σ evolves according to

$$p_{\sigma,t+1} = p_{\sigma,t} - \Delta p^s_{\sigma,t} - \Delta p^o_{\sigma,t}, \tag{C.7}$$

where $\Delta p^s_{\sigma,t}$ denotes the drop in market price due to the shareholder contagion channel, and $\Delta p^o_{\sigma,t}$ denotes the drop in market price due to the overlapping portfolio contagion channel. To demonstrate that neither channel can drive the market price below zero when both channels interaction, we first show that $\Delta p^o_{\sigma,t} \leq p_{\sigma,t} - \Delta p^s_{\sigma,t}$, i.e. overlapping portfolio contagion does not drive the market price below zero even when it has already been depressed by shareholder contagion. Second, we discuss why $\Delta p^s_{\sigma,t} \leq p_{\sigma,t}$, i.e. shareholder contagion does not drive the market price below zero even when the market price has been depressed by overlapping portfolio contagion at previous times.

Because we have assumed that all tradable securities σ sold at time t are sold against the new price $p_{\sigma,t+1}$, the liquidity recovered from the sale is reduced by the shareholder contagion $\Delta p^s_{\sigma,t}$. Hence, equation (C.5) becomes

$$\Delta p^o_{\sigma,t} = \frac{p_{\sigma,t+1}\Delta n_{\sigma i,t}}{D_\sigma} = \frac{\left(p_{\sigma,t} - \Delta p^s_{\sigma,t} - \Delta p^o_{\sigma,t}\right)\Delta n_{\sigma i,t}}{D_\sigma}, \tag{C.8}$$

and rewriting yields

$$\Delta p^o_{\sigma,t} = \frac{\left(p_{\sigma,t} - \Delta p^s_{\sigma,t}\right)\Delta n_{\sigma i,t}}{D_\sigma + \Delta n_{\sigma i,t}}, \tag{C.9}$$

so we find that $\Delta p^o_{\sigma,t} \leq p_{\sigma,t} - \Delta p^s_{\sigma,t}$. Hence, the shareholder contagion channel depresses the overlapping portfolio contagion channel by reducing the price impact per share sold, similar to Wiersema et al. [2021].

To guarantee that the shareholder contagion channel cannot drive the market price of equity shares below zero when the price has already been depressed by the overlapping contagion channel at a previous time, we should multiply equation (5.6) 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 (5.6) with the market-to-book ratio yields

$$A^{vv}_{\sigma i,t} = s_{i\sigma,t}\frac{p_{\sigma,t}}{E_{\sigma,t}}, \tag{C.10}$$

such that when institution σ suffers a loss $x^v_{\sigma,t} = E_{\sigma,t}$, the shareholder contagion suffered by i equals $s_{i\sigma,t}p_{\sigma,t}$. Hence, when a shock exhausts institution σ's equity, institution i loses the current market value of its position in shares issued by σ. The market-to-book ratio reflects that the overlapping portfolio contagion channel has depressed the equity shares' market value below their book value, such that the impact of the shareholder contagion channel is reduced (Wiersema et al., 2021).

However, 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, so the market-to-book ratio $p_{\sigma,t}S_{\sigma,t}/E_{\sigma,t} = 1$ and can be omitted from equation (5.6) for simplicity.

C.4 Baseline Market Depth Estimates

We estimate baseline market depths $\hat{D}_\sigma$ for six different classes 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. Due to data limitations, we do not distinguish between tradable securities of a specific type issued by different non-financial corporates, nor between tradable securities of a specific type issued by different domestic banks. 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.

We set our baseline estimate of a security's market depth equal to its market capitalization divided by its initial price (similar to e.g. Wiersema et al., 2019, 2021). As initial prices are normalized, division by the initial price simply serves to make the baseline market depth estimate dimensionless. For (domestic) government bonds, we use the market capitalization of South African government bonds at the end of 2016[1]. Note that the banks and funds own 21% of the market capitalization of the South African government bonds. Due to data limitations, we estimate 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 class is given by banks' and funds' aggregate holdings of the security class divided by 21%.

[1]The market capitalization of South African bonds as of Q4 2016 is sourced from the Q1 2017 SARB Quarterly Bulletin; https://www.resbank.co.za/content/dam/sarb/publications/quarterly-bulletins/quarterly-bulletin-publications/2017/7718/07Statistical-tables—Public-Finance.pdf.

Bibliography

D. Acemoglu, V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 2012.

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

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

V. V. Acharya and T. Yorulmazer. Information contagion and bank herding. *Journal of money, credit and Banking*, 40(1):215–231, 2008.

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.

T. Adrian, H. S. Shin, et al. Liquidity and financial contagion. *Banque de France Financial Stability Review: Special Issue on Liquidity*, 11:1–7, 2008.

J. Aharony and I. Swary. Additional evidence on the information-based contagion effects of bank failures. *Journal of Banking & Finance*, 20(1):57–69, 1996.

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

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

I. Aldasoro and I. Alves. Multiplex interbank networks and systemic importance: An application to european data. *Journal of Financial Stability*, 35:17–37, 2018.

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

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

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

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

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

K. Anand, G. Bédard-Pagé, and V. Traclet. Stress testing the canadian banking system: A system-wide approach. *Financial System Review*, 61, 2014.

K. Anand, C. Gauthier, and M. Souissi. Quantifying contagion risk in funding markets: A model-based stress-testing approach. Technical report, Bank of Canada Working Paper, 2015.

R. Anderson, J. Dannielson, C. Baba, U. S. Das, H. Kang, and M. Segoviano. Macroprudential stress tests and policies: Searching for robust and implementable frameworks. Technical report, IMF Working Paper No. 18/197, 2018.

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

N. Arinaminpathy, S. Kapadia, and R. M. May. Size and complexity in model financial systems. *Proceedings of the National Academy of Sciences*, 109(45): 18338–18343, 2012.

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

C. Aymanns and J. D. Farmer. The dynamics of the leverage cycle. *Journal of Economic Dynamics and Control*, 50:155–179, 2015.

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

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

W. Bagehot. *Lombard Street: A description of the money market.* London: HS King, 1873.

M. Bardoscia, S. Battiston, F. Caccioli, and G. Caldarelli. Debtrank: A microscopic foundation for shock propagation. *PloS one*, 10(6):e0130406, 2015.

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

M. Bardoscia, G. Bianconi, and G. Ferrara. Multiplex network analysis of the UK OTC derivatives market. Technical report, Bank of England Working Paper No. 726, 2018.

S. Battiston and S. Martinez-Jaramillo. Financial networks and stress testing: Challenges and new research avenues for systemic risk analysis and financial stability implications, 2018.

S. Battiston, M. Puliga, R. Kaushik, P. Tasca, and G. Caldarelli. Debtrank: Too central to fail? financial networks, the fed and systemic risk. *Scientific reports*, 2: 541, 2012.

S. Battiston, G. Caldarelli, and M. D'Errico. The financial system as a nexus of interconnected networks. In *Interconnected networks*, pages 195–229. Springer, 2016a.

S. Battiston, M. D'Errico, and S. Gurciullo. Debtrank and the network of leverage. *The Journal of Private Equity*, 20(1):58–71, 2016b.

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

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

R. J. Berndsen, C. León, and L. Renneboog. Financial stability in networks of financial institutions and market infrastructures. *Journal of Financial Stability*, 35: 120–135, 2018.

BIS. Strengthening the resilience of the banking sector. Technical report, Bank of International Settlements, 2009.

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.

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

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

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

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

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

BoE. General insurance stress test 2019 - scenario specification, guidelines and instructions. Technical report, Bank of England, 2019.

R. Bookstaber. Agent-based models for financial crises. *Annual Review of Financial Economics*, 9:85–100, 2017.

R. Bookstaber, J. Cetina, G. Feldberg, M. Flood, and P. Glasserman. Stress tests to promote financial stability: Assessing progress and looking to the future. *Journal of Risk Management in Financial Institutions*, 7(1):16–25, 2014.

R. Bookstaber, D. Y. Kenett, et al. Looking deeper, seeing more: a multilayer map of the financial system. *OFR Brief*, 16(06), 2016.

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.

M. Boss, H. Elsinger, M. Summer, and S. Thurner 4. Network topology of the interbank market. *Quantitative finance*, 4(6):677–684, 2004.

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.

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

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

K. B. Budnik, M. Balatti, G. Covi, I. Dimitrov, J. Groß, I. Hansen, M. Kleemann, T. Reichenbachas, F. Sanna, A. Sarychev, et al. Macroprudential stress test of the euro area banking system. Technical report, ECB Occasional Paper Series, 2019.

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

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

F. Caccioli, T. A. Catanach, and J. D. Farmer. Heterogeneity, correlations and financial contagion. *Advances in Complex Systems*, 15(supp02):1250058, 2012.

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. 2013.

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

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

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

B. Candelon and M. A. N. Sy. *How did markets react to stress tests?* International Monetary Fund, 2015.

CFTC. Supervisory stress test of clearinghouses. Technical report, US Commodity Futures Trading Commission, 2016.

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

R. Cont and E. Schaanning. Fire sales, indirect contagion and systemic stress testing. Technical report, Norges Bank Working Paper No. 2017/2, 2017.

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

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

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

R. Cont, A. Moussa, et al. Network structure and systemic risk in banking systems. In J.-P. Fouque and J. A. Langsam, editors, *Handbook on Systemic Risk*, chapter 13, pages 327–36. Cambridge University Press, Cambridge, 2013.

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

R. Cont, A. Kotlicki, and L. Valderrama. Liquidity at risk: Joint stress testing of solvency and liquidity. *Journal of Banking & Finance*, 118:105871, 2020.

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

B. Craig and G. Von Peter. Interbank tiering and money center banks. *Journal of Financial Intermediation*, 23(3):322–347, 2014.

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

J. Danielsson and H. S. Shin. Endogenous risk. *Modern risk management: A history*, pages 297–316, 2003.

J. de Haan. How emergence arises. *Ecological complexity*, 3(4):293–301, 2006.

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.

S. Dees and J. Henry. Stress-test analytics for macroprudential purposes: Introducing stamp€. *Satellite Models*, 13, 2017.

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. Di Gangi, F. Lillo, and D. Pirino. Assessing systemic risk due to fire sales spillover through maximum entropy network reconstruction. *Journal of Economic dynamics and control*, 94:117–141, 2018.

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. *FRB of New York Staff Report*, (645), 2018.

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

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

EBA. Adverse macro-financial scenario for the 2018 eu-wide banking sector stress test. Technical report, European Banking Authority, 2018.

ECB. Eurosystem Supervisory Data. https://www.bankingsupervision.europa.eu/banking/statistics, 2019. [Online; accessed 10-April-2019].

EIOPA. Occupational pensions stress test 2019. Technical report, European Insurance and Occupational Pensions Authority, 2019. URL https://www.eiopa.europa.eu/occupational-pensions-stress-test-2019_en.

EIOPA. Insurance stress test 2021 technical specifications. Technical report, European Insurance and Occupational Pensions Authority, 2021.

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

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

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

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

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

ESMA. Eu-wide ccp stress test 2015. Technical report, European Securities and Markets Authority, 2015.

J. D. Farmer. Market force, ecology and evolution. *Industrial and Corporate Change*, 11(5):895–953, 2002.

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.

J. D. Farmer, A. M. Kleinnijenhuis, and T. Wetzer. Stress testing the financial macrocosm. *Forthcoming in Handbook of Financial Stress Testing*, 2021a.

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

FED. Supervisory scenarios for annual stress tests required under the dodd-frank act stress testing rules and the capital plan rule 2018. Technical report, Board of Governors of the Federal Reserve System, 2018.

A. Foglia. Stress testing credit risk: a survey of authorities' approaches. *Eighteenth issue (September 2009) of the International Journal of Central Banking*, 2018.

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

J.-P. Fouque and J. A. Langsam. *Handbook on systemic risk*. Cambridge University Press, 2013.

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

D. Fricke and T. Lux. Core–periphery structure in the overnight money market: evidence from the e-mid trading platform. *Computational Economics*, 45(3):359–395, 2015.

FSB. Basel iii: Finalising post-crisis reforms. Technical report, Financial Stability Board, 2017.

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

C. H. Furfine. Interbank exposures: Quantifying the risk of contagion. *Journal of money, credit and banking*, pages 111–128, 2003.

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

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

A. Gandy and L. A. M. Veraart. Adjustable network reconstruction with applications to cds exposures. *Journal of Multivariate Analysis*, 172:193–209, 2019.

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.

P. Glasserman and G. Tangirala. Are the fed's stress test results predictable? *Columbia Business School Research Paper*, (15-23), 2015.

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

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

C. Goodhart. Ratio controls need reconsideration. *Journal of Financial Stability*, 9 (3):445–450, 2013.

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

D. Greenlaw, A. K. Kashyap, K. L. Schoenholtz, and H. S. Shin. Stressed out: Macroprudential principles for stress testing. *Chicago Booth Research Paper*, (12-08), 2012.

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

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

G. Hałaj and C. Kok. Assessing interbank contagion using simulated networks. *Computational Management Science*, 10(2-3):157–186, 2013.

A. G. Haldane. Rethinking the financial network. In *Fragile stabilität–stabile fragilität*, pages 243–278. Springer, 2013.

F. Heider, M. Hoerova, and C. Holthausen. Liquidity hoarding and interbank market spreads: the role of counterparty risk. Technical report, European Central Bank, 2009.

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

J. H. Holland. *Emergence: From chaos to order.* OUP Oxford, 2000.

A.-C. Hüser and C. Kok. Mapping bank securities across euro area sectors: comparing funding and exposure networks. 2019.

A.-C. Hüser, G. Hałaj, C. Kok, C. Perales, and A. van der Kraaij. The systemic implications of bail-in: a multi-layered network approach. *Journal of Financial Stability*, 38:81–97, 2018.

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

IMF. Canada financial system stability assessment. Technical report, International Monetary Fund, 2014.

IMF. Germany financial system stability assessment. Technical report, International Monetary Fund, 2016.

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

M. A. A. Jobst, N. Sugimoto, and T. Broszeit. Macroprudential solvency stress testing of the insurance sector. 2014.

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

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

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.

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

C. Kok and M. Montagna. Multi-layered interbank model for assessing systemic risk. Technical report, ECB Working Paper No. 1944, 2016.

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

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

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

T. Lux. A model of the topology of the bank–firm credit network and its role as channel of contagion. *Journal of Economic Dynamics and Control*, 66:36–53, 2016.

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.

R. Mastrandrea, T. Squartini, G. Fagiolo, and D. Garlaschelli. Enhanced reconstruction of weighted networks from strengths and degrees. *New Journal of Physics*, 16(4):043022, 2014.

P. Mazzarisi and F. Lillo. Methods for reconstructing interbank networks from limited information: A comparison. In *Econophysics and Sociophysics: Recent Progress and Future Directions*, pages 201–215. Springer, 2017.

M. Montagna, G. Torri, and G. Covi. On the origin of systemic risk. *Available at SSRN 3699369*, 2020.

N. Musmeci, S. Battiston, G. Caldarelli, M. Puliga, and A. Gabrielli. Bootstrapping topological properties and systemic risk of complex networks using the fitness model. *Journal of Statistical Physics*, 151(3):720–734, 2013.

M. Newman. *Networks*. Oxford university press, 2018.

A. Persaud. Macro-prudential regulation. 2009.

S. Poledna, S. Thurner, J. D. Farmer, and J. Geanakoplos. Leverage-induced systemic risk under basle ii and other credit risk policies. *Journal of Banking & Finance*, 42:199–212, 2014.

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.

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

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

M. Quagliariello. *Stress-testing the banking system: methodologies and applications*. Cambridge University Press, 2009.

J.-C. Rochet and J. Tirole. Interbank lending and systemic risk. *Journal of Money, credit and Banking*, 28(4):733–762, 1996.

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.

L. C. Rogers and L. A. Veraart. Failure and rescue in an interbank network. *Management Science*, 59(4):882–898, 2013.

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

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

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

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

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

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

T. Squartini, G. Cimini, A. Gabrielli, and D. Garlaschelli. Network reconstruction via density sampling. *Applied network science*, 2(1):3, 2017.

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

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

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, E. Kemp, and T. Wetzer. Higher-order exposures. *Available at SSRN 3914076*, 2021.