

System-Wide Stress Testing & Systemic Risk



Alissa M. Kleinnijenhuis
Christ Church
University of Oxford

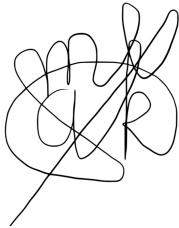
A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2019

Statement of Originality

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

Candidate: Alissa Magdaleen Kleinnijenhuis

Signature:

A handwritten signature in black ink, consisting of a series of loops and a long horizontal stroke at the bottom.

Date: 9 September 2019

This thesis is dedicated to

Rafael Baptista Ochoa

Acknowledgements

I would like to wholeheartedly thank Professor J. Doyne Farmer for providing a fertile ground to execute my research agenda and for having taught me many aspects of conducting scientific research, as well as for being ever inspiring. I am profoundly grateful to my parents, Professor Jan Kleinnijenhuis and Dr. Anne Magda Smilde (*in spe*), for raising me in an intellectual environment and for their invaluable support throughout my studies. I cannot have been more thankful to have met the light of my life Dr. Mustafa G. Aydogan in Oxford. My fond memories go out to Rafael Baptista Ochoa, a close friend and collaborator. My research would not have been possible without the generous support of the Institute of New Economic Thinking, the Mathematical Institute (Thornton scholarship) and Christ Church.

Abstract

The financial crisis of 2007-2009, which brought the entire system at the brink of collapse, renewed efforts to guard against financial instability. A key pillar of the post-crisis regulatory toolkit is ‘stress testing’. Stress tests provide a forward-looking examination of firms’ potential losses during severely adverse conditions. And enable timely action to recapitalise those firms who experience capital shortfalls in such crisis scenarios. Today’s regulatory stress tests do not heed the *key* lesson of the financial crisis: amplifications in the networked financial system *must* be taken into account to be able to assess systemic risk. Because of this, these tests are unable to assess systemic risk and *ergo* to address it – defeating their *raison d’être*.

The overarching research question in this thesis is whether new building blocks – expressing the heterogeneity of institutions, contracts, markets, constraints and behaviour in the interconnected financial system – can be supplied for system-wide stress tests to better capture the endogenous amplification of shocks in order to improve the assessment of systemic risk and the evaluation of prudential policies to address financial fragility.

The cornerstone of my thesis is the development of a *generic* network-based method, comprised of these five building blocks (i.e. institutions, contracts, markets, constraints and behaviour), for system-wide stress testing – which has gained traction from leading central banks, including the Bank of England and the European Central Bank. Using this method, I implement two data-driven models to address some of the most salient financial stability questions of today. First, we ask how the regulatory buffer size and its usability under Basel III affect systemic risk? We find that financial resilience decreases if regulatory buffers are seen to be less usable by banks. If regulatory buffers are not treated as usable, then regulatory buffers *de facto* act as capital requirements. In such case, if an adverse shock threatens an institution to breach its capital buffers constraints, it is forced to delever, which tends to

have a destabilising effect on the financial markets. We show that the size of usable regulatory buffers that is required to maintain stability is underestimated if the interaction between exposure loss contagion, funding contagion, overlapping portfolio contagion and margin call contagion is not taken into account. Second, we inquire what the systemic implications are of the bail-in design to resolve systemically important banks? First of all, we find that the bail-in design tremendously matters for whether bail-ins can be credibly executed in system-wide financial crises and cases of large systemically important bank failures, without significantly exacerbating financial distress. Our results demonstrate that an early bail-in, strong recapitalisation and fair distribution of equity compensation by means of debt-to-equity conversion rates makes bail-in a feasible option on the table for idiosyncratic cases of bank failure and limits – but not eliminates – contagion in cases of system-wide distress. We further show that excluding run-prone, short-term debt from the application of the bail-in tool, increasing the requirements on loss absorbing debt and providing investors with certainty about the bail-in design lowers contagion in system-wide crises to manageable levels. Our findings highlight that while well-designed bail-ins could be credibly administered in system-wide crises, it is not clear that the current bail-in design is in the regime of stability. Altogether, the methods and findings of this thesis emphasise the promise that system-wide stress tests hold for regulators to efficaciously assess systemic risk and calibrate prudential policies constituting the financial architecture.

Contents

1	Introduction	1
2	Models of Financial Stability and Their Application in Stress Tests	11
2.1	Summary	11
2.2	Introduction	11
2.3	Two Approaches to Modeling Systemic Risk	15
2.4	A View of the Financial System	16
2.4.1	Balance sheet composition	17
2.4.2	Balance sheet dynamics	18
2.5	Contagion in Financial Networks	20
2.5.1	Financial linkages and channels of contagion	20
2.5.2	Counterparty loss contagion	23
2.5.3	Overlapping portfolio contagion	26
2.5.4	Funding liquidity contagion	29
2.5.5	Interaction of contagion channels	30
2.6	From Models to Policy: Stress Tests	31
2.6.1	What are stress tests?	31
2.6.2	A brief history of stress test	31
2.7	Microprudential Stress Tests	33
2.7.1	Microprudential stress tests of banks	33
2.7.2	Strengths and weaknesses of current microprudential stress tests	36
2.8	Macroprudential Stress Tests	38
2.8.1	Examples of macroprudential stress tests	39
2.8.2	Comparing and evaluating macroprudential stress tests: five building blocks	45
2.8.3	Strengths and weaknesses of the current macroprudential stress tests	51
2.9	Conclusion	53

3	Foundations of System-Wide Stress Testing	54
3.1	Summary	54
3.2	Introduction	55
3.3	Relevant Literature	58
3.3.1	Modelling System-Wide Stress Dynamics	58
3.3.2	Stress Tests and Prudential Regulation	60
3.4	A Framework for System-Wide Stress Tests with Heterogeneous Institutions	61
3.4.1	Five Building Blocks to Represent Financial Systems	62
3.4.2	Contagion and Amplification	67
3.4.3	Key Design Choices for Simulation Software	67
3.5	A System-Wide Stress Test Model for the European Financial System . .	69
3.5.1	Financial Institutions and their Constraints	70
3.5.2	Financial Contracts, Markets & Contagion Mechanisms	75
3.5.3	Behaviour	81
3.6	Policy experiments and Results	90
3.6.1	Default Parameters and Visualisation	91
3.6.2	From Micro to Macro: A Macroprudential Overlay for the EBA 2018 Stress Test	93
3.6.3	Contagious Feedback Loops Between Banks and Non-Banks	97
3.6.4	Amplification of Contagion Mechanisms	99
3.6.5	‘Usability’ of Buffers and Contagion	101
3.6.6	Calibration of Buffers with System-Wide Stress Tests	102
3.7	Policy Implications and Conclusion	104
4	Systemic Implications of the Bail-In Design	107
4.1	Summary	107
4.2	Introduction	108
4.3	Introduction to Bail-In	112
4.3.1	Failure Threshold: Conditions to Resolve a Bank	113
4.3.2	The Prime Resolution Tool is the Bail-In Tool	114
4.3.3	The Bail-In Tool	114
4.3.4	Following a Bail-In: Restructuring	119
4.4	Data	120
4.5	Model	122
4.5.1	Modelling the Design of the Bail-In Mechanism	122
4.5.2	Modelling Multiple Contagion Mechanisms in the Financial System	153
4.6	Link to the Literature	168

4.7	Results	169
4.7.1	Stability Impact of the Primary Bail-In Design	175
4.7.2	Stability Impact of the Secondary Bail-In Design	183
4.7.3	Contagious Amplifications of the Bail-in Design	186
4.8	Link to the Literature	192
4.9	Discussion and Policy Implications	193
5	Concluding Remarks	199
A	Foundation of System-Wide Stress Testing	202
A.1	Further Details on Model Implementation	202
A.1.1	Notation	202
A.1.2	Initialisation	203
A.1.3	Default Configuration	207
A.2	System-Wide Stress Testing Software	208
A.2.1	Design Principles	208
A.3	Sensitivity Analysis	214
B	Systemic Implications of the Bail-In Design	221
B.1	Data	221
B.1.1	Initialisation of Balance Sheets	221
B.1.2	Initialisation of Seniority Classes	221
B.2	Notation	224
B.3	Model Specifics	225
B.3.1	Alternative Ways of Dealing with Bank Failure	225
B.4	Sensitivity Analysis	226
B.5	Detailed Descriptions	231
B.5.1	Comparing Debt Eligibility Criteria for MREL, TLAC, and Bail-In	231
B.5.2	Principles on the Debt-to-Equity Conversion Rate	234
	Bibliography	237

List of Figures

2.1	Diversification & Stability	25
2.2	Robust yet Fragile	28
2.3	Map of Financial System	33
2.4	Time Steps in RAMSI Stress Test	40
2.5	Modules in MFRAF Stress Test	40
2.6	Time Steps in MFRAF Stress Test	41
2.7	Composition of ABM for Financial Vulnerabilities Stress Test	42
3.1	Stylised balance sheet of bank	70
3.2	Balance sheets of non-banks.	75
3.3	Systemic risk as a function of the scaled 2018 EBA scenerio	94
3.4	Systemic risk as a function of the scaled 2018 EBA scenario for triple buffers	97
3.5	Impact of non-banks on systemic risk	99
3.6	Amplification among contagion mechanisms	101
3.7	The usability of buffers and systemic risk	102
3.8	Buffer sizes and systemic risk	103
4.1	Stylised balance sheet of a bank	120
4.2	Stylised balance sheet of a bank	121
4.3	Bail-inable debt composition	124
4.4	Debt Composition of each Bank	125
4.5	Aggregate Debt Domposition of EU Banks	126
4.6	Excess bail-inable debt of banks above their loss absorbing requirements	131
4.7	Fair conversion rates in bail-in	142
4.8	Scope for unfair conversion rates and necessity of the resolution financing fund	147
4.9	Bail-in process for unfair and fair conversion rates	150
4.10	Jump diffusion process to simulate the asset value and risk-weighted capital ratio	160
4.11	Sensitivity of bail-inable debt value to jump process parameters	162

4.12	Sensitivity of bail-inable debt value to the bail-in parameters	163
4.13	Sensitivity of bail-inable debt value to the priority class	164
4.14	Stability Impact of Primary Bail-In Design	175
4.15	Stability Impact of FLTF Ratio	177
4.16	Stability Impact of the Recapitalisation Target	179
4.17	Stability Impact of the Conversion Rates	181
4.18	Stability Impact of Secondary Bail-In Design	183
4.19	Stability Impact of Loss Absorbing Requirements and Debt Exclusions	184
4.20	Contagious Amplifications of the Primary Bail-In Design	186
4.21	Contagious Amplifications of the Secondary Bail-In Design	190
A.1	Class diagram of system-wide stress testing code	209
A.2	Systemic risk as a function of the scaled 2018 EBA shock turning the leverage ratio off	214
A.3	Systemic risk as a function of the scaled 2018 EBA shock turning the risk-weighted capital ratio off	214
A.4	Excess systemic risk due to amplification under Basel III	216
A.5	Amplification among contagion mechanisms under Basel III	217
A.6	Usability of buffers and systemic risk under disorderly liquidation	218
A.7	Usability of the LCR and systemic risk	219
A.8	The affect of the aggressiveness of prudent actions and systemic risk	219
A.9	Systemic risk and the size of buffers as a function of the price impact	220
B.1	Stability Impact of Bail-In Design: Small SIBs	226
B.2	Comparative Stability Impact of Failure Proceedings: Good Primary Bail-In Parameters	227
B.3	Comparative Stability Impact of Failure Proceedings: Bad Primary Bail-In Parameters	227
B.4	Stability Impact of Failure Threshold: Bad Primary Settings	228
B.5	Stability Impact of the Recapitalisation Target: Bad Primary Settings	228
B.6	Contagious Amplifications of Primary Bail-In Design: Bad Settings	229
B.7	Stability Impact of Non-Bank Leverage	229
B.8	Stability Impact of the Size of Leveraged Non-Banks	230
B.9	Asset Loss to Real Economy	230

Chapter 1

Introduction

Ben S. Bernanke served as Chairman of the Federal Reserve (Fed) from 2006 to 2014. He presided over steering the economy of the United States and arguably the wider world economy through the Great Recession that took place between 2007-2009. He wrote the following passage in his memoirs titled ‘The Courage to Act’: ‘Seen from the vantage point of early 2007, the economy’s good performance, combined with the relatively small size of the subprime mortgage market and what appeared to be a healthy banking system, led me and others at the Fed to conclude that subprime problems – though certainly a major concern for affected communities and the housing sector generally – were unlikely to cause major economic damage. But we failed to anticipate that problems in the subprime mortgage market could trigger an old-fashioned financial panic, albeit in a new, unfamiliar guise (Bernanke (2015)).’ We calculated that ‘even if subprime mortgages defaulted at extraordinarily high rates, [...], the resulting financial losses would be smaller than those from a single bad day in global stockmarkets (Bernanke (2015)).’ We came to realise that the “ultimate economic costs of the panic far outweighed the magnitude of the trigger” (Bernanke (2015)), like it had done in the 1907 Panic. Bernanke addressed the need to account for endogenous amplifications to comprehend systemic risk in the financial system.

The overarching research question in this thesis is whether new building blocks – expressing the heterogeneity of institutions, contracts, markets, constraints and behaviour in the interconnected financial system – can be supplied for system-wide stress tests to better capture the endogenous amplification of shocks in order to improve the assessment of systemic risk and the evaluation of prudential policies to address financial fragility.

A few additional statements about the nature of the 2007-2009 crisis and recently-introduced stress tests will be made, before elaborating on the novel methods that will

be developed and applied in the chapters of this thesis.

The new financial crisis (2007-2009) revived the old understanding that exogenous shocks can be endogenously amplified by the financial system to give widespread financial turmoil. Brunnermeier & Pedersen (2009) are the first to document the prevailing contagion mechanisms through which this amplification took place in the last crisis. First, borrowing institutions – such as banks – are exposed to two ‘liquidity spirals’. Both asset losses and funding shocks cause distressed sales of assets, so-called ‘fire sales’, resulting in a decline in asset prices. This erodes an institution’s capital buffers, prompting a further round of distressed sales and further tightening of lending. Second, institutions may start ‘hoarding liquidity’ when they become anxious about their future access to capital markets. Such funding withdrawals may take place irrespective of the creditworthiness of borrowers. Though such behaviour might be rational from an institution’s perspective, it is destabilising from the perspective of the system as a whole: it imposes funding shocks upon counterparties, who may in turn be prompted to liquidate assets at fire-sale prices to meet their repayment obligation. Third, runs on financial institutions – typically on those who are considered to be the ‘slowest antelopes in the herd’ (Blinder (2013)) so that these lose access to market funding – can further cause capital to collapse. Finally, defaulted institutions can afflict exposure losses upon counterparties, which can – if these losses causes counterparties to default in turn – trigger a domino chain of further exposure losses in the networked financial system.

The experience of the crisis and its lessons into the amplification mechanisms that made it so calamitous, ‘exposed a significant lack of information as well as data gaps on key financial sector vulnerabilities relevant for financial stability analysis’ (FSB (2009b)). The Financial Stability Board (FSB), an international body that monitors the resilience of the global financial system and makes recommendations to improve it, recommended in 2009 that ‘quantitative tools and indicators to (better) monitor and assess the build-up of macroprudential risks in the financial system’ be developed (FSB (2009a)).

The prime tool that emerged as part of the post-crisis regulatory toolkit to assess and address systemic risk is a ‘stress test’. Though the International Monetary Fund (IMF) had conducted early forms of stress tests since 1999, under the so-called Financial Sector Assessment Program (FSAP), to analyse the resilience of a country’s financial sector (IMF (2019)). It is only since 2009 that stress tests have become institutionalised by most major jurisdictions (as an annual or bi-annual exercise). The first such stress test was the Supervisory Capital Assessment Program (SCAP) conducted by the Federal Reserve

System in 2009. It was designed to assess whether the largest financial institutions in the United States had sufficient capital buffers to withstand the financial market turbulence and impending recession ensuing from the financial crisis. Just as later stress tests would do, it assessed the banking system's ability to survive through a 'what-if' exogenous scenario of 'severely adverse' nature and compared this situation against a 'baseline' scenario. [Greenwood et al. \(2017\)](#) note that 'the severely adverse scenario contemplated in the SCAP was actually a fair representation of the reality at the time, in the depths of the financial crisis.' It essentially marked-to-market the assets of banks inducing banks to reflect expected losses in reported equity capital. This ensured that the banks received regulatory pressure to recognise the full reality of their solvency problems and take steps to enlarge their capital buffers. Ben S. Bernanke describes the introduction of the SCAP stress test as a turning point in the crisis:

'In retrospect, the SCAP stands out for me as one of the critical turning points in the financial crisis. It provided anxious investors with something they craved: credible information about prospective losses at banks. Supervisors' public disclosure of the stress test results helped restore confidence in the banking system and enabled its successful recapitalization' ([Bernanke \(2013\)](#))

Since SCAP, the United States, the European Union (EU) and the United Kingdom – among others – have set in a place an (bi-) annual stress testing exercise in which systemically-important financial institutions are required to participate. These are called the Dodd-Frank Act stress test (taking place as part of the Comprehensive Capital Analysis and Review (CCAR)), the European Banking Authority (EBA) stress test and the Bank of England (BoE) stress test, respectively.

Though the newer stress tests have brought many benefits, such as impelling banks to holistically assess their risk rather than doing so department-by-department in a disaggregated way, some have argued that their value does not reach their potential. Instead, their performance appears far less than their powerful promise. For instance, [Glasserman et al. \(2015\)](#) show that stress test outcomes have become more predictable over time. This is problematic because it means the informational value stress tests provide is largely lost, which undercuts their purpose as an early detector of potential capital shortfalls. One potential reason stress tests have become more predictable is that stress tests are easy to game ([Sarin & Summers \(2020, forthcoming in Handbook of Financial Stress Testing\)](#)). [Bulow & Klemperer \(2013\)](#) point out that both the numerator and denominator of bank equity ratios are subject to manipulation. Banks not only have the ability to game the

stress tests, it is also in their interest to do so. Banks that suffer too large losses under the adverse scenario must raise capital or retain dividends. Those who suffer too small losses, on the other hand, are seen as too conservative by regulators. So, it is best to manipulate stress outcomes in such a way that the capital ratio under an adverse scenario is in the middle of the herd, avoiding regulatory scrutiny either way.

This is, however, not the biggest shortcoming of stress tests, undercutting their potential as a powerful monitoring and intervention tool to preserve financial stability. Today’s regulatory stress tests do not heed the *key* lesson of the financial crisis, as clearly narrated by [Bernanke \(2015\)](#) and [Brunnermeier & Pedersen \(2009\)](#): amplifications in the networked financial system *must* be taken into account to be able to assess systemic risk – where systemic risk refers to the risk of a breakdown of the entire financial system in response to (a potentially small) shock.¹ This means that these so-called ‘microprudential’ stress tests are unable to assess systemic risk and *ergo* to address it – defeating their *raison d’être*. Instead, microprudential stress tests evaluate the resilience of individual institution to specific shocks, assuming these institutions live on independent islands barring the spread of any contagion via their interconnections. Some regulators argue that microprudential stress tests ‘indirectly’ capture contagious amplifications by ensuring that the adverse scenario is severe enough so as to also capture not just initial shocks, such as the subprime collapse, but also higher-order shocks ones. We, in [Chapter 3](#), invalidate this argument. We show that two identical adverse shock scenarios with two identical first-order loss impacts applied to two different financial systems can result in widely different system-wide losses for each system, depending on the shock-amplifying tendency of that financial system. Therefore, the outcome of the microprudential stress test, which only captures first-order losses, is *not* informative on systemic risk. This result shows that ‘macroprudential’ stress tests, also called system-wide stress tests – which assess the resilience of the networked financial system as a whole – are not superfluous. They are essential.

Outside of the regulatory perimeter, researchers have been exploring network methods to assess systemic risk, which *do* incorporate contagious amplifications. The ‘financial networks and systemic risk’ field has started to flourish in the aftermath of the 2007-2009 financial crisis. At first, these models typically captured a single contagion mechanism, such as exposure loss contagion (see e.g. [Upper & Worms \(2004\)](#), [E Santos et al. \(2010\)](#))

¹A widely adopted, more elaborate, definition of systemic risk is provided by the FSB, IMF and Bank of International Settlements. They define it as ‘a risk of disruption to financial services that is caused by an impairment of all or parts of the financial system and has the potential to have serious negative consequences for the real economy’ ([FSB \(2009c\)](#)).

or overlapping portfolio contagion (also called fire sale contagion; see e.g. [Caccioli et al. \(2014, 2015\)](#), [Duarte & Eisenbach \(2015\)](#), [Greenwood et al. \(2015\)](#), [Cont & Schaanning \(2017\)](#)), and they primarily focused on post-default domino effects. Increasingly, these models encapsulate how multiple contagion mechanisms, such as those described by Markus Brunnermeier ([Brunnermeier & Pedersen \(2009\)](#)), interact and amplify one another – not only post-failure, but also in prior, institutions may act in individually rational, though systemically destabilising ways to avoid default ([Caccioli et al. \(2013\)](#), [Kok & Montagna \(2013\)](#), [Poledna et al. \(2015\)](#), [Hüser & Kok \(2019\)](#), [Wiersema et al. \(2019\)](#)).

Early-day macroprudential stress tests developed by central banks have sought to incorporate these financial stability models to explore their usefulness, as we discuss in [Chapter 2](#). Although some of these macroprudential stress tests inform policy, none has yet been employed by regulators to: (1) systemically assess systemic risk; (2) identify vulnerabilities; and (3) calibrate policies, whereas their microprudential counterparts have been. Indubitably, there are politically-motivated and data-related reasons that have thus far impeded the entrance of macroprudential stress tests onto the grand stage of macroprudential policymaking.

A scientific reason exists too. According to [Anderson et al. \(2018\)](#), regulators are ‘searching for robust and implementable frameworks’ to conduct macroprudential stress tests, which would both assess systemic risk and evaluate policies. Financial stability models so far have been largely inapplicable to serve as models underpinning a prospective regulatory system-wide stress test – which could be conducted on an annual basis. Indicated reasons ([Anderson et al. \(2018\)](#)) include, for instance, that they are too narrow in scope: focusing only on one contagion mechanism rather than interacting contagion mechanisms, on one type of institution (a bank) rather than the different type of financial institutions that operate in the financial network, and/or on one regulatory constraint rather than the swathe of regulatory policies that apply. Another reason is that they are not generic enough. Journal-focused models are often tailored to address one research question, thereby losing their general applicability to assess systemic risk in any financial system.

The lack of serious regulatory use of system-wide stress tests is a lost opportunity. System-wide stress tests models are singularly capable of evaluating the system-wide effects of microprudential policies (i.e. policies focusing on the robustness of individual institutions) and calibrating macroprudential policies (i.e. policies aiming for system-wide re-

silience). Microprudential stress tests *cannot* serve as a substitute for macroprudential stress tests to perform these tasks. [Aymanns et al. \(2016\)](#) have shown that well-intended microprudential policies, such as the Basel II capital regulation, can lead to unstable macroprudential outcomes. The Basel II capital regulation puts a cap on the amount of leverage a bank can attain, dependent on the amount of risk – measured by the value-at-risk (VaR) – it faces. From an individual microprudential perspective such regulation makes perfect sense. However, from a system-wide perspective this policy is destabilising. Under the Basel II rules, a bank’s leverage constraint tightens when its VaR increases. Because a decrease in asset prices causes both the leverage and the VaR to rise, declines in asset prices may force disorderly asset sales to delever to comply with the Basel II leverage limit. This triggers a positive feedback loop where selling prompts more selling, which further tightens the Basel II leverage constraint – leading to widespread instability. Such pitfalls can be avoided were system-wide stress tests wielded to evaluate the systemic implications of prudential policies.

This thesis develops a *generic*² framework for system-wide stress testing – suitable for assessing systemic risk and evaluating policies.

Building on the framework, several novel models are implemented that are tailored to answer two highly relevant policy questions:

- 1. How does regulatory buffer size and usability affect systemic risk?**
- 2. What are the systemic implications of the bail-in design?**

The first two questions address the stability impact of two pillars of the post-crisis regulatory design: Basel III and new resolution regime. In addition to a review of the stress test application of financial stability models, each question is addressed in a separate chapter of this thesis. For each research question, in the remainder of this introduction we will discuss: (a) the motivation for studying the question; (b) the method we drew on; (c) our main findings; (d) the contribution to the literature; and (e) the implications of our research. Herein, we will also draw some comparisons between the research methods and findings of each chapter.

Financial Stability Models and Their Application in Stress Tests

This chapter reviews heterogeneous agent models of financial stability and their application in stress tests. In contrast to the mainstream approach, which relies heavily on the rational expectations assumption and focuses on situations where it is possible to

²A stress test framework is said to be generic if it is able to coherently host a suit of different stress testing models, based on different data, and focussing on different policy questions.

compute an equilibrium, this approach typically uses stylised behavioural assumptions and relies more on simulation. This makes it possible to include more actors and more realistic institutional constraints, and to explain phenomena that are driven by out of equilibrium behaviour, such as clustered volatility and fat tails. We argue that traditional equilibrium models and agent-based models are complements rather than substitutes, and review how the interaction between these two approaches has enriched our understanding of systemic financial risk. We also review the network aspects of systemic risk, including models for three key channels of contagion: counterparty loss, overlapping portfolios and funding liquidity. We provide an overview of applications to microprudential and macroprudential stress testing.

Foundations of System-Wide Stress Testing

In a highly connected financial system, seemingly localised shocks can be propagated and amplified to take on systemic importance. However widely recognised, this reality is not yet reflected in the design of financial stress tests; notwithstanding the substantial progress that has been made, stress testing frameworks lack the holistic quality required to coherently combine multiple interacting contagion and amplification mechanisms as well as the behavioural responses of heterogeneous financial institutions to shocks. This raises the concern that current stress test results may provide false comfort to regulators, markets, and the public at large.

We propose a structural framework for the development of system-wide financial stress tests with multiple interacting contagion, amplification channels and heterogeneous financial institutions. This framework conceptualises financial systems through the lens of five building blocks: financial institutions, contracts, markets, constraints, and behaviour. These blocks can be flexibly implemented to form a dynamic multiplex network using the accompanying simulation engine and software library (the ‘Economic Simulation Library’, or ESL). Using this framework, we implement a system-wide stress test for the European financial system that incorporates amplification risks associated with default contagion, price-mediated contagion via asset sales, funding contagion, and liquidity stress via margin calls. We apply this stress test model to data provided by *S&P Global Market Intelligence*, the *ECB Statistical Warehouse*, the 2018 *European Banking Authority* (EBA) stress test results, allowing us to initialise balance sheets of European banks and non-banks. In line with [Hałaj & Kok \(2013\)](#), [Kok & Montagna \(2013\)](#), we reconstruct the interbank, secured funding and common asset holding networks, which interconnect these institutions. We compare our findings to the European Banking Authority’s stress test from 2018 and find that our system-wide approach reveals hidden weaknesses in the resilience of the financial system: we find unambiguously that shocks

are amplified by the interaction of heterogeneous agents and multiple contagion mechanisms. While current microprudential stress tests remain valuable, our findings suggest that they should be complemented by system-wide stress tests when evaluating financial stability and calibrating the size and designing the ‘usability’ of capital buffers.

Critically, we find that the willingness of banks to draw on their capital buffers to absorb losses – which we term the ‘usability’ of capital buffers – significantly affects the shock-amplifying tendency of a financial system. The actions banks take to avoid using their buffers in response to an adverse shock, which could for example be motivated by a desire to avoid regulatory restrictions on dividend payments, can generate pro-cyclical dynamics that substantially increase system-wide losses. In light of this result, regulators should evaluate how the design and enforcement of regulatory buffers may affect their ‘usability’ in times of financial stress, and be mindful of the financial stability implications of buffers that produce behavioural effects similar to those of regulatory requirements (Goodhart et al. (2008), Goodhart (2013)).

To the best of our knowledge, we are the first to propose a generic framework for system-wide stress testing with heterogeneous financial institutions. The multi-layered network model, which we implement using this structural framework, contributes to a growing body of systemic risk literature that captures multiple mechanisms with which shocks can be endogenously amplified (see e.g. Caccioli et al. (2013), Kok & Montagna (2013), Poledna et al. (2015), Hüser & Kok (2019), Wiersema et al. (2019)). In contrast to previous proposals, our model also captures liquidity stress via margin calls – on top of exposure loss contagion, overlapping portfolio contagion and funding contagion.

Our findings quantitatively examine the systemic impact of the ‘usability’ of Basel III regulatory buffers. While the system-wide consequences of the ‘usability’ of buffers have, hitherto, been merely qualitatively conceptualised, for instance by Goodhart et al. (2008), Goodhart (2013). Up to now, quantitative network models of financial stability have wrongly taken each portion of the bank’s capital pile to be equally usable or not usable (see e.g. Battiston et al. (n.d.), Caccioli et al. (2013), Greenwood et al. (2015)). This ignores that capital below the minimum capital requirements is typically less usable than the capital that comprises regulatory buffers, and it fails to recognise that banks may consider their capital pile in excess of the regulatory buffers to be more usable than their capital stack constituting the regulatory buffer.

Systemic Implications of the Bail-In Design

The 2007-2008 financial crisis forced governments to choose between the unattractive alternatives of bailing out a systemically important bank (SIBs) or having it fail in a system-wide disruptive manner (Bernanke (2017)). Bail-in has been put forward as the

primary tool to resolve a failing bank, which would end the ‘too-big-to-fail’ problem by letting stakeholders shoulder the losses, while minimising the calamitous impact of a bank’s failure on the financial system and the economy in general (e.g. [FSB \(2013\)](#), [Chennells & Wingfield \(2015\)](#), [BoE \(2017a\)](#)). Though the aptness of bail-in has been evinced in cases of relatively minor idiosyncratic bank failures, its efficacy in maintaining stability in cases of large bank failures and episodes of system-wide crises remains to be practically tested.

This paper investigates the stability implications of the bail-in design, for all these cases. We do so by developing a multi-layered network model of the European financial system, which extends [Farmer et al. \(2020\)](#) described in Chapter 3. It captures the pertinent endogenous-amplification mechanisms: exposure loss contagion, overlapping portfolio contagion, funding contagion, bail-inable debt revaluations, and bail-inable debt runs.

We apply this stress test model to data provided by *S&P Global Market Intelligence*, the *ECB Statistical Warehouse*, the 2018 *European Banking Authority* (EBA) stress test results, allowing us to initialise balance sheets of European banks and non-banks, as well as decompose banks’ liabilities in seniority classes. In line with [Halaj & Kok \(2013\)](#), [Kok & Montagna \(2013\)](#), we reconstruct the bank debt and common asset holding networks, which interconnect these institutions. The loss absorbing requirements, which set a minimum on the amount of long-term loss-absorbing debt that banks should issue (and cannot be cross-held by banks), further inform the calibration of the maturity profile and non-bank holdings of bail-inable debt.

Our results reveal that stability hinges on a set of ‘primary’ and ‘secondary’ bail-in parameters, including the failure threshold, recapitalisation target, debt-to-equity conversion rate, loss absorption requirements, debt exclusions and bail-in-design certainty. We show that bail-in usually works when relatively small European SIBs idiosyncratically fail regardless of the elected bail-in parameters – consistent with previous experience ([WBG \(2017\)](#)). However, we find that bailing-in banks in a system-wide crisis may heftily exacerbate financial fragility when bail-in policy parameters are set ‘badly’; and we observe that ill-designed bail-ins may induce widespread contagion, if large European D-SIBs idiosyncratically fail. Strikingly, we observe a phase shift from an unstable to stable system, if resolution authorities do choose ‘good’ bail-in parameters. Instability remains curbed even if systemic effects cause multiple banks to be bailed-in amid pervasive distress. Our evidence fortunately suggests that the pivot for stability is in the hands of policymakers. It also suggests, however, that the current policy parameters might be in the regime of instability.

Our contribution adds to the nascent network literature on the systemic effects of bail-in. [Klimek et al. \(2015\)](#) employ an agent-based network model to evaluate the economic

and financial ramifications of bail-in, and they compare its performance against other resolution mechanisms. [Hüser et al. \(2017\)](#) evaluate the systemic implications of bail-in in the European Union, drawing on a calibrated multi-layered network model to bank debt and equity cross-holdings. These papers neither investigate the systemic impact of the bail-in design, nor include multiple contagion mechanisms and non-banks in their analyses. Instead they take the bail-in design as is and merely explore the repercussions of exposure loss contagion. By ignoring multiple interaction contagion mechanisms, they risk underestimating the systemic footprint of the bail-in design. It is worth noting that though bail-in has been designed with systemic considerations in mind,³ it is not enough to assert its suitability on a system-wide scale. As [Aymanns et al. \(2016\)](#) have shown in the case of the Basel II leverage requirements, well-intended microprudential regulation may undermine financial resilience when systemic feedbacks are taken into account. This makes the investigation of the stability implications of the bail-in design in a networked financial system a necessary gap to fill.

Organisation of thesis

The remainder of this thesis presents each of the papers we just discussed in a separate chapter. Chapter 5 concludes.

³See: Directive 2014/59/EU of the European Parliament and of the Council.

Chapter 2

Models of Financial Stability and Their Application in Stress Tests

2.1 Summary

We review heterogeneous agent models of financial stability and their application in stress tests. In contrast to the mainstream approach, which relies heavily on the rational expectations assumption and focuses on situations where it is possible to compute an equilibrium, this approach typically uses stylized behavioral assumptions and relies more on simulation. This makes it possible to include more actors and more realistic institutional constraints, and to explain phenomena that are driven by out of equilibrium behavior, such as clustered volatility and fat tails. We argue that traditional equilibrium models and agent-based models are complements rather than substitutes, and review how the interaction between these two approaches has enriched our understanding of systemic financial risk. After presenting a brief summary of key terminology, we review models for leverage and endogenous risk dynamics. We then review the network aspects of systemic risk, including models for the three main channels of contagion: counterparty loss, overlapping portfolios and funding liquidity. Finally, we give an overview of applications to stress testing, including both microprudential and macroprudential stress tests.

Authors of Paper - Christoph Aymanns, J. Dooyne Farmer, Alissa M. Kleinnijenhuis (first author) & Thom Wetzler. The full paper is found in Aymanns et al. (2018) .

2.2 Introduction

The financial system is a classic example of a complex system. It consists of many diverse actors, including banks, mutual funds, hedge funds, insurance companies, pension funds and shadow banks. All of them interact with each other, as well as interacting directly with the real economy (which is undeniably a complex system in and of itself). The financial crisis of 2008 provided a perfect example of an emergent phenomenon, which is

the hallmark of a complex system.

While the causes of the 2008 turmoil remain controversial, the crisis has made everyone aware of the complex nature of the interactions and feedback loops in the economy, and it has driven an explosive amount of research attempting to better understand the financial system from a systemic point of view. It has also underlined the policy relevance of the complex systems approach. Systemic risk occurs when the decisions of individuals, which might be prudent if considered in isolation, combine to create risks at the level of the whole system that may be qualitatively different from the simple combination of their individual risks. By its very nature systemic risk is an emergent phenomenon that comes about due to the nonlinear interaction of individual agents. To understand systemic risk, we need to understand the collective dynamics of the system that gives rise to it.

In this paper, we review financial stability models, which aim to capture these dynamics, and their application in (regulatory) stress tests. Stress tests subject financial institutions, or the system as a whole, to an adverse, coherent, and plausible crisis scenario to assess systemic risk and evaluate prudential policies. To the best of our knowledge, we are the first to discuss how financial stability models are applied in stress tests, as well as contrast the two key types of applied models: equilibrium and agent-based models.

First, we briefly compare two basic approaches to elucidating the underlying mechanisms driving financial stability: traditional equilibrium models and agent-based models. The mainstream approach has been to focus on situations where it is possible to compute an equilibrium. This generally requires making very strong simplifications, e.g. studying only a few actors and interactions at a time. The equilibrium approach has been useful to clarify some of the key mechanisms driving financial instabilities and financial contagion, but it comes at the expense of simplifications that limit the realism of the conclusions. There is also a concern that, particularly during a crisis, the assumptions of rationality and equilibrium are too strong.

The alternative approach abandons equilibrium and rationality and replaces them with behavioral assumptions. This approach often relies on simulation, which has the advantage that it is easier to study more complicated situations, e.g. with more actors and more realistic institutional constraints. It also makes it possible to study multiple channels of interaction; even though research in this direction is still in its early stages, it is clear that this plays an important role.¹

¹We note that both equilibrium models and agent-based models use analytical as well as simulation-based methods. Agent-based based models more frequently rely on simulation-based methods, since the heterogeneity and institutional detail that these models typically capture makes analytical solutions

The use of behavioral assumptions as an alternative to utility maximization is controversial. Unlike utility, actions, on which behavioral assumptions are based, have the advantage of being directly observable, and in many cases the degree to which they are followed can be confirmed empirically. The disadvantage of this approach is that behavior may be context dependent, and as a result, such models typically fail the Lucas critique. We will show examples here where models based on behavioral assumptions are nonetheless very useful because they make it possible to directly investigate the consequences of a given set of behaviors. We will show examples where it leads to simple models that make clear predictions, at the same time that it can potentially be extended to complex real-world situations.

This review will focus primarily on the simulation approach, though we will attempt to discuss key influences and interactions with the more traditional equilibrium approach. Our view is that the two approaches are complements rather than substitutes. The most appropriate approach depends on the context and the goals of the modeling exercise. We predict that the simulation approach will become increasingly important with time, for several reasons. One is that this approach can be easier to bring to the data, and data is becoming more readily available. Many central banks are beginning to collect comprehensive data sets that make it possible to monitor the key parts of the financial system. This makes it easier to test the realism of behavioral assumptions, making such models less ad hoc. With such models it is potentially feasible to match the models to the data in a literal, one-to-one manner. This has not yet been done, but it is on the horizon, and if successful such models may become valuable tools for assessing and monitoring financial stability, and for policy testing. In addition, computational power is always improving. This is a new area of pursuit and the computational techniques and software are rapidly improving.

Second, we examine key models of financial contagion due to interconnectedness, which are applied in stress tests. Since the actors in the financial system are highly interconnected, network dynamics plays a key role in determining financial stability. The distress of one institution can propagate to other institutions, a process that is often called *contagion*, based on the analogy to disease. We discuss multiple channels of contagion, including counterparty risk, funding risk, and common assets holdings. *Counterparty risk* is caused by the web of bilateral contracts, which make one institution's assets another's liabilities. When a borrower is unable to pay, the lender's balance sheet is affected, and the resulting financial distress may in turn be transmitted to other parties, causing them

unfeasible. Given that equilibrium models more often investigate simplified set-ups with representative sectors, an analytical solution is in many cases viable.

to come under stress or default. *Funding risk* occurs when a lender comes under stress, which may create problems for parties that routinely borrow from this lender because loans that they would normally expect to receive fail to be extended. Institutions are also connected in many indirect ways, e.g. by common asset holdings, also called *overlapping portfolios*. If an institution comes under stress and sells assets, this depresses prices, which can cause further selling, etc. There are of course other channels of contagion, such as common information, that can affect expectations and interact with the more mechanical channels described above.

These channels of contagion cause nonlinear interactions that can create positive feedback loops that amplify external shocks or even generate purely endogenous dynamics, such as booms and busts. Nonlinear feedback loops can also be amplified by behavioral and institutional constraints and by bounded rationality (often in the context of incomplete information and learning).

Finally, we study how the financial stability models discussed above are applied stress tests, and give examples of such stress tests. We start by discussing microprudential stress tests, which, as we point out, do not incorporate financial stability models of contagion at all. Instead, microprudential stress tests typically focus on evaluating the *direct* impact of a crisis scenario on the solvency of individual institutions, who are treated to be *unlinked*; thereby, failing to consider the *higher-order* contagion dynamics that may arise in the *interconnected* financial system as a consequence of this initial shock. We continue to study macroprudential stress tests, whose distinctive feature is that they evaluate the resilience of the system as a whole, as well as the institutions in it, leveraging on financial stability models that allow for contagion. We dissect a number of macroprudential stress test to highlight the use of contagion models.

The remainder of this paper is organized as follows: In Section 2.3 we briefly contrast and compare traditional equilibrium models with agent-based models. In Section 2.4 we introduce the dynamical systems perspective on the financial system that will underlie many of the models of financial stability that we discuss in subsequent sections. In Section 2.5 we discuss models of financial contagion due to interconnectedness. Sections 2.6 to 2.8 consider various different stress tests. In particular, Section 2.6 gives a brief conceptual overview of stress tests; Section 2.7 introduces and critically evaluates standard, microprudential stress tests; Section 2.8 discusses examples of macroprudential stress tests and how to bring them to data; finally Section 2.9 concludes.

2.3 Two Approaches to Modeling Systemic Risk

As mentioned in the introduction, traditionally finance has focused on modeling systemic risk in highly stylized models that are analytically tractable. These efforts have improved our understanding of a wide range of phenomena related to systemic risk ranging from bank runs (Diamond & Dybvig (1983), Morris & Shin (2001)), credit cycles (Kiyotaki & Moore (1997), Brunnermeier & Sannikov (2014)), balance sheet (Allen & Gale (2000)) and information contagion (Acharya & Yorulmazer (2008)) over fire sales (Shleifer & Vishny (1992)), to the feedback between market and funding liquidity (Brunnermeier & Pedersen (2009)). A comprehensive review that does justice to this literature is beyond the scope of this paper. However, we would like to make a few observations in regard to the traditional modeling approach and contrast it with the agent-based approach.

Traditional models place great emphasis on the incentives and information structure of agents in a financial market. Given those, agents behave strategically, taking into account their beliefs about the state of the world, and other agents' strategies. The objects of interest are then the game theoretic equilibria of this interaction. This allows for studying the effects of properties such as asymmetric information, uncertainty or moral hazard on the stability of the financial system. While these models provide valuable qualitative insights, they are typically only tractable in very stylized settings. Models are usually restricted to a small number or a continuum of agents, a few time periods and a drastically simplified institutional and market set up. This can make it difficult to draw quantitative conclusions from such models.

Agent-based models typically place less emphasis on incentives and information, and instead focus on how the dynamic interactions of behaviorally simple agents can lead to complex aggregate phenomena, such as financial crises, and how outcomes are shaped by the structure of this interaction and the heterogeneity of agents. From this perspective, the key drivers of systemic risk are the amplification of dynamic instabilities and contagion processes in financial markets. Complicated strategic interactions and incentives are often ignored in favor of simple, empirically motivated behavioral rules and a more realistic institutional and market set up. Since these models can easily be simulated numerically, they can in principle be scaled to a large number of agents and, if appropriately calibrated, can yield quantitative insights.

Two common criticisms leveled against heterogeneous agent-based models are that they often lack of strategic interactions and do not provide sufficiently robust results with respect to parameter choices. The first criticism is fair and, in many cases, highlights an important shortcoming of this approach. Hard-wired behavioral rules need to be carefully calibrated against real data, and even when they are, they can fail in new situations

where the behavior of agents may change. For computer simulations to be credible, their parameters need to be calibrated and the sensitivity of outcomes to those parameters needs to be understood. The latter in particular is more challenging in computational models than in tractable analytical models. Agent-based models of financial systems that capture heterogeneity often require more granular data to initialise and calibrate the models. This is becoming less of an obstacle, now that increasingly more big and granular datasets are being collected by central banks.

In our view, what is not fair is to regard computer simulations as inherently inferior to analytic results. Analytic models have the benefit of the relative ease with which they can be used to understand the concepts driving structural cause-and-effect relationships. But many aspects of the economic world are not simple, and in most realistic situations computer simulations are the only possibility. Good practice is to make code freely available and well documented, so that results are easily reproducible.

Traditional and heterogeneous agent-based models are complements rather than substitutes. Some heterogeneous agent-based models already use myopic optimization, and in the future the line between the two may become increasingly blurred.² As methods such as computational game theory or multi-agent reinforcement learning mature, it may become possible to increasingly introduce strategic interactions into computational heterogeneous agent-based models. Furthermore, as computational resources and large volumes of data on the financial system become more accessible, parameter exploration and calibration should become increasingly feasible. Therefore, we are optimistic that, provided technology progresses as expected,³ in the future heterogeneous agent-based models will be able to overcome some of the shortcomings discussed above. And as we demonstrate here, they have already led to important new results in this field, that were not obtainable via analytic methods.

2.4 A View of the Financial System

At a high level, it is useful to think of the financial system as a dynamical system that consists of a collection of institutions that interact via centralized and bilateral markets. An institution can be represented by its balance sheet, i.e. its assets and liabilities, together with a set of decision rules that it deploys to control the state of its balance sheet in order to achieve a certain goal. Within this framework, a market can be thought

² In fact, this is already the case in the literature on financial and economic networks, see for example [Goyal \(2018\)](#).

³ It seems unlikely that scientists' ability to analytically solve models will improve as quickly as numerical techniques and heterogeneous agent-based simulations, which benefit from rapid improvements in hardware and software.

of as a mechanism that takes actions from institutions as inputs and changes the state of their balance sheets based on its internal dynamics. Anyone wishing to construct an agent model of the financial system therefore has to answer three fundamental questions: (i) what comprises the institutions' balance sheets, (ii) what determines their actions conditional on the state of the world, and (iii) how do markets respond to these actions? In the following, we will sketch the balance sheet of a generic leveraged investor, which will serve as the fundamental building block of the models of financial stability that we will discuss in this review. We will also briefly touch on (ii) and (iii) when discussing the important channels through which leveraged investors interact. In the subsequent sections we will then discuss concrete models of financial stability that fall within this general framework.

2.4.1 Balance sheet composition

When developing a model of a financial system, it is useful to distinguish between two types of agents which we refer to as “active” and “passive” agents. Active agents are the objects of interest and their internal state and interactions are carefully modeled. Passive agents represent parts of the financial system that interact with active agents but are not the focus of the model, and are therefore typically represented by simple stochastic processes. For the remainder of this review, consider a financial system that consists of a set \mathcal{B} of active leveraged investors and a set of passive agents which will remain unspecified for now. We are particularly interested in systemic risk that is driven by borrowing, and thus we focus on agents that use leverage (defined as purchasing assets with borrowed funds). However, the setup below is sufficiently general to accommodate unleveraged investors as a special case with leverage equal to one.

Leveraged investors need not be homogeneous and may differ, among other aspects, in their balance sheet composition, strategies or counterparties. In practice, a leveraged investor might be a bank or a leveraged hedge fund and other active investors might include unleveraged mutual funds. Passive agents could be depositors, noise traders, fund investors that generate investment flows or banks that lend to hedge funds. The choice of active vs. passive investors of course varies from model to model.

The balance sheet of an investor $i \in \mathcal{B}$ is composed of assets A_i , liabilities L_i and equity E_i , such that $A_i = L_i + E_i$. The investor's leverage is simply the ratio of assets to equity $\lambda_i = A_i/E_i$. It is useful to decompose the investor's assets into three classes: bilateral contracts A_i^B between investors, such as loans or derivative exposures; traded securities A_i^S , such as stocks; and external assets A_i^R , whose value is assumed exogenous. Throughout this review, we assume that the value A_i^S of traded securities is marked to

market.⁴ That is, the value of a traded security on the investor's balance sheet will be determined by its current market price. Of course we must have that $A_i = A_i^B + A_i^S + A_i^R$. Each asset class can be further decomposed into individual loan contracts, stock holdings and so on.

The investor's liabilities can be decomposed in a similar fashion. For now, let us decompose the investor's liabilities simply into bilateral contracts L_i^B between investors, such as loans or derivative exposures, and external liabilities L_i^R which can be assumed exogenous. In the case of a bank, these external liabilities might be deposits. Again we must have that $L_i = L_i^B + L_i^R$, and bilateral liabilities can be further decomposed into individual bilateral contracts. Bilateral assets and liabilities might be secured, such as repurchase agreements, or unsecured such as interbank loans. Naturally, bilateral liabilities are just the flip side of bilateral assets such that summing over all investors we must have $\sum_i A_i^B = \sum_i L_i^B$.

2.4.2 Balance sheet dynamics

Of all the factors that affect the dynamics of the investors' balance sheets, three are of particular importance for financial stability: leverage, liquidity and interconnectedness. Below, we discuss each factor in turn.

Leverage: Leverage amplifies returns, both positive and negative. Therefore, investors typically face a leverage constraint to limit the investors' risk.⁵ However, at the level of the financial system, binding leverage constraints can lead to substantial instabilities. On short time scales, a leveraged investor may be forced to sell into falling markets when she exceeds her leverage constraint. Her sales will in turn depress prices further as we explain in the next paragraph on market liquidity. Leverage constraints can thus lead to an unstable feedback loop between falling prices and forced sales. On longer time scales dynamic leverage constraints that depend on backward looking risk estimates can lead to entirely endogenous volatility – so called leverage cycles.⁶

⁴ The term *marked to market* means that the value of assets is recomputed in every period based on current market prices. This is in contrast to valuing assets based on an estimate of their fundamental value.

⁵ This constraint may be imposed by a regulator, a counterparty or internal risk management.

⁶ Beyond leverage, investors may also face other constraints. Regulators have imposed restrictions on the liquidity of assets that some investors may hold (with a preference for more liquid assets) and the stability of their funding (with a preference for more long term funding). The effect of these constraints on systemic risk is much less studied than the effect of leverage constraints. A priori however, one would expect these constraints to improve stability. This is because of the absence of feedback loops similar to the leverage-price feedback loop that drives forced sales.

Liquidity: Broadly speaking, one can distinguish between two types of liquidity: market liquidity and funding liquidity.

Market liquidity can be understood as the inverse to price impact. When market liquidity is high, the market can absorb large sell orders without large changes in the price. If markets were perfectly liquid it would always be possible to sell assets without affecting prices and most forms of systemic risk would not exist.⁷ Leverage is dangerous both because it directly increases risk, amplifying gains and losses proportionally, but also because the market impact of liquidating a portfolio to achieve a certain leverage increases with leverage. This point was stressed by [Caccioli, Bouchaud & Farmer \(2012\)](#), who showed how, due to her own market impact, an investor with a large leveraged position can easily drive herself bankrupt by liquidating her position. They showed that this can be a serious problem even under normal market conditions, and recommended taking market impact into account when valuing portfolios in order to reduce this problem. The problem can become even worse if investors are forced to sell too quickly, inducing *fire sales* in which a market is overloaded with sell orders, causing a dramatic decrease in liquidity for sellers.⁸ Fire sales can be induced when investors hit leverage constraints, forcing them to sell, which in turn causes leverage constraints to be more strongly broken, inducing more selling.

Funding liquidity refers to the ease with which investors can borrow to fund their balance sheets. When funding liquidity is high, investors can easily roll over their existing liabilities by borrowing again, or even expand their balance sheets. In times of crises, funding liquidity can drop dramatically. If investors rely on short term liabilities they may be forced to liquidate a large part of their assets to pay back their liabilities. This forced sale can trigger fire sales by other investors.

Interconnectedness: Investors are connected via their balance sheets and so are not isolated agents. Connections can result from direct exposures due to bilateral loan contracts, or from indirect exposures due to investments into the same assets. Interconnectedness together with feedback loops resulting from binding leverage constraints and endogenous liquidity can lead to financial contagion. In analogy to epidemiology, financial contagion refers to the process by which “distress” may spread from one investor to another, where distress can be broadly understood as an investor becoming uncomfortably close to insolvency or illiquidity. Typically financial contagion arises when, via some mechanism or channel, a distressed investor’s actions negatively affect some subset of other

⁷ Prices would of course still change to reflect the arrival of new information.

⁸ There is always market impact from buying or selling. The term “fire sales” technically means selling under stress, but often means simply a case where the sale of assets is forced (even when markets remain orderly). See the discussion in Section 5.3.

investors. This subset of investors is said to be connected to the distressed investor. A simple example of such connections are the bilateral liabilities between investors. Taken together, the set of all such connections form a network over which financial contagion can spread. For an in-depth review of financial networks see [Iori & Mantegna \(2018\)](#).

The aim of the subsequent section is to discuss models of financial contagion as they form the scientific bedrock of the stress testing models that will be discussed in Sections 2.6 and beyond. While liquidity is of great importance, we will only discuss it implicitly in Section 2.5, rather than dedicating a separate section to it. This is because, unfortunately, there are currently only few dedicated models on this topic, see [Bookstaber & Paddrik \(2015\)](#) for an example. We will not be able to provide a complete overview of the agent-based modelling literature devoted to various aspects of financial stability. Important topics that we will not be able to discuss include the role of heterogeneous expectations or time scales in the dynamics of financial markets, see for example [Hommes \(2006\)](#), [LeBaron \(2006\)](#) for early surveys and [Dieci & He \(2018\)](#) for a recent overview.

2.5 Contagion in Financial Networks

2.5.1 Financial linkages and channels of contagion

A channel of contagion is a mechanism by which distress can spread from one financial institution to another. Often the channel of contagion is such that distress can only spread from one institution to a subset of all institutions in the system. These susceptible institutions are said to be linked to the stressed institution. The set of all links then forms a financial network associated with the channel of contagion.⁹ Depending on the channel, links in this network may arise directly from bilateral contracts between banks, such as loans, or indirectly via the markets in which the banks operate. In the literature, one typically distinguishes between three key channels of contagion: counterparty loss, overlapping portfolios and funding liquidity contagion.¹⁰ Counterparty loss and overlapping portfolio contagion affect the value of the assets on the investors' balance sheet while funding liquidity contagion affects the availability of funding for the investors' balance sheets. In the following we will first introduce the investor's balance sheet relevant for this section. We will then give a brief overview of the channels of contagion before discussing each in more detail.

⁹ See [Iori & Mantegna \(2018\)](#) for a review of financial networks.

¹⁰ Information contagion (cf. [Acharya & Yorulmazer \(2008\)](#)) is another channel of contagion but won't be discussed in this section.

Balance sheet: Throughout this chapter we will consider a set \mathcal{B} of leveraged investors (banks for short) whose assets can be decomposed into three classes: bilateral interbank contracts A_i^B , traded securities A_i^S that are marked to market and external, unmodeled assets A_i^R . Furthermore bank liabilities can be decomposed into bilateral interbank contracts L_i^B , and external, unmodeled liabilities L_i^R such that $L_i = L_i^B + L_i^R$. Note that bilateral interbank contract need not be loans, they can also be derivative contracts for example. For simplicity however, we will think of bilateral interbank contracts as loans for the remainder of this section.

Counterparty loss: Suppose bank i has lent an amount C to bank j such that $A_i^B = L_j^B = C$. Now suppose the value of bank j 's external assets A_j^R drops due to an exogenous shock. As a result the probability of default of bank j is likely to increase, which will affect the value of the claim A_i^B that bank i holds on bank j . If bank i 's interbank assets are marked to market, a change in bank j 's probability of default will affect the market value of A_i^B . In the worst case, if bank j defaults, bank i will only recover some fraction $r \leq 1$ of its initial claim A_i^B . If the loss of bank i exceeds its equity, i.e. $(1 - r)A_i^B > E_i$, bank i will default as well.¹¹ Now, how can this lead to financial contagion? To elaborate on the above stylised example, suppose that bank i in turn borrowed an amount C from another bank k such that $A_k^B = L_i^B = C$.¹² In this scenario, it can be plausibly argued that an increase in the probability of default of j increases the probability of default of i which in turn increases the probability of default of k . If all banks mark their books to market, an initial shock to j can therefore end up affecting the value of the claim that bank k holds on bank i . Again, in the extreme scenario, the default of bank j may cause bank i to default which may cause bank k to default. This is the essence of counterparty loss contagion. Naturally, in a real financial system the structure of interbank liabilities will be much more complex than in the stylised example outlined above. However, the conceptual insights carry over: the financial network associated with the counterparty loss contagion channel is the network induced by the set of interbank liabilities.

Overlapping portfolios: The overlapping portfolio channel is slightly more subtle. Suppose bank i and bank j have both invested an amount C in the same security l such that $A_{il}^S = A_{jl}^S = C$, where we have introduced the additional index to reference the security.¹³ Now, suppose the value of bank j 's external assets A_j^R drops due to some

¹¹ In reality, this scenario is excluded due to regulatory large exposure limits which require that $A_i^B < E_i$.

¹² We assume that the contract between i and j as well as i and k has the same notional purely for expositional simplicity and all conceptual insights carry over for heterogenous notionals.

¹³ Again, we assume that both banks invest the same amount purely of expositional simplicity.

exogenous shock. How will bank j respond to this loss? In the extreme case, when the exogenous shock causes bank j 's bankruptcy ($E_i < 0$), the bank will liquidate its entire investment in the security in a fire sale. However, even if the bank does not go bankrupt, it may wish to liquidate some of its investment. This can occur for example when the bank faces a leverage constraint. Bank j 's selling is likely to have price impact. As a result, the market value of A_{il}^S will fall. If bank i also faces a leverage constraint, or even goes bankrupt following the fall in prices, it will liquidate part of its securities portfolio in response. How will this lead to contagion? Suppose that bank i also has invested an amount C into another security m and that another bank k has also invested into the same security, such that $A_{im}^S = A_{km}^S = C$. If bank i liquidates across its entire portfolio, it will sell some of security m following a fall in the price of security l . The resulting price impact will then affect the balance sheet of bank k which was not connected to bank j via an interbank contract or a shared security. This is the essence of overlapping portfolio contagion. Banks are linked by the securities that they co-own and the fact that they liquidate with market impact across their entire portfolios. Empirical evidence from the 2007 Quant meltdown for this contagion channel has been provided in [Khandani & Lo \(2011\)](#).

Funding liquidity contagion often occurs when a lender is stressed, and so often occurs in conjunction with overlapping portfolio contagion and counterparty loss contagion. To see this, let us reconsider the scenario we discussed for counterparty loss contagion. Suppose bank i has lent an amount C to bank j such that $A_i^B = L_j^B = C$. As before, suppose the value of bank j 's external assets A_j^R drops due to some exogenous shock and as a result, the probability of default of bank j increases. Now, suppose that every T periods bank i can decide whether to roll over its loan to bank j . Further assume that bank i is bank j 's only source of interbank funding and L_j^R is fixed. Given bank j 's increased default probability, bank i may choose not to roll over the loan at the next opportunity. Ignoring interest payments, if bank i does not roll over the loan, bank j will have to deliver an amount C to bank i . In the simplest case, bank j may choose not to roll over its own loans to other banks which in turn may decide against rolling over their loans. This is the essence of funding liquidity contagion. As for counterparty loss contagion, the associated financial network is induced by the set of interbank loans. Empirical evidence on the fragility of funding markets during the past financial crisis has been provided for example by [Afonso et al. \(2011\)](#), [Iyer & Peydro \(2011\)](#). In a further complication, bank j may also choose to liquidate part of its securities portfolio in order to pay back its loan. Funding liquidity contagion can therefore lead to fire sales and overlapping portfolio contagion and vice versa. This interdependence of contagion chan-

nels makes the funding liquidity and overlapping portfolio contagion processes the most challenging from a modeling perspective.

In the remainder of this section, we will discuss models for counterparty loss, overlapping portfolio and funding liquidity contagion, as well as models for the interaction of all three contagion channels.

2.5.2 Counterparty loss contagion

Let P denote the matrix of nominal interbank liabilities such that banks hold interbank assets $A_i^B = \sum_j P_{ij}^T$, where T denotes the matrix transpose. In addition, banks hold external assets A_i^R which can be liquidated at no cost. Banks have interbank liabilities $L_i^B = \sum_j P_{ij}$ only. Assume all interbank liabilities mature at the same time and have the same seniority. We further assume that all banks are solvent initially. There is only one time period. At the end of that period all liabilities mature, external assets are liquidated and banks pay back their loans if possible. Now suppose banks are subject to a shock $s_i \geq 0$ to the value of their external assets such that $\hat{A}_i^R = A_i^R - s_i$. Given an exogenous shock, we can ask a number of questions. First, which loan payments are feasible given the exogenous shock? Second, which banks will default on their liabilities? And finally, how do the answers to the first two questions depend on the structure of the interbank liabilities P ? There is a large literature that studies counterparty loss contagion in a set up similar to the above, including Eisenberg & Noe (2001), Gai & Kapadia (2010), May & Arinaminpathy (2010), Elliott et al. (2014), Acemoglu et al. (2015), Battiston et al. (n.d.), Amini et al. (2013) and Capponi et al. (2015). In the following, we will briefly introduce the seminal contribution by Eisenberg & Noe (2001), who provide a solution to the first two questions. We will then consider a number of extensions of Eisenberg & Noe (2001) and alternative approaches to addressing the above questions.

Define the relative nominal interbank liabilities matrix as $\Pi_{ij} = P_{ij}/L_i^B$ for $L_i^B > 0$ and $\Pi_{ij} = 0$ otherwise. The relative liabilities matrix corresponds to the adjacency matrix of the weighted, directed network \mathcal{G} of interbank liabilities. Let $\mathbf{p} = (p_1, \dots, p_N)$ denote the vector of total payments made by the banks when their liabilities mature, where $N = |\mathcal{B}|$. Naturally, a bank pays at most what it owes in total, i.e. $p_i \leq L_i^B$. However, it may default and pay less if the value of its external assets plus the payments it receives from its debtors is less than what it owes. The individual payments that bank i makes are given by $\Pi_{ij}p_j$ since by assumption all liabilities have equal seniority. The vector of payments, also known as the clearing vector, that satisfies these constraints is the solution to the following fixed point equation

$$p_i = \min\{L_i^B, \hat{A}_i^R + \sum_j \Pi_{ij}^T p_j\}. \quad (2.1)$$

Eisenberg & Noe (2001) show that such a fixed point always exists. In addition, if within each strongly connected component of \mathcal{G} there exists at least one bank with $\hat{A}_i^R > 0$, Eisenberg & Noe (2001) show that the fixed point is unique.¹⁴ In other words, there exists a unique way in which losses incurred due to the adverse shock $\{s_i\}$ are distributed in the financial system via the interbank liabilities matrix. The clearing vector and the set of defaulting banks can be found easily numerically by iterating the fixed point map in Eq. (2.1). As the map is iterated, more and more banks may default, resulting in a default cascade propagating through the financial network.

It is important to note that in this set up losses are only redistributed and the system is conservative – contagion acts as a distribution mechanism but does not, in the aggregate, lead to any further losses to bank shareholders beyond the initial shock. To see this, define the equity of bank i prior to the exogenous shock as $E_i = A_i^B + A_i^R - L_i^B$ and after the exogenous shock as $\hat{E}_i = \hat{A}_i^B(\mathbf{p}) + A_i^R - s_i - \hat{L}_i^B(\mathbf{p})$. Note that post-shock both bank i 's assets and liabilities depend on the clearing vector \mathbf{p} . Taking the difference and summing over all banks we obtain $\sum_i E_i - \hat{E}_i = \sum_i A_i^R - (A_i^R - s_i) = \sum_i s_i$ since $\sum_i A_i^B = \sum_i L_i^B$ and $\sum_i \hat{A}_i^B(\mathbf{p}) = \sum_i \hat{L}_i^B(\mathbf{p})$. Also note that, while bank shareholder losses are not amplified, losses to the total value of bank assets are amplified due to indirect losses, i.e. losses not stemming from the initial exogenous shock but due to revaluation of interbank loans. This can be seen by taking the difference between pre- and post-shock total assets in the system. The total pre-shock assets of bank i are $A_i = A_i^B + A_i^R$ and its total post-shock assets are $\hat{A}_i = \hat{A}_i^B(\mathbf{p}) + A_i^R - s_i$, then $\sum_i A_i - \hat{A}_i = \sum_i A_i^B - \hat{A}_i^B(\mathbf{p}) + s_i \geq \sum_i s_i$. Some authors argue that this total asset loss can be useful measure of systemic impact of the exogenous shock, see Glasserman & Young (2015). Finally, note that the mechanism of finding a clearing vector ignores any potential frictions in the financial system and ensures that the maximal payment is made given the exogenous shocks. Several authors have argued that this is too optimistic and assume instead that once a default has occurred, some additional bankruptcy costs are incurred, see for example Rogers & Veraart (2013) and Cont et al. (2010).¹⁵¹⁶ In this case, aggregate bank shareholder losses may be larger than the aggregate exogenous shock. Further shortcomings of the Eisenberg and Noe model include the lack of heterogeneous seniorities or maturities and the lack of the possibility of strategic default.

¹⁴ In a strongly connected component of a directed graph there exists a directed path from each node in the component to each other node in the component. The strongly connected component is the maximal set of nodes for which this condition holds.

¹⁵ Such bankruptcy cost might for example capture the cost of forced liquidation of the banks external assets.

¹⁶ Papers that do not assume bankruptcy costs are essentially treating the system as if were conservative: losses to one party are gains to the other, but there is no deadweight loss that ravages welfare. Hence, the fail to capture the negative externalities imposed by the banking system on society.




Figure 2.1: Expected number of defaults as a function of diversification in Elliott et al. (2014).

The extent of the default cascade triggered by an exogenous shock depends on the structure of the financial network induced by the matrix of interbank liabilities P . One key property of the financial network is the average degree of a bank, i.e. the number of other banks it lends to. A well-known result is that, as banks' interbank lending A_i^B becomes more diversified over \mathcal{B} , i.e. the average degree increases, the expected number of defaulting banks first increases and then decreases, see Fig. 2.1. If banks lend only to a very small number of other banks, the network is not fully connected. Instead, it consists of several small and disjoint components. A default in a particular component cannot spread to other components, hence limiting the size of the default cascade. As banks become more diversified, the network will become fully connected and default cascades can spread across the entire network. As banks diversify further, the size of the individual loans between banks declines to the point that the default of any one counterparty becomes negligible for a given bank. Thus default cascades become unlikely. However, if they do occur, they will be very large. This is often referred to as the “robust-yet fragile” property of financial networks and has been observed for specifications of the financial network and the default cascade mechanism, see for example Elliott et al. (2014), Gai & Kapadia (2010), Battiston et al. (n.d.) or Amini et al. (2013). However, not only the average of the network's degree distribution is important for the system's stability. Caccioli, Catanach & Farmer (2012) show that if the degree distribution is very heterogeneous, i.e. there are a few banks that lend to many banks while most only lend to a few, the system is more resilient to contagion triggered by the failure of a random bank, but more fragile with respect to contagion triggered by the failure of highly connected nodes. In addition, Capponi et al. (2015) show that the level of concentration of the liability matrix, as defined by a majorization order, can qualitatively change the system's loss profile.

The models and solution methods discussed above tend to be simple to remain tractable and usually reduce to finding a fixed point.¹⁷ However, these equilibrium models

¹⁷ Gai & Kapadia (2010) for example make similarly restrictive assumptions on the structure of bank

often form useful starting points for heterogeneous agent models that try to incorporate additional dynamic effects and more realism into the counterparty loss contagion process. See for example [Georg \(2013\)](#) where the effect of a central bank on the extent of default cascades is studied.

Finally, note that it is widely believed that large default cascades are quite unlikely for reasonable assumptions about the distribution of the exogenous shock and nominal interbank liabilities matrix, see for example [Glasserman & Young \(2015\)](#). For larger cascades to occur, default costs or additional contagion channels are necessary. Nevertheless, the existence of a counterparty loss contagion channel is important in practice as it affects the decisions of agents, for example in the way they form lending relationships. In other words, while default cascades are unlikely to occur in reality, they form an “off-equilibrium” path that shapes reality, see [Elliott et al. \(2014\)](#).

2.5.3 Overlapping portfolio contagion

In the following we will formally discuss the mechanics of overlapping portfolio contagion. To this end, consider again our set of banks \mathcal{B} . There is an illiquid asset whose value is exogenous and a set of securities \mathcal{S} , with $M = |\mathcal{S}|$, traded by banks at discrete points in time $t \in \mathbb{N}$. Let $\mathbf{p}_t = (p_{1t}, \dots, p_{Mt})$ denote the vector of prices of the securities and let the matrix $\mathbf{S}_t \in \mathbb{R}^{N \times M}$ denote the securities ownership of all banks at time t . Thus S_{ijt} is the position that bank i holds in security j at time t . The assets of bank i are then given by $A_{it} = \mathbf{S}_{it} \cdot \mathbf{p}_t + A_i^R$, where A_i^R is the bank’s illiquid asset holding. Let E_{it} and $\lambda_{it} = A_{it}/E_{it}$ denote bank i ’s equity and leverage, respectively. There are no interbank assets or liabilities.

As mentioned above, overlapping portfolio contagion occurs when one bank is forced to sell and the resulting price impact forces other banks with similar asset holdings to sell. What might force banks to sell? In an extreme scenario, a bank might have to liquidate its portfolio if it becomes insolvent, i.e. $E_{it} < 0$. But even before becoming insolvent, a bank might be forced to liquidate part of its portfolio if it violates a leverage constraint $\bar{\lambda}$.¹⁸ Both of these were considered by [Caccioli et al. \(2014\)](#) and by [Cont & Schaanning \(2017\)](#). In fact [Caccioli et al. \(2014\)](#) showed that such pre-emptive liquidations only make the problem worse due to increasing the pressure on assets that are already stressed. (This is closely related to the problem that liquidation can in and of itself cause default as

balance sheets as [Eisenberg & Noe \(2001\)](#). In addition several technical assumptions on the structure of the matrix of liabilities are necessary to solve for the fixed point of non-defaulted banks via a branching process approximation.

¹⁸ Other “constraints” might also lead to forced sales and overlapping portfolio contagion. For example, investor redemptions that depend on past performance, as in [Thurner & Poledna \(2013\)](#), can force a fund to liquidate which may result in an overlapping portfolio contagion similar to the one induced by leverage constraints.

studied by [Caccioli, Bouchaud & Farmer \(2012\)](#)). Other papers that discuss the effects of overlapping portfolios include [Duarte & Eisenbach \(2015\)](#), [Greenwood et al. \(2015\)](#), [Cont & Wagalath \(2016, 2013\)](#). An important early contribution to this topic is [Cifuentes et al. \(2005\)](#).

Let us first discuss the simpler case where liquidation occurs only upon default. Suppose bank i is subject to an exogenous shock $s_i > 0$ that reduces the value of its illiquid assets to $\hat{A}_{it}^R = A_{it}^R - s_i$. If $s_i > E_{it}$, the bank becomes insolvent and liquidates its entire portfolio. Let $Q_{jt} = \sum_{i \in \mathcal{J}_t} S_{ijt}$ denote the total amount of security j that is liquidated by banks in the set \mathcal{J}_t of banks that became insolvent at time t . The sale of the securities is assumed to have market impact such that $p_{jt+1} = p_{jt}(1 + f_j(Q_{jt}))$, where $f_j(\cdot)$ is the market impact function of security j . [Caccioli et al. \(2014\)](#) assume an exponential form $f_j(x) = \exp(-\alpha_j x) - 1$, where x is volume liquidated and $\alpha_j > 0$ is chosen to be inversely proportional to the total shares outstanding of security j . In the next period, banks reevaluate their equity at the new securities prices. The change in equity is equal to $\Delta E_{it+1} = \sum_j S_{ijt} p_{jt} f_j(Q_{jt}) - s_i$. Note that in this setting we hold S_{ijt} fixed unless a bank liquidates its entire portfolio. Thus, banks who share securities with the banks that were liquidating in the previous period will suffer losses due to market impact. These losses may be sufficiently large for additional banks to become insolvent. If this occurs, contagion will spread and more banks will liquidate their portfolios, leading to further losses. Over the course of this default cascade, banks may suffer losses that did not share any common securities with the initially insolvent banks.

The evolution of the default cascade can be easily computed numerically by following the procedure outlined above until no further banks default. [Caccioli et al. \(2014\)](#) also show that the default cascade can be approximated by a branching process, provided suitable assumptions are made about the network structure. For their computations, [Caccioli et al. \(2014\)](#) assume that a given bank i invests into each security with a fixed probability μ_B/M , where μ_B is the expected number of securities that a bank holds. The bank distributes a fixed investment over all securities it holds. When μ_B/M is high, the portfolios of banks will be highly overlapping, i.e. banks will share many securities in their portfolios. Similar to the results for counterparty loss contagion, the authors find that as banks become more diversified, that is μ_B increases while M is held fixed, the probability of default (blue circles) first increases and then decreases, see [Fig. 2.2](#). The intuition for this result is again similar to the counterparty loss contagion case. If banks are not diversified, their portfolios are not overlapping and price impact from portfolio liquidation of one bank affects only a few banks. As banks become more diversified, their portfolios become more overlapping and price impacts spreads throughout the set of banks leading to large default cascades. Eventually, when they become sufficiently diversified, the losses




Figure 2.2: Blue circles: probability of contagion. Red squares: conditional on a contagion, the fraction of banks that fail (i.e. the extent of contagion). Taken from [Caccioli et al. \(2014\)](#).

resulting from a price change in an individual security become negligible and large default cascades become unlikely. However, when they do occur, they encompass the entire set of banks. Thus, here again the financial network displays the robust-yet fragile property. Interestingly, the authors also show that for a fixed level of diversification, there exists a critical bank leverage λ_{it} at which default cascades emerge. The intuition for this result is that, when leverage is low, banks are stable and large shocks are required for default to occur, as leverage grows banks become more susceptible to shocks and defaults occur more easily.

As mentioned above, banks are likely to liquidate a part of their portfolio even before bankruptcy, if an exogenous shock pushes them above their leverage constraint. This is the setting studied in [Cont & Schaanning \(2017\)](#) and [Caccioli et al. \(2014\)](#). In this case, the shocks for which banks start to liquidate as well as the amount liquidated are both smaller than in the setting discussed above. If banks breach their leverage constraint due to an exogenous shock s_i to the value of their illiquid assets, [Cont & Schaanning \(2017\)](#) require that banks liquidate a fraction Γ_i of their entire portfolio such that $((1 - \Gamma_i)\mathbf{S}_{it} \cdot \mathbf{p}_t + \hat{A}_{it}^R)/E_{it} = \bar{\lambda}$. The corresponding liquidated monetary amount for a security j is then $Q_{jt} = \sum_{i \in \mathcal{B}} \Gamma_i S_{ijt} p_{jt}$. Again, the sale of the securities is assumed to have market impact such that $p_{jt+1} = p_{jt}(1 + f_j(Q_{jt}))$. In contrast to [Caccioli et al. \(2014\)](#), the authors assume that the market impact function $f_j(x)$ is linear in x , where x is the total monetary amount sold rather than the number of shares. Similar market impact functions are used by [Greenwood et al. \(2015\)](#) and [Duarte & Eisenbach \(2015\)](#).

The shape and parameterization of the market impact function is crucial for the practical usage of models of overlapping portfolio contagion. There is a large body of market microstructure literature addressing this question. This literature begins with [Kyle \(1985\)](#), who derived a linear impact function under strong assumptions. More recent theoretical and empirical work indicates that under normal circumstances market

impact is better approximated by a square root function [Bouchaud et al. \(2008\)](#).¹⁹ A good example of an event in which firesales took place was the quant meltdown in August 2007. In this case, multiple participants were selling off assets at, often, discounted prices ([Cont & Wagalath \(2013\)](#)).

[Cont & Schaanning \(2017\)](#) calibrate their model to realistic portfolio holdings and market impact parameters and obtain quantitative estimates of the extent of losses due to overlapping portfolio contagion. This provides a good starting point for more sophisticated financial system stress tests that will be discussed in the following sections. The models outlined above can be improved in many ways. [Cont & Wagalath \(2016\)](#) study the effect of overlapping portfolios and fire sales on the correlations of securities in a continuous time setting, where securities prices follow a stochastic process rather than being assumed fixed up to the price impact from fire sales.

2.5.4 Funding liquidity contagion

Funding liquidity contagion has been much less studied than overlapping portfolio or counterparty loss contagion. In the following we will briefly outline some of the considerations that should enter a model of funding liquidity contagion.

In modeling funding liquidity contagion, it is useful to partition an investor's funding into long term funding as well as short term secured and unsecured funding. Only short term funding should be susceptible to funding liquidity contagion as long term funding cannot be withdrawn on the relevant time scales. The availability of secured and unsecured short term funding may be restricted via two channels: a deleveraging channel and a default anticipation channel. The deleveraging channel applies equally to secured and unsecured funding: when a lender needs to deleverage, she can refuse to roll over short term loans, which may in turn force the borrower to deleverage, resulting in a cascade. This channel can be modeled using the same tools applied to overlapping portfolio and counterparty loss contagion. A paper that studies this channel is [Gai et al. \(2011\)](#). The default anticipation channel is different for secured and unsecured funding. In the case of secured funding, a lender might withdraw funding if the quality of the collateral decreases so that the original loan amount is no longer adequately collateralized. In the case of unsecured funding, a lender that questions the credit quality of one of its borrowers might anticipate the withdrawal of funding of other lenders to that borrower and therefore withdraw her funding. This mechanism is similar to a bank run and therefore should be modeled as a coordination game, see [Diamond & Dybvig \(1983\)](#) and [Morris & Shin \(2001\)](#). This poses a challenge for heterogeneous agent models and might explain

¹⁹ The market impact function takes the form $k\sigma\sqrt{\Delta V/V}$, where σ is volatility, ΔV is the size of the trade, V is market trading volume and k is a constant of order one, whose value depends on the market.

the relative scarcity of the literature on this topic. One notable exception that tries to combine both mechanisms is [Anand et al. \(2015\)](#).

2.5.5 Interaction of contagion channels

So far we have focused on counterparty loss and overlapping portfolio contagion in isolation. Of course, focusing on one channel in isolation only provides a partial view of the system and thus ignores important interaction effects. Indeed, it has been shown by a number of authors that the interaction of contagion channels can substantially amplify the effect of each individual channel (e.g. [Poledna et al. \(2015\)](#), [Caccioli et al. \(2015\)](#), [Kok & Montagna \(2013\)](#), [Arinaminpathy et al. \(2012\)](#)). Although constructing models with multiple contagion channels is difficult, some progress has been made.

[Cifuentes et al. \(2005\)](#) and [Caccioli et al. \(2015\)](#) study the interaction of counterparty loss and overlapping portfolio contagion by combining variants of the contagion processes outlined above into a comprehensive simulation model. In particular, using data from the Austrian interbank system, [Caccioli et al. \(2015\)](#) show that the expected size of a default cascade, conditional on a cascade occurring, can increase by orders of magnitude if overlapping portfolio contagion occurs alongside counterparty loss contagion, rather than in isolation.

In an equilibrium model [Brunnermeier & Pedersen \(2009\)](#) show that market liquidity and funding liquidity can be tightly linked. In particular, consider a market in which intermediaries trade a risky asset and use it as collateral for their secured short term funding. A decline in the price of the risky asset can lead to an increase in the haircut applied on the collateral. An increase in the haircut can be interpreted as a decrease in funding liquidity and can force intermediaries to sell some of their assets. This in turn can lead to a decrease in market liquidity of the asset. [Aymanns et al. \(2017\)](#) show that a similar link between market and funding liquidity can also result from the local structure of liquidity in over-the-counter markets (OTC). The authors show that, when the markets for secured debt and the associated collateral are both OTC, the withdrawal of an intermediary from the OTC markets can cause a liquidity contagion through the networks formed by the two OTC markets. Similar to the [Caccioli et al. \(2015\)](#), the authors show that under certain conditions the interaction of two contagion channels – funding and collateral – can drastically amplify the resulting cascade.

Finally, [Kok & Montagna \(2013\)](#) construct a model that attempts to combine counterparty loss, overlapping portfolio and funding liquidity contagion. Such comprehensive stress testing models are the subject of the remainder of this chapter and will be discussed in detail in the following sections.

2.6 From Models to Policy: Stress Tests

2.6.1 What are stress tests?

The insights from the models discussed so far are increasingly used in the tools designed to assess and monitor financial stability. After the crisis of 2008, maintaining financial stability has become a core objective of most central banks.²⁰ One example of such a tool, which has become increasingly prominent over the past years, has been the stress test.²¹ Stress tests assess the resilience of (parts of) the financial system to crises (Siddique & Hasan (2012), Scheule & Roesh (2008), Quagliariello (2009), Moretti et al. (2008)). The central bank designs a hypothetical but plausible adverse scenario, such as a general economic shock (e.g. a negative shock to house prices or GDP) or a financial shock (e.g. a reduction in market liquidity, increased market volatility, or the collapse of a financial institution). Using simulations, the central bank then evaluates how this shock – in the event this scenario were to take place – would affect the resilience of the institution or financial system it tests. Say, for example, that the central bank submits a bank to a stress test. In this case, it would provide the bankers with a hypothetical adverse scenario, and ask them to determine the effect this scenario would have on the bank’s balance sheet. If a bank’s capital drops below a given threshold, it must raise additional capital. Stress tests evaluate resilience to shocks and link that evaluation to a specific policy consequence intended to enhance that resilience (e.g. raising capital). The process also provides valuable information to regulators and market participants, and helps both to better identify and evaluate risks in the financial system.

2.6.2 A brief history of stress test

Stress tests are a relatively novel part of the regulatory toolkit. The potential utility of stress tests had been extensively discussed in the years preceding the financial crisis, and were already used by the International Monetary Fund to evaluate the robustness of countries’ financial systems. Banks already designed and conducted stress tests for internal risk management under the Market Risk Amendment of the Basel I Capital Accord, but it was only during the financial crisis that regulators introduced them on a large scale and took a more proactive role in their design and conduct (Armour et al. (2016)).

²⁰For example, the mission statement of the US Federal Reserve (FED): ‘The Federal Reserve promotes the stability of the financial system and seeks to minimize and contain systemic risks through active monitoring and engagement in the U.S. and abroad’ <https://www.bankofengland.co.uk>.

²¹Timothy Geithner, who played a key role in fighting that crisis as President of the New York Fed and U.S. Secretary of the Treasury, has named his memories after the tool he helped introduce, see Geithner (2014).

In February 2009 the U.S. Treasury Department introduced the Supervisory Capital Assessment Program (SCAP). This effort was led by Timothy Geithner, at a time when uncertainty about the capitalization of banks was still paramount (Schuermann (2014), Geithner (2014)). Under the auspices of this program the Federal Reserve Board introduced a stress test and required the 19 largest banks in the U.S. to apply it. The immediate motivation was to determine how much capital a bank would need to ensure its viability even under adverse scenarios, and relatedly, whether capital injections from the U.S. tax payer were needed. A secondary motivation was to reduce uncertainty about the financial health of these banks to calm markets and restore confidence in U.S. financial markets (Anderson (2016), Tarullo (2016)).

In later years, SCAP was replaced by the Comprehensive Capital Analysis and Review (CCAR) and the Dodd-Frank Act Stress Test (DFAST), which have been run on an annual basis since 2011 and 2013, respectively (FED (2017b,a)). These early stress tests gave investors, regulators and the public at large insight into previously opaque balance sheets of banks. They have been credited with restoring trust in the financial sector and thereby contributing to the return of normalcy in the financial markets (Bernanke (2013)).

Across the Atlantic European authorities followed suit and introduced a stress test of their own (EBA (2017a)). This resulted in the first EU stress tests in 2009, overseen by the Committee of European Banking Supervisors (CEBS) (Acharya et al. (2014)). Due to concerns about their credibility, the CEBS stress test was replaced in 2011 by stress tests conducted by the European Banking Authority (EBA) (see Ong & Pazarbasioglu (2014)). These have been maintained ever since (EBA (2017b)).

In 2014 the Bank of England also introduced stress tests in line with the American example ((Ban (2014)). Around that time, stress tests became a widely used regulatory tool in other countries too (Boss et al. (2007)). Now stress tests are regarded as a cornerstone of the post-crisis regulatory and supervisory regime. Daniel Tarullo, who served on the board of the U.S. Federal Reserve from 2009 to 2017 and was responsible for the implementation of stress tests in the U.S., has hailed stress tests as ‘the single most important advance in prudential regulation since the crisis’ (Tarullo (2014)).

Stress tests are not a uniform tool. They can take a variety of forms, which can be helpfully classified along two dimensions. The first dimension concerns their *object*, or the types of agent that the stress test covers; does the stress test only cover banks, or non-banks as well? In the early days of stress testing, only banks were considered, but now there is an increasing trend towards including non-banks. Given the composition of the financial system in most advanced economies, and the importance of non-banks in these financial systems, it is increasingly acknowledged that excluding non-banks from

stress tests would leave regulators with a partial picture of financial stability risks in their jurisdiction. In the United Kingdom, for example, almost half of the assets in the financial system are held by non-banks (Burrows et al. (2015)), as is illustrated by a stylized map of the UK financial system depicted in Figure 2.3.

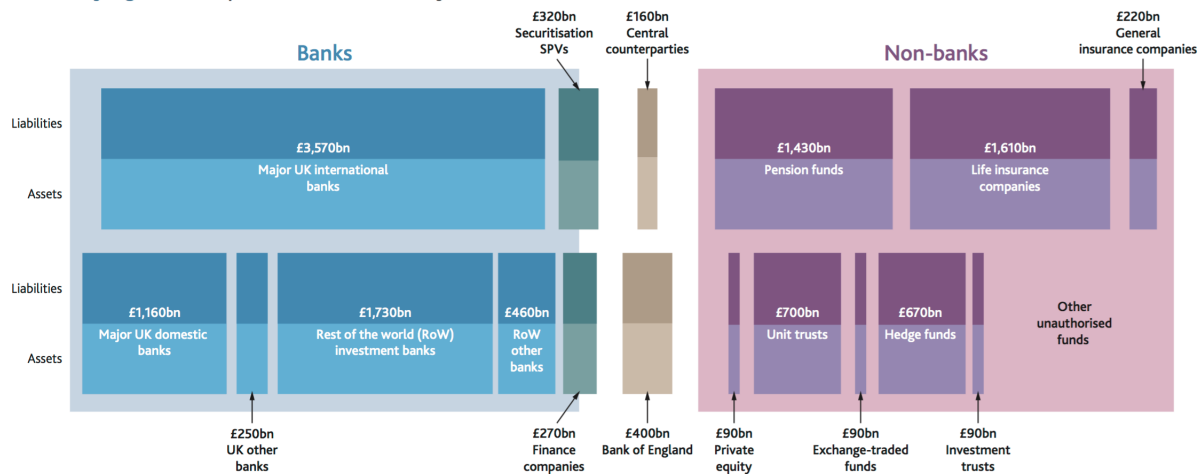


Figure 2.3: Map of the UK financial system. Source: Burrows et al. (2015).

The second dimension concerns the *scope* of the stress test. Generally speaking, stress tests can be used to evaluate the resilience of individual institutions (*microprudential* stress tests), but could also assess the resilience of a larger group of financial institutions or even of the financial system as a whole (*macroprudential* stress tests) (Cetina et al. (2015), Bookstaber, Cetina, Feldberg, Flood & Glasserman (2014)). Methodologically speaking, the key difference is that *macroprudential* stress tests take the feedback loops and interactions between (heterogeneous) financial institutions - as described in section 2.5 of this chapter - into account, whereas the *microprudential* stress tests do not.

Perhaps more than any other financial stability tool, stress tests rely explicitly on the models introduced so far. The following sections will cover micro- and macroprudential stress tests in depth. In each instance, we will first review some representative stress tests and subsequently conclude with an evaluation of their strengths and weaknesses.

2.7 Microprudential Stress Tests

2.7.1 Microprudential stress tests of banks

As noted, microprudential stress tests evaluate the resilience of an individual institution, in this case a bank. Regulators subject the bank to an adverse scenario and evaluate

whether a bank has sufficiently high capital buffers²² (and, in some cases, liquid assets²³) to withstand it.²⁴ If this is not the case, regulators can require the bank to raise additional capital (or liquidity) to enhance its buffers. The idea is that this will make the bank more resilient, and by implication the resilience of the financial system as a whole.

Given this general approach, microprudential stress tests for banks tend to follow three steps. First, the regulator designs the adverse scenario the bank is subjected to. As noted, this scenario usually involves an economic and/or financial shock. In some cases, the scenario consists of multiple (exogenous) shocks operating at the same time, sometimes with specified ripple effects affecting other variables, which together create a ‘crisis narrative’ for the bank. The hypothetical scenario a bank is subjected to must be adverse, plausible and coherent. That is, it cannot consist of a set of shocks that, taken together, violate the relationships among variables historically observed or deemed conceivable. Typically, the exogenous shocks affect a set of macro-variables (such as equity prices, house prices, unemployment rate or GDP) as well as financial variables (such as interest rates and credit spreads).

Second, the effect of this scenario on the bank’s balance sheet is determined²⁵. This determination primarily relates to the effect of the scenario on the bank’s capital (and liquidity) buffer, usually expressed as a ratio of capital (liquidity) buffers to assets,²⁶ and profits. This calculation is based on an evaluation of how the shocks change the values of the assets and liabilities on the bank’s balance sheet, as well as on the bank’s expected income. Value changes on the balance sheet materialize either through a re-evaluation of the market value (if the asset or liability is marked-to-market), or through a credit shock re-evaluation. These effects are captured by market risk models and credit risk

²²The simplest measure of a capital buffer is that of a bank’s net assets – the value of its assets minus its liability. This represents a buffer that protects the bank against bankruptcy when its assets decline in value. In most models described earlier in this chapter, this buffer corresponds to a bank’s equity. When describing whether a bank has a sufficiently high buffer, the term ‘capital adequacy’ is commonly used. For a more comprehensive overview, see [Armour et al. \(2016\)](#), Chapter 14.

²³A liquidity buffer is intended to ensure that, when liquidity risk of the type discussed in section 2.4.2 materialize, a bank has sufficient liquid assets to meet demands for cash withdrawals. Although microprudential liquidity stress tests for banks have been developed, they are currently not yet widely used for regulatory purposes. Hence, we will focus on microprudential capital stress tests here.

²⁴Note that this capital buffer is an example of a regulatory leverage constraint as introduced in section 2.4.2.

²⁵Depending on the regulatory regime, this determination is made either by the regulator or by banks themselves.

²⁶When capital ratios are computed as capital over total (unweighted) assets, this amounts to the inverse of the leverage ratio. Regulators typically use a more complex measure of the capital buffer to account for the fact that some assets are riskier than others. Suppose a bank holds two assets with the same value, but one (asset Y) is riskier than the other (asset X). When regulators take the riskiness of these assets into account, to meet regulatory requirements the bank would have to hold a higher capital buffer for asset Y than for asset X, corresponding to their relative riskiness. This process is referred to as ‘risk-weighting’, and the resulting capital buffer is commonly expressed relative to ‘risk-weighted assets’ (RWA).

models (such as those described in Siddique & Hasan (2012), Scheule & Roesh (2008), Quagliariello (2009), Moretti et al. (2008)). Credit losses for specific assets or asset classes are commonly computed by multiplying the probability of default (PD), the exposure at default (EaD), and the loss given default (LGD). Estimating these variables is therefore key to the credit risk component of stress testing (Foglia (2008)). Value changes to the expected income stream result largely from shocks that affect income on particular assets or asset classes, such as interest rate shocks. This determination matters in the context of the stress test because such income can, in the form of retained earnings, feed-back into capital buffers.²⁷ Usually, these microprudential stress test models therefore equate the post-stress regulatory capital buffer to the sum of post-stress retained earnings plus regulatory capital²⁸ over the post-stress (risk-weighted²⁹) assets.

Third, once the bank's post-stress capital buffer has been determined, regulators compare it to a hurdle rate. This hurdle rate is usually set at such a level that, when passing it, the bank would withstand the hypothetical scenario without being at risk of bankruptcy. Consequently, if the bank does not meet this hurdle rate, it fails the stress test and is said to be 'undercapitalized' (that is, its capital buffer is insufficient). When that happens, the regulator commonly has the authority to require the bank to raise extra capital to increase its buffer, so as to leave it better prepared for adverse scenarios. Microprudential stress tests are thus used as a tool to recapitalize undercapitalized banks, thereby reducing their leverage and increasing their resilience.

Given the importance of non-bank financial institutions to the financial system³⁰, it was only a matter of time before the scope of microprudential stress tests would be extended beyond banks. The rationale for doing so is similar to the one that applies to banks: regulators want to understand the resilience of non-bank financial institutions, and where they find fragility they want to be able to amend it. So far, at least three types of non-bank financial institutions are subjected to stress tests: insurers, pension funds and central clearing parties (CCPs).

²⁷ This is true unless part of this income is being paid to shareholders as dividends, which stress tests commonly assume not to be the case.

²⁸ If the scenario results in a loss to the bank's equity and lowers its income, the capital buffer drops (*ceteris paribus*).

²⁹ In most cases the model also updates the assets' risk-weights to reflect that the adverse scenario has altered the riskiness of the asset (class). For an overview of the methodologies commonly used by banks, see: Capgemini (2014). The 'standard' approach as proposed by regulators is set out in BIS (2015)

³⁰ See e.g. FSB (2015a), ECB (2015), Burrows et al. (2015), Pozsar et al. (2010), Pozsar & Singh (2011), Mehrling et al. (2013), Pozsar (2013)

2.7.2 Strengths and weaknesses of current microprudential stress tests

Microprudential stress tests are valuable from at least three perspectives. First, they give market participants more insight into the opaque balance sheets of the financial institutions being evaluated (Bookstaber, Cetina, Feldberg, Flood & Glasserman (2014)). Opacity coupled with asymmetric information can, especially in times of financial distress, lead to a loss of confidence (Diamond & Dybvig (1983), Brunnermeier (2008)). If the type and quality of a financial institution's assets and liabilities are unclear, outsiders may conceivably fear the worst and, for example, pull back their funding.³¹ Such responses feed speculative runs which can turn into self-fulfilling prophecies and, ultimately, (further) destabilize the financial system at the worst possible time (He & Xiong (2012), Diamond & Dybvig (1983), Martin et al. (2014), Copeland et al. (2014)). Credibly executed microprudential stress tests provide insight into an institution's balance sheet, can signal confidence about the institution's ability to withstand severe stress, and create a separating equilibrium that allows solid banks to avoid runs (Ong & Pazarbasioglu (2014), Bernanke (2013)).³²

Second, microprudential stress tests help financial institutions to improve their own risk-management. By forcing them to assess their resilience to a variety of novel scenarios, stress tests require banks to take a holistic look at their own risk-management practices (Bookstaber, Cetina, Feldberg, Flood & Glasserman (2014)). As a consequence, more banks are now also engaged in serious internal stress tests (Wackerbeck et al. (2016)).

Third, microprudential stress tests have proven to be an effective mechanism to recapitalize banks (Armour et al. (2016)). In the EU, the stress tests have forced banks to raise their capital by 260 billion euros from 2011 to 2016 (Arnold & Jenkins (2016)), and in the US the risk-weighted regulatory ratio of the banks that took part in the stress test went up from 5,6 percent at the end of 2008 to 11,3 at the end of 2012 (Bernanke (2013)). Against a backdrop of frequent questions about the adequacy of banks' capital buffers³³, in part due to the gaming of risk weights (Behn et al. (2016), Fender et al. (2015), Groendahl (2015)), many regulators have welcomed the role that stress tests have played to enhance the resilience of banks. Even if microprudential stress tests are not, strictly speaking, designed to assess and evaluate systemic risk, their role in raising capital adequacy standards can have the effect of enhancing resilience (Greenwood et al. (2015)).

³¹The general economic principle at play is that of asymmetric information causing market failures, see: Akerlof (1970)

³²Weaker banks, however, may be exposed by the stress test. But regulators would learn this information first, giving them an opportunity to intervene before the information reaches the market.

³³See, for example, Admati & Hellwig (2014).

Despite their strengths in specific areas, the current microprudential stress tests have been criticized on at least four grounds. First, and most importantly from the perspective of this chapter, microprudential stress tests ignore the fact that economies are complex systems (as noted in section 2.2) and therefore are ill-suited to capture systemic risk. As discussed in section 2.5 of this chapter, systemic risk materializes due to interconnections between heterogeneous agents (for example due to overlapping portfolios and funding liquidity contagion). By considering institutions in isolation, microprudential stress tests (largely³⁴) ignore the interconnections and interaction between financial institutions that serve to propagate and amplify distress caused by the initial shock resulting from the adverse scenario. Empirical research suggests that this approach substantially underestimates the losses from adverse scenarios (Bookstaber, Paddrik & Tivnan (2014), also see section 2.4). Bernanke (2015), for example, notes that the majority of the losses in the last financial crisis can be traced back to such interactions as opposed to the initial shock emerging from credit losses in subprime mortgage loans.

Second, microprudential stress tests tend to impose an unrealistically large initial shock. Because regulators are aware of the fact that a microprudential modelling strategy does not capture the higher order losses on the balance sheets of individual financial institutions, they use a more severe initial scenario that causes direct losses to compensate for that. To generate a sufficiently large initial shock, the scenario tends to depart quite strongly from reality. Often, the initial scenario posits a substantial increase in the unemployment rate as well as a sharp drop in GDP.³⁵ In reality, however, it is uncommon for these conditions to *precede* a financial crisis, so the stress test might be testing for the wrong type of scenario.³⁶ Imposing an unrealistic shock – and excluding higher-order effects – can also affect the outcome of the stress test in unexpected ways. In particular, while stress tests with large initial shocks might get the overall losses right, they might fail to accurately capture the distribution of losses across institutions, which ultimately determines which banks survive and which do not. For an investigation of this issue, see for example Cont & Schaanning (2017).

Third, the value of the information produced by microprudential stress tests is increasingly being questioned. The outcomes of stress tests have converged (Glasserman et al. (2015)), perhaps because banks seem increasingly able to ‘train to the test’. This has left some to wonder what the information produced by the stress tests is actually worth (Hirtle et al. (2016)), and others to conclude that the value of such information has

³⁴In some cases a proxy for such contagious effects is included in the microprudential stress test, but this is rare.

³⁵See, for example, FED (2016), BoE (2016), ESRB (2016).

³⁶Instead, exogenous shocks such as declining house prices or stock markets precede financial crises. These are commonly also part of the initial scenario.

declined over time (Candelon & Sy (2015)). Such concerns have been further fuelled by the apparent willingness of some regulators to allow banks to pass the test on the basis of dubious assumptions.³⁷

Finally, the stress tests are commonly calibrated to the losses incurred during the last financial crisis, raising questions about their relevance in relation to current, let alone future, scenarios – not least because the financial system constantly changes.

2.8 Macroprudential Stress Tests

Because the financial system is a complex system (see section 2.2), the whole is different from the sum of its parts (Anderson et al. (1972), Farmer (2012), Battiston et al. (2016)). In other words, measures focused on the health of individual institutions (as microprudential stress tests would prescribe) will not necessarily guarantee the health of the financial system as a whole. In fact, such measures might destabilize the system. To understand the system as a whole - and, by implication, systemic risk - stress tests have to account for feedback loops and non-linearities.

The inability of microprudential stress tests to appropriately account for systemic risk has prompted the development of a specific type of stress tests focused on this goal; the macroprudential stress test. Macroprudential stress tests aim to assess the resilience of a whole sector, or even the whole financial system, rather than that of one particular institution. To do so, they extend the microprudential stress test by including contagion effects between interconnected financial institutions that can arise following the initial adverse scenario. This means that the regulators must not only assess the effect of the initial shocks on the individual balance sheets, but must capture how the balance sheets are interlinked (see section 2.5). They should also address what consequences such interlinkages have for the potential of financial distress to propagate throughout the system. The contagion models discussed in section 2.5 can help inform regulators on how to model these higher order spill-over effects.

In this section, we focus on discussing two macroprudential models for banks, and examining one that combines banks and non-banks. We also briefly touch on recently developed system-wide stress testing models. The first two models, the Bank of England's 'Risk Assessment Model of Systemic Institutions' (RAMSI) and the Bank of Canada's 'MacroFinancial Risk Assessment Framework' (MFRAF), have been used in stress tests.

³⁷Deutsche bank, which has seen its share price fall significantly in 2016 on fears that it could face a US fine of up to USD 14bn, was given special treatment by the European Central Bank in the 2016 EBA stress tests, so that it could use the result of the stress test as evidence of its healthy finances (Noonan et al. (2016)).

The last model, U.S. Office of Financial Research’s (OFR’s) ‘Agent-Based Model for Financial Vulnerabilities’ (ABMFV), has not.

The ABMFV and the RAMSI are examples of cases where heterogeneous agent models have been applied to macroprudential stress tests.³⁸ The MFRAF is an example of another (neoclassical) approach.

After introducing these three models, their differences and similarities are outlined. The section ends with a discussion of the strengths and weaknesses of these macroprudential stress tests.

2.8.1 Examples of macroprudential stress tests

RAMSI stress test of the Bank of England

The Bank of England has pioneered the development and use of a macroprudential banking stress test, called the RAMSI model. Though the model has been phased-out, we discuss this model to showcase the strengths and weaknesses of one of the earliest systemic stress test models that leverages models of contagion. The model evaluates how adverse shocks transmit through the balance sheets of banks and can cause further contagion effects (Burrows et al. (2012)). It is based on earlier research that has been conducted by Bank of England researchers and others (Aikman et al. (2009), Kapadia et al. (2013), Alessandri et al. (2009)).

The RAMSI stress test begins as a microprudential stress test. Subsequently, possible feedback effects within the banking system are considered. If the initial shocks have caused a bank to fall below its regulatory capital ratio, or have caused the bank to be shut off of all unsecured funding³⁹ markets, the bank respectively suffers an insolvency or illiquidity default. Subsequently, the default causes two interbank contagion effects: common asset holding contagion and interbank contagion. The combined effect of the marked-to-market losses and the credit losses can cause other banks to default through insolvency or illiquidity by being shut out of the funding market. If this happens, the loop is repeated. If this does not happen, each bank’s net operating expenses are invested in assets such that the bank targets its regulatory risk-weighted target ratio. The credit losses persist, but the marked-to-market losses are assumed to disappear as each asset

³⁸Indeed, these models combine the contagion models discussed in section 2.5

³⁹This causes funding liquidity contagion. The bank is shut off of all unsecured funding based on a rating system. Based on the shocked balance sheets and profit and losses (PL), the credit score for the bank is computed, which the authors assume affects the funding cost of the bank and its ability to access the long-term and short-term funding market. This credit score takes into account liquidity and solvency characteristics of the bank’s balance sheet, but also system-wide market distress. If its credit score is above a certain threshold, the bank is shut out of the unsecured funding markets altogether (both long-term and short-term) and is assumed to default.

price returns to its fundamental value. Then, the next time step starts, and the process can be repeated, starting with a balance sheet that includes the credit losses incurred in the previous time step.

Thus, the RAMSI stress test turns a microprudential foundation into a macroprudential model by including interbank contagion effects via common asset holdings, interbank losses and funding liquidity contagion. Figure 2.4 summarizes what happens at each step of the RAMSI model.

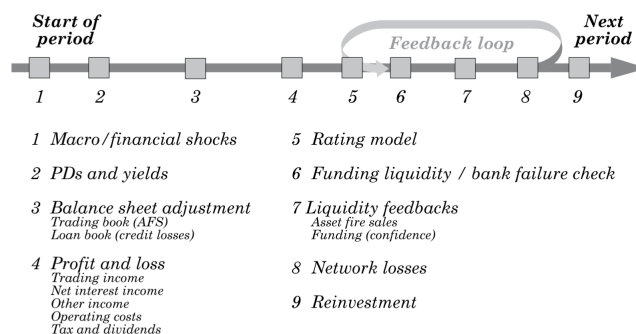


Figure 2.4: Description of the RAMSI stress test of the Bank of England. Source: [Aikman et al. \(2009\)](#)

MFRAF stress test of the Bank of Canada

Contrary to the RAMSI model, the Bank of Canada’s MacroFinancial Risk Assessment Framework (MFRAF) is at its core not a heterogeneous agent model, but a global games model, such as those described in [Morris & Shin \(2001\)](#). In the way it sets up funding runs (i.e. as a global coordination game) it is similar to the seminal model of [Diamond & Dybvig \(1983\)](#) (discussed in section 2.5). It captures three sources of risk that banks face ([Anand et al. \(2014\)](#), [?](#), [BoC \(2012\)](#)): solvency, liquidity, and spill-over risk (see Figure 2.5).




Figure 2.5: Description of the MFRAF stress test of the Canada. Source: [Anand et al. \(2014\)](#).

The MFRAF stress test has been applied to the Financial Sector Stability Assessment (FSAP) of the Canadian financial sector conducted by the International Monetary Fund (IMF) in 2014 (IMF (2014)). The 2014 FSAP stress test, which considers the direct effects of adverse shocks on the solvency of banks, is microprudential. When extending it to capture system-wide effects (i.e. liquidity effects and spill-over effects) using MFRAF, overall losses to the capital of the Canadian banks rose with 20 percent. This again shows that microprudential stress tests significantly underestimate system-wide losses. We will now discuss the theoretical underpinnings of the MFRAF stress tests, which builds on research at the Bank of Canada and elsewhere (Anand, Gauthier & Souissi 2015, Gauthier, Lehar & Souissi 2012, Gauthier, Souissi, Liu et al. 2014).

The theoretical model that underpins the MFRAF stress test is described in Anand et al. (2015) and will be discussed here.⁴⁰ The model captures how solvency risks, funding liquidity risks, and market risks of banks are intertwined. In essence, this works as follows: a coordination failure between a bank’s creditors and adverse selection in the secondary market for the bank’s assets interact, leading to a vicious cycle that can drive otherwise solvent banks to illiquidity. Investors’ pessimism over the quality of a bank’s assets reduces the bank’s access to liquidity, which exacerbates the incidence of runs by creditors. This, in turn, makes investors more pessimistic, driving down other banks’ access to liquidity. The model does not capture interbank contagion upon default, although this is captured in MFRAF (IMF (2014)).

The key components of the model according to the evolution of the model over time is summarized in Figure 2.6.

$t = 0$	$t = 1$ (round 1)	$t = 1$ (round 2)	$t = 2$
1. Debt issuance	1. Interim shock	1. Belief updated	1. Investment matures
2. Investments	2. Private signals	2. New pooling price	2. Final shock
	3. Debt withdrawals	3. New private signals	3. Debts honored
		4. Debt withdrawals	

Figure 2.6: Time steps in theoretical model of MFRAF stress test of the Bank of Canada. Source: Anand et al. (2015).

ABM for Financial Vulnerabilities

⁴⁰The degree to which the theoretical model of Anand et al. (2015) is in unaltered form translated into the MFRAF stress tests is not made explicit in the IMF (2014) documentation of the MFRAF stress test.

The final system-wide stress testing model that will be discussed, the Agent-Based Model (ABM) for Financial Vulnerabilities (Bookstaber, Paddrik & Tivnan (2014))^{41,42}, captures similar contagion mechanisms as MFRAF, but it does so using a different methodology. The model is designed to investigate the vulnerability of the financial system to asset- and funding-based firesales that can lead to common asset holding contagion.

The financial system is modelled as a combination of banks that act as intermediaries between the cash provider (a representative agent for various types of funds) and the ultimate investors (i.e. the hedge funds). Hedge funds can receive funding from banks for long positions in return for collateral. Banks, in turn, receive funding from the cash provider in return for collateral. Funding and collateral therefore flow in opposite directions, as is illustrated in Figure 2.7.

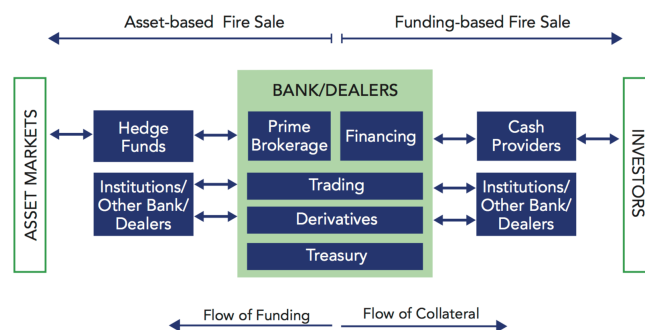


Figure 2.7: Map of the financial system and its flows, as considered in the ABM for Financial Vulnerabilities. Source: OFR (2014).

The role of the cash provider c in the model is to provide secured funding to banks.⁴³ Although the cash provider is not actively modelled, it can take two actions. First, it can set the haircut (this can force the hedge fund to engage in fire sales), and second it can pull funding from the banks (this may lead the bank to contribute to pre-default contagion or default).

Hedge funds have a balance sheet that consists of cash and tradable assets on the asset side, and secured loans and equity (and possibly short positions) on the liability side. A hedge fund funds its long positions in assets using funding from banks in the form

⁴¹The model was developed by researchers who at the time worked at the U.S. Office of Financial Research.

⁴²A further discussion of some agent-based models of the financial crisis and stress testing can be found in Bookstaber & Kirman (2018).

⁴³The cash provider is a representative agent that represents financial institutions that typically provide funding to banks, such as asset managers, pension funds, insurance companies, and security lenders, but most importantly, money market funds.

of repurchase contracts (often referred to as repos).⁴⁴ When funding themselves this way, hedge funds receive cash in return for collateral they pledge to the bank. Although the hedge fund does not face a regulatory leverage constraint, it faces an implicit leverage constraint based on the haircut it receives on its collateral. The haircut determines how much equity a hedge fund needs for a given amount of repo funding. If the haircuts on all types of collateral (i.e. on all types of assets that can be pledged as collateral) is the same, and assuming that the bank passes on the haircut it receives from the cash provider, the maximum leverage $\bar{\lambda}_{jt}$ of the hedge fund j at time t is given by $\bar{\lambda}_{jt} = \frac{1}{h_{cjt}}$. If the leverage of the hedge fund exceeds the maximum leverage⁴⁵, the hedge fund is forced to de-lever. It will do so by fire selling assets. This can cause marked-to-market losses for other banks or hedge funds who hold the same assets.

The banks act as an intermediary between buyers and sellers of securities and between lenders and borrowers of funding.⁴⁶ On the whole, the bank can contribute to financial distress pre-default and post-default in various ways. Pre-default, the bank may have to fire sell assets or to pull funding from the hedge fund (which consequently may also have to engage in firesales) in order to raise cash, de-lever, or pay back funding to the cash provider (if the cash provider pulled its funding). In addition, by passing on an increased haircut to the hedge fund, it can trigger a hedge fund to engage in firesales. Post-default, the bank contributes to exposure losses and further firesale losses.

Recent developments in macroprudential stress testing

Though, we have focussed on comparing three models: RAMSI, ABMFV and MFRAF, other relevant macroprudential stress tests exist. Here, we will discuss relevant systemic

⁴⁴In a repo, one party sells an asset to another party at one price at the start of the transaction and commits to repurchase the fungible assets from the second party at a different price at a future date. If the seller defaults during the life of the repo, the buyer (as the new owner) can sell the asset to a third party to offset his losses. The asset therefore acts as collateral and mitigates the credit risk that the buyer has on the seller. Although assets are sold outright at the start of a repo, the commitment of the seller to buy back the fungible assets in the future means that the buyer has only temporary use of those assets, while the seller has only temporary use of the cash proceeds of the sale. Thus, although repo is structured legally as a sale and repurchase of securities, it behaves economically like a collateralized loan or secured deposit. For an overview, see: <https://www.icmagroup.org/Regulatory-Policy-and-Market-Practice/repo-and-collateral-markets/icma-ercc-publications/frequently-asked-questions-on-repo/1-what-is-a-repo/>.

⁴⁵A hedge fund's leverage can exceed the maximum due to asset prices depreciations (as a consequence of firesales, for example) or increases in the haircut (due to the cash provider's downward assessment of the bank's solvency and/or liquidity). If the hedge fund is forced to de-lever, it will attempt to go back to a 'buffer leverage' level, which is below the maximum leverage value.

⁴⁶ In its role, it facilitates maturity, liquidity, and risk transformations. The banks have various desks that play a role in these processes: the prime broker, the finance desk, the trading desk, the derivatives desk, and the treasury. The various equations associated with the functioning of the bank dealer and its various subdesks can be found in [Bookstaber, Paddrik & Tivnan \(2014\)](#).

stress test models that have recently been developed.

[Baranova et al. \(2017\)](#), from the BoE, study market liquidity in the corporate bond market by modelling broker-dealers, hedge funds and (unlevered) fund managers. This stress test focusses on simulating common asset holding contagion; it does not model any of the other contagion mechanisms that have been identified in the financial stability literature. [Baranova et al. \(2017\)](#) show that unlevered non-banks, such as mutual funds, may be the a source of distress, by acting procyclically in response to distress. Importantly, they demonstrate that continued financial resilience, in such case, depends on the ability of intermediaries, such as market makers, to buy, when some institutions, specifically funds, are forced to sell. They also show the role that hedge funds can play in buttressing financial stability, by buying under-priced assets, which puts a floor under asset prices and offers market liquidity by acting as a buyer when others are (forced) sellers. Overall, [Baranova et al. \(2017\)](#) improve upon earlier stress tests, such as RAMSI, ABMFV and MFRAF, by modelling and highlighting how the different types of financial institutions may contribute to or quench distress depending on their incentives, resources, and binding constraints.

Rather than simulating stress dynamics in a network model, as [Baranova et al. \(2017\)](#) do, [Aikman et al. \(2019\)](#), also from the BoE, put forward a general equilibrium model to conduct system-wide stress tests of the market-based-finance sector. At the centre of this model is a set of representative agents, which represent the core financial sectors: dealers, open-ended investment funds, hedge funds, and long-term investors, such as insurance companies and pension funds. By focussing on representative sectors, their model is silent about within-sector heterogeneity, which is an important omission. Their model incorporates two key contagion channels, and their interaction, that are emphasised in the literature: overlapping portfolio contagion and margin-call contagion. In their model, the interaction between agents' solvency and liquidity constraints can prompt fire sales by one agent that can create falls in market prices, reducing the asset value of other agents, which in turn may prompt them to sell. They also focus on the repo market, in which they consider that cash-providers could either pull short-term funding or increase haircuts, forcing those reliant on repo funding to deleverage, including by fire selling assets or reducing their own provision of repo. A similar type of contagion in the model can also take place via derivatives markets via margin calls. This general equilibrium stress test explores how tipping points from stability to instability may depend on the size of the adverse shock and distance to institutions' constraints, among others. In many ways, this model is the equilibrium version of the ABMFV, likewise focusing on the interaction

between overlapping portfolio and margin-call contagion. Unlike the ABMFV, they include open-ended investment funds.

Dees & Henry (2017), from the ECB, have developed STAMP€: Stress-Test Analytics for Macroprudential Purposes in the euro area. Dees & Henry (2017) offer a host of modular, macroprudential stress testing tools. Importantly, STAMP€ may be used as a macroprudential extension to a microprudential stress test, as in the case of MFRAF. Two financial stability models are incorporated in STAMP€’s stress test toolkit: overlapping portfolio contagion and interbank contagion. STAMP€ also integrates liquidity and solvency stress, and it allows for dynamic balance sheet adjustments of institutions in response to an adverse scenario. Other than all the other stress tests we have discussed in this section, it encompasses – in a stylised form – feedbacks between the financial system and the real economy.

2.8.2 Comparing and evaluating macroprudential stress tests: five building blocks

To comprehensively design, study and evaluate macroprudential stress tests, we introduce a general framework consisting of five building blocks that allow us to break down each stress test in discrete components: (1) types of financial institutions (agents), (2) financial contracts, (3) markets, (4) constraints, and (5) behavior.⁴⁷ This framework also offers an analytically coherent way to combine the various heterogeneous agent models discussed in section 2.5 in order to capture their interactions (see section 2.5.5). With such a framework one can capture *critical features*⁴⁸ necessary to be able to capture systemic risk. This section covers these five building blocks and compares the three macroprudential stress tests discussed above⁴⁹ as we go along (these findings are summarized in table 2.1). We will see that these stress tests implement each building block with varying degrees of fidelity to the real world.

⁴⁷With these five building blocks, many relevant features of a financial system can be captured by initialising bespoke implementations for each building block. Once financial institutions and financial contracts are defined, a multi-layered network can be initialized. When, subsequently, markets, constraints and behavior are chosen, the dynamics of system can be studied. For a more elaborate description of how these building blocks can be used to develop a generic nesting model for system-wide stress tests, see: Farmer et al. (2020).

⁴⁸E.g. all relevant contagion channels and sectors.

⁴⁹See section 2.8.1, section 2.8.1, and section 2.8.1.

Table 2.1: Comparison between the three macroprudential stress tests (RAMSI, MFRAF, ABMFV) regarding the (system-wide stress test) building blocks: (1) financial institutions; (2) financial contracts; (3) markets; (4) constraints; and (5) behavior. Note that rc, cc, mc stand for regulatory, contractual and market-based constraints respectively. Remark that MFRAF captures unsecured interbank loans, counterparty loss contagion and a leverage constraint, the theoretical model of [Anand et al. \(2015\)](#) does not. We list the behavior that impacts the state of the system.

	RAMSI	MFRAF	ABMFV
<hr/> (1) Financial institutions: <hr/>			
Banks	✓	✓	✓
Creditors (exogeneous)	✓	✓	✓
Hedge funds	-	-	✓
<hr/> (2a) Financial contracts: <hr/>			
Traded securities	✓	✓	✓
Unsecured interbank loans	✓	✓	-
Unsecured term deposits	-	✓	-
Secured interbank loans (repos)	-	-	✓
<hr/> (2b) Channels of contagion: <hr/>			
Overlapping portfolios	✓	✓	✓
Counterparty loss	✓	✓	-
Funding liquidity	✓	✓	✓
Margin spirals	-	-	✓
<hr/> (3) Modeled markets: <hr/>			
Traded securities	✓	✓	✓
<hr/> (4) Constraints: <hr/>			
Leverage constraints (rc)	✓	✓	✓
Liability payment obligations (cc)	✓	✓	✓
Margin call obligations (cc)	-	-	✓
Funding run (mc)	✓	✓	✓
<hr/> (5) Behavior: <hr/>			
Pre-default			
- no action (banks)	✓	✓	-
- action (banks, hedge funds)	-	-	<ul style="list-style-type: none"> • control leverage • meet contractual obligations • maximize profits
- exogenous action (creditor run)	✓	✓	✓
Post-default			
- Default procedure	<ul style="list-style-type: none"> • fire sales • exposure losses 	<ul style="list-style-type: none"> • fire sales (of collateral) 	<ul style="list-style-type: none"> • fire sales (implicit)

Financial institutions

Financial institutions are at the heart of any financial stability analysis and form a key component of macroprudential stress tests. In most models they are represented by balance sheets filled out with a collection of financial contracts that are unique to that institution. Moreover, each institution comes with its own set of constraints and behavioral rules. By endowing an institution with its unique collection of financial contracts,

combination of constraints, and behavioral rules, various types of heterogeneous financial institutions (e.g. banks, insurance companies, hedge funds, unlevered funds, central clearing parties) can be characterized. This allows for the inclusion of the many types of financial institutions that need to be studied to capture the dynamics of a financial system under stress.

None of the macroprudential models discussed in this chapter capture all relevant financial institutions, which limits their claim to be a truly ‘system-wide’ macroprudential stress test. Specifically, the RAMSI and MFRAF model only capture the banking system, and though the ABMFV also considers non-banks it only covers a subset (hedge funds and cash providers).⁵⁰

Financial contracts: interlinkages and associated contagion channels

Contracts sit on the balance sheet of each institution, but because contracts are between institutions, they also stipulate the interconnections between institutions. Taking institutions as the nodes in the network the contracts define the edges of the network. (Common asset holdings also define connections, though a more accurate approach is to treat these as bipartite networks). Contagion dynamics, such as those described in section 2.5, operate over these financial contracts to jump from institution to institution. It is therefore important to ensure that the models representing these contracts capture the features that create the interconnections between institutions (e.g. contractual counterparties) and enable contagion (e.g. valuation method, contractual obligations).

The three macroprudential stress tests capture these three contractual characteristics for a subset of contracts (leaving out some relevant contractual types), but do study how the contagion dynamics operating over them can interact. Specifically, models capture the interaction between contagion channels discussed in section 2.5: common asset holding contagion, counterparty loss contagion, and funding liquidity contagion. The ABMFV also captures ‘collateral contagion’.⁵¹

⁵⁰Each model also considers exogenous creditors. The balance sheets of exogenous agents is not explicitly modelled. As such exogenous agents cannot default. When exogenous creditors withdraw a loan, the cash exists the system.

⁵¹‘Collateral contagion’ refers to the contagious spill-overs that can arise from margin calls associated to repo contracts (e.g. secured funding contracts). Institutions receive a margin call when the asset collateral value drops (or haircuts increase) so that it is not enough to cover the loan amount. If institutions are not able to meet the margin call they may be forced to engage in fire sales. Collateral contagion is especially relevant as it interacts with common asset holding contagion. Indeed, price falls due to fire sales can trigger collateral contagion.

Markets

In most models (as in reality), markets are the places where asset prices are determined, as well as the place where new contracts are agreed upon and existing ones modified or terminated. It is their role in the price formation process and the provision of liquidity that makes the modelling of markets particularly relevant to macroprudential stress tests. Markets are diverse in their institutional characteristics; they can be bilateral (such as the interbank loan market), exchange-based (like the stock market), intermediated (like a dealer-based market for, say, corporate bonds), or centrally cleared (i.e. by a CCP). Typically, there is a specific market for each financial contract on the balance sheet of an institution.

However, although each of our three macroprudential models consider multiple types of contractual linkages, they only model one market: the market for common asset holdings.⁵² Moreover, although all three models consider a bilateral funding market for (un)secured funding, they do not consider a bilateral funding market for these contracts. Therefore, when an (un)secured loan is not rolled over, institutions have no opportunity to seek funding elsewhere. That potentially causes these models to overestimate financial distress.

Because financial stability critically depends on price formation and the ability of institutions to forge contractual links (or break them), it is important to model the markets that exist for each type of contract (and do so with sufficient realism).⁵³ An understudied challenge is thus to determine whether and how the dynamics in a given market contribute to financial (in)stability, and to reflect that in stress testing models. This is complicated, because ideally it would require an understanding of the supply and demand functions for each market.⁵⁴

Constraints

Financial institutions typically face four types of constraints: regulatory constraints, contractual constraints, market-based constraints, and internal risk limits. Regulatory

⁵²The modelling of price formation is approached differently in the three models. In the case of the RAMSI and MFRAF model a price impact function is used. The MFRAF model updates prices based on the investors' beliefs about the quality of the assets.

⁵³E.g. [Baranova et al. \(2017\)](#) show for the case of corporate bond markets that market liquidity (and common asset holding contagion) critically depends on the ability of intermediaries to make markets.

⁵⁴To capture price formation (or counterparties for a bilateral contract), the model must produce well-balanced supply and demand as observed in normal times and allow for imbalances in times of distress.

constraints are constraints set by the regulator. Most regulatory constraints are specific to a type of institution; banks face different regulatory constraints than insurers, for example. The models capture a subset of the regulatory constraints that banks face⁵⁵ and do not capture the regulatory constraints that non-banks confront. As has been discussed in that section, insurers face a Solvency II constraint, pension funds face a coverage ratio, and CCPs must fulfil default fund requirements..

Contractual constraints arise out of contractual obligations. Because, as noted before, each financial institution holds a unique collection of contracts, the contractual constraints of each institution are unique too. Each model covers repayment obligations, because they capture (un)secured funding contracts. The ABMFV also considers margin call obligations as part of the secured funding contracts. Because each of the macroprudential stress tests discussed above only captures a subset of the relevant contracts, the contractual constraints they capture are incomplete as well. Banks, for example, typically hold derivatives contracts (e.g. credit default swaps) that can give liquidity shocks that may foster pre- or post-default contagion.

Market-based constraints (commonly referred to as ‘market discipline’) are those that are enforced by market participants. Sometimes, market participants set higher standards than regulators do; a bank might, for example, be cut off from funding markets because its leverage is judged to be too high, even though it still meets the regulatory leverage requirements. In this case, the market constraint could be formalized as a leverage constraint that is stricter than the regulatory leverage constraint. The most relevant market-based constraint, which entails that creditors run if the liquidity and/or solvency characteristics of a bank are sufficiently negative, is captured by all three models.⁵⁷

Finally, internal risk limits are set by the financial institutions themselves, as part of their risk-management practices. An example could be a value-at-risk (VaR) constraint on a portfolio.⁵⁸

Taken together, these constraints (and their various interactions) can drive an institution’s behavior, especially under stress. First, institutions may act in a precautionary manner to avoid breaching constraints in order to avoid defaults. These actions, which

⁵⁵The RAMSI model and the ABMFV consider one regulatory constraint for banks, an (unweighted) leverage ratio and a risk-weighted leverage ratio respectively. The theoretical model of [Anand et al. \(2015\)](#) that underpins the MFRAF stress test does not consider a regulatory leverage constraint.⁵⁶ Banks default when they no longer meet their minimum (risk-weighted) leverage constraint. The models do not capture other regulatory constraints of banks that may affect financial stability, such as liquidity constraints (e.g. the liquidity coverage ratio) for banks.

⁵⁷The models consider the creditors to be exogenous to the system. A more realistic approach would be to make these creditors endogenous to the system. That way, cash does not leave the system but ends up in an institution’s pockets.

⁵⁸None of the macroprudential models discussed here consider internal risk limits.

are often prudent for each institution separately, may contribute to pre-default contagion (e.g. firesales in order to meet payment obligations). Second, institutions may fail to avoid breaching a constraint and default, which then leads to post-default contagion (e.g. due to exposure losses). Given their vital role in driving interactions under stressed conditions, it is important to consider whether the constraints included in a given stress test model represent those most relevant to the description of the system or sector that is being studied. More specifically, for any given institution the nature of its contribution to contagion will be critically determined by the set of constraints it faces. In sum, a failure to consider the relevant constraints makes it unlikely that the stress test model will correctly identify which channels of contagion operate and which institution are affected (Cetina et al. (2015)).

Behavior

Behavior drives the dynamics of the financial system and the evolution of the multi-layered network representation thereof. It therefore critically affects the inherent stability of the financial system and can be an important driver of contagion. behavior of institutions is typically not known and must thus be reasonably estimated.

Institutions can *affect the state of the system* when they default (*i.e. post-default*) or when they are still alive (*i.e. pre-default*). When institutions are alive they act for two reasons: *to fulfil objectives* (e.g. seek profits) and *to avoid default*.⁵⁹ When institutions default either through insolvency (*i.e. breaching regulatory constraints*) or illiquidity (*i.e. when an institution does not meet its contractual obligations*) they also affect the system. Through these pre- and post-default actions institutions can contribute to contagion.⁶⁰

The three macroprudential stress testing models capture the critical drivers of financial stability dynamics to various degrees. The ABMFV most realistically simulates a financial market and its (contagious) dynamics. It captures that institutions can contribute to ‘pre-default contagion’ when they aim to avoid default⁶¹, but can also contribute to ‘post-default contagion’ once they have defaulted⁶². In addition, the ABMFV captures

⁵⁹Note that many financial stability models (see section 2.5) abstract away from profit-seeking behavior. This may be a reasonable abstraction because in times of distress behavior is typically mostly driven by the wish to avoid default. However, by doing so, these models might overestimate contagion. As in crises times, the institutions who are not under pressure (e.g. do not experience binding constraints) can stabilize the market.

⁶⁰Or act as a stabilizer.

⁶¹E.g. institutions who must meet the contractual obligation to repay a loan, may engage in fire sales to do so.

⁶²To capture the contagion consequences that may ensue following a default, the relevant aspects of a default procedure must be modelled. For example, models must not only capture one contagion effect (e.g. exposure loss contagion), but all relevant contagion effects (e.g. including common asset holding contagion, etc.).

normal-time behavior, presumably to ensure that contagion is not overestimated (e.g. some may be willing to buy when others are forced to sell). The MFRAF and the [Aikman et al. \(2009\)](#), [Alessandri et al. \(2009\)](#) versions of the RAMSI model⁶³ assume that institutions are largely *passive*: they do not act until they default (only when they do default, institutions affect the system). Barring any defaults, these models thus only capture dynamics to a limited extent. By not capturing pre-default contagion, these may significantly underestimate losses (see e.g. [Bardoscia et al. \(2017\)](#)). Table 2.1 summarizes the implementation of behavior in the three macroprudential stress testing models.

2.8.3 Strengths and weaknesses of the current macroprudential stress tests

Macroprudential stress tests are strongly complementary to microprudential stress tests, because they allow regulators to assess the resilience of the financial system as a whole (or a larger subset of it) rather than that of individual financial institutions. The current macroprudential stress tests have three related strengths.

First, they provide insights into the interlinkages between financial institutions, mapping out how financial shocks transmit through individual balance sheets and affect other institutions. The data-driven methodology to establish the model setup (as well as the subsequent calibration) provide a promising avenue for future stress tests, but also for further data-driven research into the structure of the financial system ([Aikman et al. \(2009\)](#)).

Second, they capture the interactions between various financial institutions and contagion channels that can drive distress, and therefore capture (some of) the feedback effects that characterize the complex nature of the financial system (see section 2.5). Especially the ABM for Financial Vulnerabilities makes an important contribution by including heterogeneous financial institutions, which is key to allow for emergent phenomena ([Bookstaber \(2017\)](#)).

Third, in addition to capturing solvency risk, or separately investigating solvency and liquidity risk, the current macroprudential stress tests capture funding liquidity risk and the interactions between solvency and liquidity (the interaction between contagion channels has been discussed in section 2.5.5). The RAMSI model, for example, not only considers defaults through insolvency, but also through illiquidity, and takes their interaction into account. In case of the MFRAF, a particular strength is that market risk and funding liquidity are endogenously determined. Market risk is based on the degree of adverse selection. Because of asymmetric information, investors offer banks a pooling price for their assets. The pooling price (and hence the market liquidity) lowers

⁶³The [Kapadia et al. \(2013\)](#) version of the RAMSI model does capture pre-default contagion.

if investors become more pessimistic and the quality of the assets is lower. Funding liquidity risk is determined as a function of the bank's credit and market losses (based on general market confidence, and thus as a function of information contagion), its funding composition and maturity profile, and concerns that creditors may have over its future solvency.

Despite these strengths, there is substantial scope for improvement. First, most macroprudential stress tests only cover banks and their creditors, and therefore fail to capture interactions with non-banks that make up a substantial part of the financial system. Non-banks have played an important role in amplifying distress to the banking sector during the 2007-2009 financial crisis (Bernanke 2015). Therefore, failing to capture non-banks does not just exclude many institutions from the analysis, but also leaves regulators less well-equipped to understand the resilience of the subset of financial institutions they do study. The ABM for Financial Vulnerabilities is an exception, since it does include multiple types of financial institutions, but contrary to the RAMSI and the MFRAF models it is not used as a regulatory stress test.

Second, and relatedly, most macroprudential stress tests capture only a few types of interconnections, even though it is clear that the multiplicity of channels and interconnections between financial institutions plays a critical role in spreading distress (Brunnermeier 2008) (see also section 2.5.5). Notable examples of such contractual linkages include securitized products and credit default swaps.

Third, most current macroprudential stress tests only capture post-default contagion. However, in financial crises pre-default contagion is rampant, often resulting from actions that are prudent from a firm-specific risk-management perspective, but destabilizing from a system-wide perspective. A bank, for example, might engage in precautionary deleveraging to avoid insolvency (i.e. breaking a leverage constraint), which can add to further negative price spirals. Not capturing such dynamics implies that the total size of contagion, as well as the timing of contagion, is misunderstood.

These three areas of improvement essentially come down to the same point: the current macroprudential stress tests insufficiently capture the diversity of agents and interactions that make up the financial system, and therefore do not do justice to the complex nature of the financial system (or, for that matter, to the insights of the heterogeneous agent model literature, see section 2.5). One of the important challenges is to devise a modelling strategy that can capture these various effects, and the ABM for Financial Vulnerabilities offers a promising start; the model could easily be extended to capture more types of financial institutions (e.g. central clearing parties, pension funds), financial contracts (e.g. derivative contracts, securitized products), and constraints that drive behavior under stressed circumstances (Cetina et al. (2015), Farmer et al. (2020)).

Finally, macroprudential stress tests must be more data-driven⁶⁴ and more carefully calibrated to be credible. So suitably designed system-wide stress tests are enabled to become more credible as regulators collect better (contract-level) data.

2.9 Conclusion

Computational agent-based models provide a useful complement to more traditional equilibrium based methods. They have already been shown to be essential for understanding the dynamics of systemic risk and for investigating the network properties of the financial system. Their role is likely to become even more important in the future, as increasingly comprehensive fine-grained data becomes available, making it possible to carefully calibrate such models so that they can yield more quantitative conclusions. Due to the inherent complexity of the financial system, and in particular its nonlinear feedback loops, analytic methods are unlikely to be sufficient.

We expect that computational and simulation methods will soon begin to go beyond hard wired behavioral rules and move increasingly toward myopic optimization. Models of boundedly rational heterogeneous agents, who learn and adapt their behavior in response to observed market realizations and newly adopted policies, withstand the Lucas critique. Behavioral economists have documented more and more situations in which people are not fully rational, emphasizing the obvious point that realistic behavior lies somewhere between full rationality and zero intelligence. Computational models offer the possibility of implementing realistic levels of strategic behavior, while allowing one to model the complex institutional structure of the financial system. We think that computational models will play an expanding role for understanding financial stability and systemic risk.

⁶⁴This depends on data availability.

Chapter 3

Foundations of System-Wide Stress Testing

3.1 Summary

We propose a framework for the development of system-wide financial stress tests with multiple interacting contagion, amplification channels and heterogeneous financial institutions. This framework conceptualises financial systems through the lens of five building blocks - financial institutions, contracts, markets, constraints, and behaviour. These blocks can be flexibly implemented to form a dynamic multiplex network using the accompanying simulation engine and software library (the ‘Economic Simulation Library’, or ‘ESL’).

Using this *framework*, we implement a system-wide stress test *model* for the European financial system that incorporates amplification risks associated with default contagion, price-mediated contagion via asset sales, funding contagion, and liquidity stress via margin calls. We apply this stress test model to data provided by *S&P Global Market Intelligence*, the *ECB Statistical Warehouse*, the 2018 *European Banking Authority* (EBA) stress test results, allowing us to initialise balance sheets of European banks and non-banks. In line with [Hałaj & Kok \(2013\)](#), [Kok & Montagna \(2013\)](#), we reconstruct the interbank, secured funding and common asset holding networks, which interconnect these institutions.

Our results evince that system-wide stress tests are *necessary* complements to microprudential stress tests: for a given outcome of a microprudential stress test, which only focusses on the first-order losses, the financial system, when higher-order losses are taken into account, can be either stable or unstable depending on the financial architecture. Hence, the outcome of the microprudential stress test does *not* give insight into systemic risk: macroprudential stress tests are required. Specifically, we show this in the context of the 2018 European Banking Authority (EBA) stress test. Given the 2018 EBA stress

test outcome, we find that the European financial system may be stable or not depending on the resolution regime, usability and size of regulatory buffers, among others. Next, we demonstrate that financial resilience decreases if regulatory buffers are seen to be less usable by banks. If regulatory buffers are not treated as usable, then regulatory buffers *de facto* act as capital requirements. In such case, if an adverse shock threatens an institution to breach its capital buffers constraints, it is forced to delever, which tends to have a destabilising effect on the financial markets. We reveal that the size of usable regulatory buffers that is required to maintain stability is underestimated if the interaction between exposure loss contagion, funding contagion, overlapping portfolio contagion and margin call contagion is not taken into account.

In sum, while current microprudential stress tests remain valuable, our findings suggest that they should be complemented by system-wide stress tests when evaluating financial stability and calibrating capital buffers.¹

Authors of Paper - J. Doyne Farmer & Alissa M. Kleinnijenhuis (first author), Paul Nahai-Williamson (Bank of England) & Thom Wetzer. The paper will be available here: [Farmer et al. \(2020\)](#).

3.2 Introduction

In a highly connected financial system, seemingly local shocks can be amplified and propagated to take on systemic importance. The salience of this observation is powerfully illustrated in Brunnermeier’s review of the dynamics of the global financial crisis ([Brunnermeier et al. \(2009\)](#)). Problems that started in the real economy with increasing sub-prime mortgage defaults quickly spread throughout the financial system through various amplification channels. Asset price falls on mortgage-backed securities prompted margin calls that put pressure on hedge funds, leading to a round of correlated selling that further depressed prices and impaired market liquidity ([Gorton & Metrick \(2012\)](#)). Banks’ common exposures to these assets put further pressure on their solvency, leading to the wholesale funding run on Lehman Brothers ([Copeland et al. \(2014\)](#)). Its subsequent default triggered solvency contagion to hedge funds, banks and money market funds as well as a freeze in interbank markets. Given these dynamics, the challenge for regulators is to constantly evaluate these risks to the resilience of individual institutions and the

¹We thank David Aikman, Christoph Aymanns, Luca Enriques, Co-Pierre Georg, Anne-Caroline Hüser, Esti Kemp, Alan D. Morrison, James Paulin, Anton Pichler, Maarten Scholl, Garbrand Wiersema, and participants at (internal) seminars at the Bank of England, European Central Bank, Institute for New Economic Thinking at the Oxford Martin School, Oxford-Man Institute of Quantitative Finance, South African Reserve Bank, University of Cape Town, Office of Financial Research and Bank of Mexico. We dedicate this paper to Rafa Baptista Ochoa (in memoriam), who made invaluable contributions at the early stages of this project. The usual disclaimers apply.

financial system as a whole. In doing so, they must understand the various financial institutions involved, their interconnections, and their interactions under stress from multiple contagion mechanisms.

In this paper, we address those challenges and propose a generic framework for the development of system-wide financial stress tests with multiple interacting contagion and amplification channels as well as heterogeneous financial institutions. This framework conceptualises financial systems through the lens of five building blocks - financial institutions, contracts, markets, constraints, and behaviour. These blocks can be flexibly implemented to form a dynamic multiplex network using the accompanying software engine and library (the ‘Economic Simulation Library’, or ‘ESL’). Depending on the needs of regulators and researchers and the data they have access to, this framework (and the software that implements it) supports both stylised stress testing models as well as large-scale, data-driven models that map out the financial system with a high degree of verisimilitude.²

Using this framework, we implement a system-wide stress test model for the European financial system. This stress test model captures the solvency-liquidity nexus, incorporates four interacting amplification channels³, and takes account of the heterogeneity of the financial institutions⁴. To evaluate the complementary value of this system-wide approach, we implement our stress-testing model as a ‘macroprudential layer’ on top of the regular micro-prudential European Banking Authority (EBA) stress test from 2018 and compare the stress test results.

This comparison yields three findings, which are robust to extensive sensitivity and robustness checks.⁵ First, depending on shock-amplifying⁶ capacity of the financial system, the system may be stable or unstable for a given microprudential stress test outcome. This strongly suggests that there is a complementary role for system-wide stress tests when evaluating financial stability: system-wide stress tests can elucidate how the same

²This software package, as well as accompanying documentation, sample implementations, and robustness checks, is freely accessible at: <https://github.com/ox-inet-resilience/resilience>.

³We include amplification associated with default contagion, price-mediated contagion via asset sales, funding contagion, and liquidity stress via margin calls.

⁴We limit ourselves to three classes of financial institutions - banks, asset managers, and hedge funds - and allow for heterogeneity within these classes.

⁵Our findings are robust to a range of modelling assumptions for institutional behaviour, the severity of the initial shock to the financial system, the price impacts of asset sales, and the number of contagion channels in operation. Our robustness and sensitivity checks are outlined in detail in Appendix A.3.

⁶We speak of amplification of exogenous shocks if the systemic risk measure including endogenous shocks is higher than without.

set of initial shocks may be endogenously amplified to starkly different degrees depending on the characteristics of different financial systems. We find, for example, that for the same microprudential stress test outcome of $\mathbb{E} = 0\%$ a financial system can be highly fragile (systemic risk approximately equal to $\mathbb{E} \approx 90\%$)⁷ or relatively resilient ($\mathbb{E} \approx 10\%$) depending on assumptions for the efficacy of resolution regime for banks.

In obtaining this result, we show that the outcome of a system-wide stress test will differ depending on the (interacting) contagion channels that are taken into account. We confirm the result that interacting contagion channels can produce significantly higher rates of bank failure (by as much as 300%) than suggested by the sum of failures when they act in isolation (Caccioli et al. (2013), Kok & Montagna (2013), Poledna et al. (2015), Hüser & Kok (2019), Wiersema et al. (2019)). Our model can serve as a tool to evaluate, under different market conditions and for different financial systems, which set of amplification mechanisms is most destabilising. For example, we show that when markets for institutions' tradable assets are liquid, solvency contagion risk is the most significant mechanism, whereas when markets are less liquid price-mediated contagion via asset sales becomes more dominant and amplifies the risks associated with other channels. We also show that the inclusion of heterogeneous financial institutions, and in particular non-banks, changes the magnitude of systemic risk.

Our second finding is that the willingness of banks to draw on their capital buffers to absorb losses - which we term the 'usability' of capital buffers - significantly affects the shock-amplifying tendency of a financial system. The actions banks take to avoid using their buffers in response to an adverse shock, which could for example be motivated by a desire to avoid regulatory restrictions on dividend payments, can generate pro-cyclical dynamics that substantially increase system-wide losses. In light of this result, regulators should evaluate how the design and enforcement of regulatory buffers may affect their 'usability' in times of financial stress, and be mindful of the financial stability implications of buffers that produce behavioural effects similar to those of regulatory requirements (Goodhart et al. (2008), Goodhart (2013)).

Finally, we find that microprudential stress tests may underestimate the (usable) regulatory buffer that is required to ensure the resilience of individual institutions and the financial system as a whole. Currently, regulators mostly use microprudential stress tests to calibrate the discretionary time-varying capital requirements under Pillar II of

⁷In line with Gai & Kapadia (2010), Gai et al. (2011), Paulin et al. (2018), we measure systemic risk \mathbb{E} by the average fraction of defaults in a systemic event.

the Basel capital adequacy framework and, in the United Kingdom, the countercyclical capital buffer. Our findings suggest that system-wide stress tests can meaningfully complement microprudential stress tests when calibrating capital buffers.

Based on the method's foundations, credible system-wide stress test models can be built to crown the macroprudential toolkit – as the only tool that explicitly captures endogenous dynamics critical to systemic risk (Danielsson & Shin (2003)). This opens up avenues to propel system-wide stress tests from the verge to the heart of macroprudential policymaking, complementing other policy tools. Furthermore, the framework supports flexibly tailoring of system-wide stress test models to answer specific policy and research questions. The European financial system model presented here and the bail-in stress test described in our companion paper (Goodhart et al. (2020) presented in Chapter 4) are cases in point.

The paper proceeds as follows. Section 3.3 specifies our contribution to the literature. Section 3.4 sets out the foundations of our generic framework for system-wide stress tests, and in Section 3.5 we use this framework to develop the model for the system-wide stress test of the European financial system. Section 4.7 presents the results of the experiments we ran on the our system-wide stress test, and we discuss the policy implications of these findings in Section 3.7.

3.3 Relevant Literature

3.3.1 Modelling System-Wide Stress Dynamics

We are by no means the first to attempt tackling the challenge of developing system-wide stress testing models (for an overview, see Aymanns et al. (2018)). Central banks have been at the vanguard (Burrows et al. (2012), Figue (2017), Kok & Montagna (2013), Dees & Henry (2017), Aikman et al. (2019)), and have been joined by academics (e.g. Cont & Schaanning (2017)). However important these contributions may be, they do not propose a generic approach to modelling system-wide dynamics. On that front, we make three contributions.

First, we outline a generic framework that allows for the systematic modelling of interacting contagion mechanisms. We distinguish two modes of contagion - node and contract amplification (see Section 3.4.2) - which we use to study four specific contagion mechanisms (Section 3.5.2): overlapping portfolio contagion, exposure loss contagion, funding contagion, and collateral contagion (via margin calls). Existing literature tends to cover subsets of these (interacting) contagion mechanisms using modelling approaches

that cannot be easily generalised to include other contagion channels. For example, [Kok & Montagna \(2013\)](#) consider the first three contagion mechanisms, [Caccioli et al. \(2013\)](#) and [Hüser & Kok \(2019\)](#) explore the first two, [Poledna et al. \(2015\)](#) investigate funding contagion for different contracts, and [Brunnermeier & Pedersen \(2009\)](#) consider the interaction of funding and market liquidity (a form of collateral contagion). Because the interaction of contagion mechanisms may amplify systemic risk (see e.g. [Kok & Montagna \(2013\)](#)), this modelling innovation has practical value for those looking to evaluate the resilience of the financial system.

Our second contribution is of a similar nature. Our generic framework allows for the joint modelling of heterogeneous financial institutions. That heterogeneity means that we can account not only for differences between various types of institutions (e.g. banks and non-banks), but also for differences within these groups (e.g. account for differences between banks). As we outline in [Section 3.4.1](#), we do so by characterising each institutions on the basis of the contracts (or, if the data is not so granular, contractual types) that each institution has on its balance sheets, the constraints (contractual, regulatory, or otherwise) it is subject to, and the behavioural assumptions we adopt. Because institutions in our generic framework can be distinguished along these dimensions, the framework can host stress testing models that reflect the heterogeneity of behavioural objectives, constraints and balance sheet resources that characterises the financial system ([Danielsson & Shin \(2003\)](#)).

Given the important interactions and interdependencies between banks and non-banks in modern finance (see e.g. [Burrows et al. \(2015\)](#), [ECB \(2017\)](#), [Pozsar & Singh \(2011\)](#)), capturing the heterogeneity of financial institutions is central to the success of system-wide stress tests. As bank/non-bank linkages continue to become more significant (e.g. [BIS \(2019\)](#)), the importance of this modelling innovation to regulators is likely to grow. This is especially true because the capacity of existing models to capture heterogeneity remains, however, limited ([Halaj \(2018\)](#), [Baranova et al. \(2017\)](#)).

Our third modelling contribution relates to the modelling of multiple interacting constraints arising from regulation and contracts. It is clear that such constraints can drive behaviour in times of financial stress (see e.g. [Greenwood et al. \(2015\)](#), [Duarte & Eisenbach \(2015\)](#), [Aymanns et al. \(2016\)](#), [Caccioli et al. \(2014\)](#)) and, moreover, that financial institutions face an increasingly complex plethora of interacting and overlapping regulatory constraints (see e.g. [Armour et al. \(2016\)](#)). Despite that reality, existing contagion models typically model either the leverage ratio⁸ or the risk-weighted capital ratio (see e.g. [Kok & Montagna \(2013\)](#), [Cifuentes et al. \(2005\)](#)), and they rarely implement the

⁸Most consider leverage targeting (e.g. [Greenwood et al. \(2015\)](#), [Duarte & Eisenbach \(2015\)](#)), where the buffer and target *de facto* coincide, some consider a distinct buffer and target for the leverage constraint ([Cont & Schaanning \(2017\)](#), [Bookstaber, Paddrik & Tivnan \(2014\)](#)).

Basel III liquidity constraints (see e.g. [De Haan & van den End \(2013\)](#), [Aldasoro et al. \(2017\)](#)). Where models only consider one constraint, they usually also consider one common rule or ‘pecking order’ to determine the actions that institutions take in response to shocks (e.g. proportional liquidation of assets as done by [Greenwood et al. \(2015\)](#)), or liquidation of the most-liquid assets first as in [Halaj \(2018\)](#)). A recent paper by [Coen et al. \(2019\)](#) is a notable exception: the authors model banks’ decisions to sell tradable assets in response to solvency and liquidity shocks by optimizing asset sales to minimise losses while meeting three regulatory constraints (i.e. the leverage ratio, risk-weighted capital ratio and liquidity coverage ratio).

In our generic framework, we propose an approach to modelling multiple interacting constraints that can, again, be easily generalised. For each regulatory ratio, we propose that the institution sets a self-imposed buffer value. Once it reaches this buffer value, the institution acts to either comply with a regulatory buffer standard or to move towards a self-chosen target value (see Section 3.5.3). This approach is consistent with the empirical findings of [Adrian & Shin \(2010\)](#) and has intuitive appeal. In line with the approach in [Coen et al. \(2019\)](#), we also employ different pecking orders for institutions’ actions depending on the constraint that binds. This approach reflects the reality that not all constraints can be (effectively) adhered to by taking the same set of actions. For instance, the pecking orders for the leverage ratio and risk-weighted capital ratio should be different, because liquidating non-cash, zero risk-weighted assets can reduce the leverage ratio but will not improve the risk-weighted capital ratio.

3.3.2 Stress Tests and Prudential Regulation

The development of a generic framework allows us to develop a system-wide stress test of the European financial system.⁹ Using this system-wide stress test, we run a number of experiments that yield three main takeaways for policymakers.

First, we find that system-wide stress tests are necessary complements to microprudential stress tests. A large body of literature has shown that systemic risk may be underestimated if non-linear contagion effects that may amplify initial shocks are not considered (see e.g. [Cont & Schaanning \(2017\)](#)). Moreover, various authors have applied a ‘macroprudential overlay’ to regulatory microprudential stress test (see e.g. [Burrows et al. \(2012\)](#), [Dees & Henry \(2017\)](#), [Paddrik et al. \(2016\)](#), [Paddrik & Young \(2017\)](#)). However, we are the first to systematically compare the system-wide (including interacting contagion channels and heterogeneous agents) and microprudential stress test results

⁹We stress that this model implementation by no means exhausts the options offered by the generic framework. We have developed this system-wide stress test model purely with a view to showcase the generic framework and to study important policy questions.

for different (scaled) regulatory stress test scenarios. We confirm unambiguously that interacting contagion channels can produce significantly higher rates of bank failure (by as much as 300%) than suggested by the sum of failures when they act in isolation. This suggests that microprudential stress tests alone will overstate resilience and give false comfort to regulators, financial markets, and the public at large.

Second, we contribute to the literature on the design of regulatory capital requirements. Existing literature recognises that capital requirements may lead to pro-cyclical responses if they cause financial institutions to act in ways that are individually rational but collectively destabilising – for example by deleveraging during crises [BIS \(2008\)](#), [Aymanns et al. \(2016\)](#). Regulators replaced strict capital and liquidity *requirements* by *buffers* – which institutions can (temporarily) draw on without breaching their regulatory obligations – so that institutions can absorb shocks and refrain from taking pro-cyclical actions ([BIS \(2008, 2009, 2013\)](#), [Drehmann et al. \(2010\)](#)). [Goodhart et al. \(2008\)](#) and [Goodhart \(2013\)](#) have emphasised that these buffers should be usable: ‘required liquidity is not true, usable liquidity. Nor might I add, is required minimum capital fully usable capital from the point of view of a bank’. We show, in a system-wide setting, how such usability affects the resilience of the financial system.

Third, we show that the calibration of these buffers should be based not only on microprudential stress tests, but also on system-wide stress tests. In discussing the calibration of the capital (or liquidity) frameworks, the existing literature does not differentiate between *requirements* and *buffers*, and also does not consider how usability of capital (or liquidity) would affect resilience (e.g. [Battiston et al. \(n.d.\)](#), [Greenwood et al. \(2015\)](#), [Cont & Schaanning \(2017\)](#), [Duarte & Eisenbach \(2015\)](#)). [Aymanns & Farmer \(2015\)](#) show that higher capital requirements may be *destabilising*. Using the ability of the generic framework to capture pre-default contagion that arises from interacting contagion mechanisms, we show that the size of regulatory *buffers* required to maintain financial stability will be underestimated if microprudential stress tests are used for calibration.

3.4 A Framework for System-Wide Stress Tests with Heterogeneous Institutions

In this section, we outline a generic framework for system-wide stress tests with heterogeneous agents. At the core of our framework are five building blocks that we use to represent financial systems. We start with *financial institutions and their balance sheets*, which are populated by *financial contracts* that connect them. Together, these two building blocks - when implemented at a level of granularity that corresponds to the available data and the needs of the modeller - create a multiplex network, with a separate layer for

Table 3.1: Shows the time evolution of the multi-layered network in a system-wide stress test consisting of five building blocks (highlighted in italics). The adverse stress scenario is applied once at time t_0 . The (contagious) endogenous dynamics are iteratively generated in discrete-event simulation by repeating the substeps $t_{x,1}$, $t_{x,2}$ and $t_{x,3}$ in a timestep t_x . The inner time steps represent a series of rounds, which take an infinitesimal amount of time and are therefore said to occur in an instant. Once, the substeps are completed, the outer time step increases from t_x to t_{x+1} through fixed-increment time progression, repeating the substeps for another round as long as the stopping condition has not been satisfied.

Time step			
t_0	Initial, adverse scenario		
$t_{x,1}$	Impact on market	Impact on <i>balance sheets</i>	Affects <i>contractual obligations</i> Affects the valuation of <i>contracts</i>
$t_{x,2}$	Observations	<i>Contractual obligations</i> Variables relative to their regulatory, market or internal-risk <i>constraints</i> Performance relative to other objectives, such as profit objectives (if any)	
$t_{x,3}$	Behavioural Actions	Honour <i>contractual obligations</i> Move away from regulatory, market or internal <i>constraints</i> Execute strategy to meet other objectives, such as profit objectives (if any)	
Next time step		If the system has not stabilised and if the maximum number of simulation time steps, T^S , has not been exceeded, increase the timestep counter, x , to $x = x + 1$ ($x=1$ initially) and repeat the three substeps per time step t : $t_{x,1}$, $t_{x,2}$, $t_{x,3}$. Else, stop the stress test.	

each contractual type, that represents the financial system. Studying the topology of this network can already yield valuable insight about systemic risk (see e.g. [Battiston et al. \(2012\)](#)), but to be able to also study the dynamics operating on that network we add three more building blocks: the *markets* in which contracts are traded, the *constraints* - whether arising from contractual obligations, market pressure, or regulatory requirements - that institutions are subjected to, and the *behavioural assumptions* that stipulate how, in the decision-space left by the constraints, each institution will act. Table 3.1 sets out the various steps based around which the static network evolves. This static and dynamic representation of the financial system is operationalised using a newly built simulation engine that can host large-scale data-driven models.

Section 3.4.1 outlines the five building blocks used to represent financial systems in greater detail. In Section 3.4.2, we discuss the endogenous (amplifying) dynamics that can arise in this generic framework, and how we conceptualise them. We conclude by highlighting some important design principles of the software that we developed to host these stress-test models in Section 3.4.3.

3.4.1 Five Building Blocks to Represent Financial Systems

Our generic framework uses five building blocks to represent financial systems and to, subsequently, study the systemic risk that is endogenously created by heterogeneous financial institutions, as called for by [Danielsson & Shin \(2003\)](#). We discuss these five

building block in turn.

Financial Institutions & Financial Contracts We represent financial institutions at a representative or individual level to reflect the importance of institutional and sectoral heterogeneity. Each institution has a unique balance sheet that is composed of a *collection of financial contracts* (assets and liabilities), rather than merely a list of aggregate values per asset class or aggregate exposures to a specific counterparty. Our generic framework allows us to model each individual contract and include information on (1) the *parties* to the contract, (2) the contract’s *value*, ‘valuation function’ (under the applicable accounting regime), and the inputs to that valuation function¹⁰, and (3) the set of (contingent) *liquidity* obligations, including the contract’s ‘liquidity function’ and its inputs.¹¹

By modelling financial institutions and their contracts in this way, we achieve at least two valuable results. First, we allow for significant heterogeneity between institutions, because institutions are characterised - in the model as in real life - by the institution-specific collection of financial contracts it holds. Second, we can construct the network of interconnections between financial institutions, both in a static and dynamic sense. The information on counterparties enables the software to create edges between different institutions (nodes) in the financial network, which gives us the static network. Moreover, when studying the dynamic network, the contract-specific information coupled with basic accounting¹² makes it possible to update contract valuations and balance sheet variables following initial or endogenous shocks, and allows to model the liquidity pressures that institutions may face due to margin calls or decisions by creditors not to roll over funding.¹³

The level of granularity that can be adopted in a specific model implementation will largely depend on the granularity of the data on which the model can be calibrated. Since the financial crisis, the mandate for regulators to gather contract-specific data has increased (see e.g. [Abad et al. \(2016\)](#)). However, in case such data is not available, the

¹⁰This tells us how the price of the contract is determined, the contractual maturity and whether a contract is secured or not, etc.

¹¹For example, the valuation function of tradable assets takes the market price as an input, and multiplies this by the unit of assets held to determine the balance sheet value of the asset (see equation 3.5.2). Similarly, the collateral price is an input to the ‘liquidity function’ of repurchase agreements: margin calls are determined based on the difference between the notional of the repo and the haircutted collateral price times the units of the collateral placed (see equation 3.8).

¹²Our generic framework supports various accounting standards, which are made available in the online repository of the Economic Simulation Library.

¹³In Section 3.4.2 we discuss why understanding the set of valuation and liquidity shocks that an institution faces at each point in time is important to understand contagion and its spread via the network of financial contracts.

second-best approach is to rely on network reconstruction methods to estimate contract-level information from aggregate data. We discuss such methods in Appendix [A.1.2](#).

Markets Contracts originate in financial markets, and it is there that their price is determined by interacting market participants. Different types of contracts are traded within distinct markets; equities, for example, are typically traded on exchanges, while interbank contracts originate in interbank markets ([Heider et al. \(2009\)](#)). Each of these markets has its own dynamics and characteristics, and the degree to which these are taken into account depend on the modeller’s objective. Reduced-form price impact functions may be sufficient to capture the impact of forced sales on asset prices, as is indeed commonly done ([Duarte & Eisenbach \(2015\)](#)), but more detailed modelling of order books that process buy and sell orders from institutions to set prices may be needed when studying price-formation and market liquidity in greater detail.¹⁴ In our generic framework, every asset or contractual type can have its own associated market, so that users can build in the appropriate market mechanism(s) for each asset or contract - in a level of detail they consider optimal - and study the associated risks of those markets.

Constraints Institutional behaviour is governed by *rules and constraints*. As asset values evolve, the financial network changes and/or exogenous shocks are applied, financial institutions update their balance sheets. In particular, institutions assess whether they have breached, or are close to breaching, regulatory¹⁵, market¹⁶, or contractual¹⁷ constraints, or their internal risk limits¹⁸. These rules and constraints limit the time-dependent set of actions available to each financial institution. They could include rules for operating under normal times – for example optimising rules to determine portfolio allocation, and internal risk limits that impact on trading behaviour – but most importantly will include constraints that drive behaviour in periods of stress.

¹⁴The market mechanism for the price formation of each type of contract is typically in the public domain and can thus be modelled, or else can be usually reasonably estimated based on the standard market mechanism for such a contract. While we can generally know the way a market functions – for example exchange-trading via an order book, or intermediated by a dealer – to model the market *dynamics*, we need to combine this with the *behaviour* of market participants. In dealer-intermediated markets for example, the behaviour of the dealers in response to buy and sell orders is an important part of the price-setting mechanism (see e.g. [Baranova et al. \(2017\)](#)).

¹⁵Examples of regulatory constraints include minimum leverage and risk-based capital ratios for banks

¹⁶Market-based constraints are implicit minima that the market sets on, for example, capital ratios, for an institution to maintain access to market-based funding ([Burrows et al. \(2012\)](#), [Bookstaber, Paddrik & Tivnan \(2014\)](#)). Such limits may be stricter than those imposed by regulators.

¹⁷Contractual constraints include obligations to exchange margin or to repay liabilities at maturity.

¹⁸Internal-risk limits are institution-specific limits, such as value-at-risk (VaR) limits ([Berkowitz & O’Brien \(2002\)](#)), which are typically set by the risk managers of the institution.

Importantly, these constraints can act both to trigger action during balance sheet distress, and to limit institutions' ability or appetite to take actions that could limit distress and support market functioning. To take banks as an example: falling leverage ratios may cause some banks to fire sell assets or reduce provision of client funding; and the ability of other banks to step in to buy discounted assets or meet clients' funding needs (and so reduce systemic stress) could be restricted by their own regulatory constraints. This is consistent with the observation that the state of balance sheet capacity within and across sectors – and the degree of similarity between sectors in the constraints they face – is likely to be a key determinant of systemic vulnerability to shocks.

The constraint that binds most drives behaviour. When constraints bind, the set of actions that institutions can take becomes limited to those consistent with the behavioural objective of not breaching binding constraints. Therefore, constraints add further heterogeneity to the stress testing framework. Regulatory constraints, for example, differ between institution types, with banks facing different regulatory constraints than hedge funds. Similarly, because institutions have different collections of contracts on their balance sheets, they will also face different contractual constraints. Which constraint binds can differ from one institution to another, and can even be context-dependent¹⁹ – which in turn affects firm-specific behaviour and ultimately system-level dynamics.

Behaviour Behaviour is central to understanding systemic risk; it is also the most challenging aspect to pin down and model (Farmer & Lo (1999), Farmer (2002), Lo (2017), Aymanns et al. (2018)). To fully capture the build-up and crystallisation of systemic risk, ultimately an understanding of behaviour under both 'normal times' and stressed times is important. Understanding how agents optimise their balance sheets subject to the constraints described above to meet their business targets (for example, to maximise return on equity or shareholder value), can give insights into the potential impacts of policies on latent risk in the financial system. Also, understanding the types of behaviour that institutions may forcibly display when under severe stress is key to modelling the dynamics that could occur when risks start to crystallise.

Behaviour in our generic framework means making decisions on buying and selling assets; and opening, continuing or terminating contractual relationships (for example by choosing not to roll-over a funding relationship). Institutions can also choose not to honour contractual commitments, with the potential outcome that they default. To understand the propagation and amplification of stress, we particularly focus on how the constraints that institutions face can limit their options and force certain behaviours.

¹⁹Regulators and contractual counterparties may, for example, loosen constraints in times of crisis when they fear that rigid enforcement might lead to default, see e.g. Pistor (2013) and Awrey (2019). We do not consider such dynamics here.

These behaviours will be institution specific – but generally speaking, they will relate to issues of solvency/profitability and/or liquidity.

To illustrate, consider how banks may take different actions when risk-based capital ratios bind than when the leverage ratio binds. Consider a bank that has significant trading activities, and a business extending secured funding to clients. Following a severe shock:

1. If the risk-based ratio binds, the bank will need to deleverage in risk-weighted asset space. Reducing the secured funding it extends to clients may achieve this to some extent, but the collateralised nature of the exposure limits the risk-weighted assets reduction that can be achieved. A more effective way can be to sell trading assets with high risk-weights – assets that tend to be less liquid. In this case, the actions of the bank may have an impact on market liquidity.
2. If on the other hand the leverage ratio binds, the bank can rapidly deleverage by cutting its provision of secured funding to clients; or by reducing low risk-weighted assets such as cash or government bonds. If the bank has surplus liquidity, the latter option would cause no or limited spillovers to the rest of the financial system; if it has to pull funding from its clients however, the bank may be forced to liquidate assets to address the funding shortfall, again potentially leading to a market impact.

In reality, where multiple constraints are at play, some form of optimisation will be required to meet all relevant constraints, and the decision-making becomes more complex. Not only that, but institutions are likely to at least attempt to take into account the impacts of their own actions and those of other financial market participants. While our framework can support such optimisation in principle, implementing it in large-scale system-wide stress testing models remains challenging.

For these reasons, understanding institutional behaviour is a key area for ongoing research and model development. Behaviour also represents the biggest unknown we face²⁰; in principle, data on institutions' balance sheets and constraints could be sourced, and the mechanics of market functioning could be modelled. System-wide stress tests build using our generic framework are, however, explicitly conditional on the behavioural assumptions made. We therefore set up our framework explicitly to enable users to easily explore the impacts of different assumptions on behaviour, giving them the flexibility to investigate outcomes conditional on plausible assumptions, and their sensitivity to these assumptions. Such sensitivity analyses can themselves convey valuable information, for

²⁰That does not mean that these assumptions are completely uninformed: market surveys, for example, can be helpful.

example about the types of behaviour that are most destabilising and should therefore be avoided.

3.4.2 Contagion and Amplification

In our model, a financial system as represented using the five building blocks can face two types of contagion and amplification. The first, ‘node amplification’, takes place within the nodes (financial institutions) of the financial network, whereas the second, ‘edge amplification’, takes place along the edges (financial contracts).

Amplification within nodes takes place when an incoming shock – whether that is a valuation or a liquidity shock – is passed on to a new outgoing shock. How a shock is passed on depends on the internal operation of the financial institution (the node), specifically its behavioural response (subject to its set of constraints and available balance sheet resources). Overlapping portfolio contagion due to asset sales is an example of node amplification. A detailed treatment of node amplification can be found in [Wiersema et al. \(2019\)](#).

Edge amplification, on the other hand, occurs when shocks to one type of contract cause shocks to another type of contract. This mechanism is more mechanical in nature, and can be captured by modelling how shocks to one contract may act as inputs to the ‘valuation function’ and ‘liquidity function’ of another contract. For example, valuation shocks to tradable assets can lead to margin calls (liquidity shocks) on repurchase agreements; and mortgage loan defaults can result in valuation shocks to mortgage-backed securities. In Section 3.5.2, we implement this approach for a stress test of the European financial system.

3.4.3 Key Design Choices for Simulation Software

To operationalise the approach outlined above, we have developed object-oriented modelling software (a simulation engine and accompanying software library) that can support modelling of the financial system with a potentially high degree of verisimilitude (for example when using transaction-level data). Otherwise stated, by embedding the generic framework in simulation software it becomes possible to take full advantage of rapid digitisation and standardisation of regulatory and market data (see e.g. [Judge & Berner \(2019\)](#)) to run advanced, data-driven system-wide stress testing models at scale.²¹

When designing the software, we have applied core concepts of sound software engineering. This not only makes the software easier to use, but also ensures that our

²¹The fully documented simulation engine, and its accompanying software library, is available at <https://github.com/ox-inet-resilience/resilience>. The contributors to this library are: Alissa M. Kleinnijenhuis, Rudy Tanin and Rafa Baptista Ochoa.

generic framework supports flexible and modular models whose operations are transparent. Modularity means that the underlying code is divided into the building blocks described above, which clarifies the structure of the stress testing model that is used. It allows users to examine the full network, but also to examine various contagion channels in isolation, or to only examine specific sectors or institutions. Flexibility means that the building blocks that make up a specific model, as well as the modelling assumptions more generally, can be easily adjusted. It also allows for implementation of models at various levels of abstraction and realism, depending on what is most appropriate. It is intended to be transparent; empowering the user to track the operations of the simulation by producing (intermediate) outputs in readily understandable forms, in order to avoid a ‘black box’ problem.²²

A major challenge in implementing system-wide stress testing models is to capture the concurrency of financial markets, with many different institutions acting simultaneously. Stress test models that are implemented using simulation-based software are often sequential, which means that the order of computations becomes a key determinant of the outcome of the stress test, thereby artificially skewing these results. One way to address this problem is to randomly shuffle the order in which institutions act (see eg [Fique \(2017\)](#)). Although this takes away systematic biases, it does not rid the simulation from biases within a time-step. In a fire sale scenario, for example, institutions that happen to be first in line could have a substantial advantage and may therefore appear more resilient than they are in reality. An alternative approach, to use parallel computer code to run system-wide stress tests, is unappealing because parallel code is error-prone.

We therefore propose a novel way to solve the problem of order dependence – which we refer to as the ‘mailbox system’. Each institution has its own mailbox. Whenever an institution acts (e.g. pulls funding, gives a margin call), the notification of that action ends up in the ‘unread mailbox’ of the relevant counterparty (or counterparties). This message will only be ‘read’ *after* every institution in a given time step has acted, at which point the simulation engine will execute all these actions at once.²³ Accordingly, actions of institutions that affect markets (such as fire sales) will only be executed at the end of the time step, even though notifications of undertaken actions will be collected during

²²In addition, we have also followed the design principles of readability (the reader should be able to read and understand the implementation in a short amount of time), performance (the code should execute as fast as possible so long as this does not come at the cost of readability), and reproducibility (the reader should be able to run the simulations and obtain an identical outcome).

²³Of course, it is theoretically possible to account for speed differentials between institutions by making some institutions slower to send or open their messages, so that they would take multiple timesteps to complete a task that other institutions can complete in one time step. The practical effect would be that this institution responds more slowly to market developments. We do not explore this option here.

the sequence of the acting institutions in each time step.²⁴

To illustrate why a *messaging-mailbox system* is necessary and random shuffling is not sufficient to achieve order independence, we ran comparative benchmark on our simple stress test model which only consists of overlapping portfolio contagion.²⁵ When we use the random shuffling, we find that the standard deviation of the average extent of systemic event \mathbb{E} (see Appendix A.1.3) soon reaches a certain minimum amount that cannot decrease no matter how high the number N of simulation runs is. The reason that the systemic risk outcome is severely affected by which specific group of institutions gain an artificial advantage in a specific time step, leaving clusters of outcomes due to which a markedly positive standard deviation is maintained. On the other hand, when we use the messaging-mailbox system, we find that the standard deviation of our systemic risk measure \mathbb{E} soon decays to zero, since in simultaneous version the same outcome is the same in a simulation run n , regardless of the shuffle.

3.5 A System-Wide Stress Test Model for the European Financial System

Using our generic framework and simulation software, we implement a system-wide stress testing model for the European financial system to study its contagion dynamics. This model combines multiple interacting contagion mechanisms and constraints, and allows us to assess how institutional behaviour under stress can amplify an initial adverse shock. Moreover, we will use the stress-test model to study the usability and size of regulatory capital buffers that is needed to mitigate systemic risk. The model showcases the power of the generic framework. We stress, however, that this framework can be, and is in fact designed to be, used to support different models that focus on different research or policy questions and utilise different data types.

The stress test model includes three types of financial institutions: banks, asset managers and hedge funds. Because we have relatively detailed institution-level data for banks, we can model them at the institution-level. We model asset managers and hedge funds using representative institutions, however, due to a lack of suitable data. The banks in our model are directly connected via unsecured interbank lending and borrowing, and are indirectly connected via common holdings of tradable assets. Asset managers and hedge funds also hold these tradable assets, and hedge funds are also directly connected

²⁴An alternative implementation of the *messaging-mailbox system*, which enables execution of institutions' actions to be distributed across multiple CPUs, can be found in abcEconomics. See: <https://github.com/ab-ce/abce>.

²⁵See: https://github.com/ox-inet-resilience/firesale_stresstest/blob/master/other_simulations/random_shuffling.py.

with banks via repo funding. In this section, we discuss the setup of the model by outlining how to model (1) each type of institution, (2) the contracts they trade in, (3) the various constraints they face, (4) the markets in which they trade, and (5) their behaviour. In Appendix A.1.2, we provide further information on the initialisation of the model. We include a table of notation in Appendix B.2.

3.5.1 Financial Institutions and their Constraints

Banks We consider the most systemically important banks in the European Union (those that took part in the 2018 stress test from European Banking Authority (EBA)) and initialise their heterogeneous balance sheets using end-2017 data obtained from *S&P Global Market Intelligence*.²⁶ The stylised balance sheet of banks $i \in \mathcal{B}$, where \mathcal{B} is the set of banks, is depicted in Figure 4.1.

C_i , Cash	D_i , Deposits
Y_i , External Assets	\tilde{I}_i , Interbank Liabilities
T_i , Tradable Assets	\tilde{R}_i , Repos
I_i , Interbank Assets	\tilde{O}_i , Other Liabilities
R_i , Reverse Repos	
O_i , Other Assets	E_i , Equity

Figure 3.1: Stylised balance sheet of a bank $i \in \mathcal{B}$.

Figure 4.1 shows that the bank's assets A_i are given by the sum of its cash C_i , external assets Y_i , tradable assets T_i , interbank assets I_i , reverse repos R_i , and other assets O_i . It also shows that the bank's liabilities L_i are given by the sum of its deposits D_i , interbank liabilities \tilde{I}_i , repos \tilde{R}_i , and other liabilities \tilde{O}_i . The bank's book equity is defined by $E_i =: A_i - L_i$. Book equity is defined in the same way for other financial institutions.

Regulatory constraints Banks face a set of regulatory capital requirements and buffer standards as well as liquidity buffer standards, which we will discuss in turn.

We calculate two key Basel III capital ratios for our banks: the risk-weighted common tier I (CET1) capital ratio ρ_i and the leverage ratio λ_i . The risk-weighted capital ratio ρ_i

²⁶Due to data limitations (e.g. missing fields), we exclude a handful of banks and end up with a total of 42 banks.

is given by a bank's CET1 equity \tilde{E}_i over its risk-weighted assets Ω_i (RWAs), where the numerator is taken from data and the denominator is calculated by assigning risk weights ω_p to each asset type A_{ip} (where $p \in \mathcal{P}$ and \mathcal{P} is the set of assets) based on standard Basel III risk weights.²⁷ The leverage ratio is given by a bank's Tier 1 capital \tilde{E}_i^{T1} (equal to the sum of CET1 equity \tilde{E}_i and additional tier I (AT1) capital \tilde{E}_i^{AT1}) over its leverage exposure \hat{A}_i , both taken from the data.²⁸

Banks are required to meet the minimum CET1 capital ratios of 4.5% and minimum leverage ratio requirements of 3%²⁹ at all times:

$$\rho_i := \frac{\tilde{E}_i}{\Omega_i} = \frac{\tilde{E}_i}{\sum_{p \in \mathcal{P}} \omega_p A_{ip}} \geq \rho^M = 4.5\%. \quad (3.1)$$

and

$$\lambda_i := \frac{\tilde{E}_i^{T1}}{\hat{A}_i} \geq \lambda^M = 3\%. \quad (3.2)$$

A main objective of *capital requirements* is to ensure that banks have sufficient gone-concern loss absorbing capacity (Goodhart (2013)). Compliance with minimum capital requirements is a condition for doing business; a bank that falls below a capital requirement is likely to be closed down by regulators (Armour et al. (2016)). Given this, and in line with Kok & Montagna (2013), we assume that banks which breach minimum capital requirements will fail and are either liquidated or resolved (see Section 3.5.3 for details).

In addition to *capital requirements*, banks are subject to several different regulatory capital *buffers*, with the size of the combined buffer (CB) being heterogeneous across banks. The combined buffer is intended to ensure that banks have sufficient going-concern loss absorbing capacity to withstand a stress and can continue operating (Goodhart (2013)). To achieve this goal, buffers should be 'usable' in the sense that banks can absorb losses without failing or engaging in damaging or destabilising behaviour such as fire sales.

We note that when regulatory *buffers* have an effect that is *de facto* equivalent to *requirements*, they are not 'usable' from the point of view of the bank (Goodhart et al. (2008)). In other words, the regulatory buffers for capital and liquidity are only true

²⁷We use the standard Basel III risk-weights for all asset classes, except for the 'other asset' class O_i , where we choose the risk-weight such that it acts as a balancing item to ensure that total RWAs Ω_i match the data.

²⁸When modelling the impacts of stress on banks' capital ratios, we assume that CET1 equity \tilde{E}_i falls one-to-one with book equity E_i and that leverage exposure \hat{A}_i falls one-to-one with book assets A_i . That is to say, we assume a change in book equity or book assets leads to an equal change in CET1 equity or leverage exposure, and ensure that the difference matches the data at the start of the stress test.

²⁹We note that UK banks must meet a leverage ratio of 3.25%, with the leverage exposure measure excluding central bank claims matched by deposits in the same currency and of identical or longer maturity. For simplicity, we do not include these UK-specific requirements in this model

buffers if a bank considers these to be ‘usable’, else the regulatory buffer acts as a requirement. To illustrate this point, [Goodhart et al. \(2008\)](#) uses the metaphor of “the weary traveler who arrives at the railway station late at night, and, to his delight, sees a taxi there who could take him to his distant destination. He hails the taxi, but the taxi driver replies that he cannot take him, since local bylaws require that there must always be one taxi standing ready at the station”. Similarly, capital (liquidity) *requirements* are not *usable*.

While authorities have made clear that buffers are intended to be usable,³⁰ banks may still seek to avoid using them for various reasons. A BIS review, for example, notes that “only if supervisors allow banks to use buffers and banks do not resist their use, can buffers work to protect banks against macroeconomic downturns and taxpayers against bailouts. Supervisory discretion, excessive market discipline, and stigma attached to the use of buffers are some of the hurdles that may undermine their effectiveness” ([BCBS \(2016\)](#)). No empirical evidence exists yet, as to whether banks consider their regulatory buffers to be usable in a system-wide crisis, since the Basel III buffers have only been introduced since the last financial crisis of 2007-2008 and a new one has yet to occur.

Given these complications to draw on buffers, we investigate the impact that banks’ willingness to use these combined buffers (see Section [3.5.3](#) and [3.6.5](#)), and the actions they take to avoid having to use their buffers, has on systemic risk. In addition, for some experiments we consider a counterfactual in which regulatory buffer standards are larger (and assume that banks meet these buffers and hold correspondingly more capital resources) in order to assess how having higher buffers would impact systemic risk.

Each bank in our system has two combined buffers: a combined CET1 capital buffer ρ_i^{CB} and a single leverage ratio buffer λ_i^{CB} . These are composed of the buffer components discussed below, and given by

$$\rho_i^{CB} := \rho_i^{CCoB} + \rho_i^{CCyB} + \max\{\rho_i^{GSIB}, \rho_i^{DSIB}, \rho_i^{SR}\}; \quad (3.3)$$

$$\lambda_i^{CB} := \frac{1}{2}\rho_i^{G-SIB}. \quad (3.4)$$

The aim of the *capital conservation buffer* (CCoB) ρ_i^{CCoB} (set at 2.5%) is to promote capital conservation in the banking sector ([BIS \(2009\)](#)). Its introduction was prompted in part by the observation that many banks kept paying dividends during the financial crisis ([Greenwood et al. \(2017\)](#)) despite questions about their financial health, which

³⁰See e.g. [Prudential Regulation Authority \(2017\)](#), which explains that the parallel operation of the risk-weighted capital and leverage regimes in the UK “creates a ‘usable’ buffer, which is the amount of CET1 that a firm subject to both the risk-weighted capital and leverage regimes would currently be able to lose before breaching a minimum going-concern requirement.”

unnecessarily weakened their capital positions. Usage of the capital conservation buffer leads to increasing restrictions on dividend payments and staff bonus payments but is not forbidden; the buffer therefore attempts to create incentives for banks to maintain or rebuild their capital positions when they can, but also to draw on that capital when they must (Armour et al. (2016)).

The time-varying *countercyclical capital buffer* (CCyB) ρ_i^{CCyB} is set by regulators with the aim of counteracting procyclicality by building up a buffer in good times that can be drawn upon in bad times (Drehmann et al. (2010), Armour et al. (2016)). BIS (2010) recommend that the CCyB should be deployed when excess aggregate credit growth is judged to be associated with a build-up of system-wide risk, in order to ensure that the banking system has an additional capital buffer (on top of the capital conservation buffer and other requirements) to protect it against future potential losses. To align with Basel III the CCyB should vary between 0% and 2.5%, although national authority have discretion to increase the buffers further if they deem it necessary to do so to meet macroprudential objectives (Drehmann et al. (2010), BCBS (2011)). When risks materialise and the banking system is under stress, regulators can cut the buffer to 0%.

On top of the CCoB and the CCyB, a bank may have to hold additional risk-weighted buffers, including the globally systemically important bank (G-SIB) $\rho_i^{GSIB} \in [0\%, 3.5\%]$ surcharge,³¹ the domestically important bank surcharge (D-SIB) $\rho_i^{D-SIB} \in [0\%, 2\%]$,³² and the systemic risk buffer $\rho_i^{SR} \geq 1\%$.³³ Furthermore, G-SIBs also face an unweighted leverage buffer λ_i^{CB} , set at 50% of its G-SIB surcharge (FSB (2017)).

Banks also face liquidity buffers. We monitor banks' Liquidity Coverage Ratios (LCRs), Λ_i .³⁴ The LCR encourages banks to maintain an adequate stock of unencumbered high-quality liquid assets (HQLA) Q_i that can be converted easily and immediately in private markets, relative to its net outflows Θ_i in a thirty-day period of distress (BIS (2013), Gorton & Muir (2016)). BIS (2013) stipulates that net outflows Θ_i must be calculated as a function of the stressed asset inflows Θ_i^I and stressed liability outflows Θ_i^O , subject to a cap on the recognition of inflows at 75% of outflows (see denominator in

³¹The intention of the globally systemically important bank surcharge ρ_i^{G-SIB} is to limit negative externalities imposed on the global financial system associated with the most globally systemic banking institutions (BIS (2014)). The G-SIB surcharge ρ_i^{G-SIB} applies to globally systemically important institutions; other banks are given a G-SIB surcharge where $\rho_i^{G-SIB} = 0$.

³²The domestically systemically important bank surcharge is designed to address the negative externalities that domestically important banks pose on the domestic financial system and economy (BIS (2012, 2014)).

³³The objective of the systemic risk buffer ρ_i^{SR} is to prevent and mitigate long-term non-cyclical systemic or macroprudential risks not covered by Regulation (EU) No 575/2013.

³⁴Banks are also subject to a Net Stable Funding Ratio (NSFR). Because Cecchetti & Anil (2018) show that the LCR and NSFR typically do not bind simultaneously, and because our focus is on short-term contagion dynamics rather than longer-term funding risks (which are the focus of the NSFR), we do not consider the NSFR.

equation 3.5 below). The stressed asset inflows Θ_i^I and liability outflows Θ_i^O are computed by assigning stressed inflow $\tilde{\omega}_p$ and outflow rates $\tilde{\omega}_l$ to assets A_{ip} (for types $p \in \mathcal{P}$) and liabilities L_{il} (for types $l \in \mathcal{L}$) with maturities below 30 days.

We set the HQLA Q_i of bank i equal to its cash C_i and government bonds T_{ia} and apply inflow and outflow rates consistent with those specified under Basel III (see Appendix A.1.2). Under normal times, a bank is expected to have an LCR Λ_i that complies with the LCR buffer standards $\Lambda^S = 100\%$ (BIS (2013)):

$$\Lambda_i := \frac{Q_i}{\Theta_i} = \frac{Q_i}{\Theta_i^O - \min\{\Theta_i^I, 0.75 \cdot \Theta_i^O\}} \geq \Lambda^S = 100\%. \quad (3.5)$$

Asset Managers We extend our model to include four representative asset managers - a bond fund, an equity fund, a mixed fund and an ‘other’ fund - initialised using 2017Q4 aggregate data from the *ECB Statistical Data Warehouse*. The stylised balance sheet of an asset manager $i \in \mathcal{M}$, where \mathcal{M} is the set of asset managers, is shown in Figure 3.2a. The assets A_i of an asset manager consist of cash C_i , tradable assets T_i , and other assets O_i . It has one form of representative liability L_i and also has equity E_i consisting of σ_i number of outstanding shares held by investors.

The critical *constraint* that (open-ended) asset managers face is that they must fulfil redemption requests from investors. Empirical evidence shows that asset managers tend to experience investment inflows or outflows (i.e. redemptions) based on their performance as measured by net asset value (NAV) (see e.g. Coval & Stafford (2007), Baranova et al. (2017)). In line with the empirical evidence of Coval & Stafford (2007), we assume that the asset manager investors redeem shares proportional to the relative loss of their NAV in our simulations. Asset managers have an obligation to pay back these shares at their prevailing NAV - we set out how they do so in section 3.5.3.

Hedge Funds Finally, we add a number of representative hedge funds (\mathcal{H} is the set of hedge funds), which use repo funding from banks to fund their holdings of tradable assets. Because we do not have detailed data on hedge funds, we introduce these institutions in a stylised way. We assume that each hedge fund receives its repo funding from one bank, to the extent that its balance sheet has a financial leverage λ_i (defined as book equity E_i over assets A_i) of 43%³⁵ (based on the FCA (2015) survey).³⁶ A hedge fund’s asset holdings are calibrated to data from the *ECB Statistical Data Warehouse* (see initialisation details in Appendix A.1.2). The stylised balance sheet of a hedge fund is displayed in Figure 3.2b.

³⁵Or, equivalently, 2.3 if leverage is defined as assets A_i over book equity E_i .

³⁶ As explained in more detail in Appendix A.1.2, we do not consider synthetic leverage (that attained by derivatives, for instance).

C_i , Cash	L_i , Liabilities
T_i , Tradable Assets	E_i , Equity
O_i , Other Assets	

(a) Stylised balance sheet of an asset manager $i \in \mathcal{M}$.

C_i , Cash	\tilde{R}_i , Repos
T_i , Tradable Assets	E_i , Equity
O_i , Other Assets	

(b) Stylised balance sheet of a hedge fund $i \in \mathcal{H}$.

Figure 3.2: Balance sheets of non-banks.

A hedge fund's assets A_i are composed of cash C_i , tradable assets T_i , and other assets O_i . Its liabilities L_i are made up of repos \tilde{R}_i and equity E_i .

Hedge funds must meet their contractual obligations – in this case to meet margin calls and repay maturing funding. We also monitor their leverage ratios, and consider how leverage targeting behaviour may impact systemic amplification risks.

3.5.2 Financial Contracts, Markets & Contagion Mechanisms

Our model includes a variety of financial contracts, which are in turn associated with a number of contagion mechanisms that operate on the networks these contracts create. We explicitly model: (i) tradable assets, T_i ; (ii) interbank contracts, I_i and \tilde{I}_i ; and (iii) repurchase agreements, R_i and \tilde{R}_i . We do not explicitly model other assets O_i and \tilde{O}_i , external assets Y_i , and deposits D_i , though shocks can be applied to these assets and liabilities.³⁷

As discussed in Section 3.4, the generic framework allows us to model the market associated with each contractual type. In our model, this would imply (1) modelling price formation in tradable asset markets, (2) modelling how the formation of new repo- and interbank contracts takes place, and (3) modelling how their prices (e.g. their interest rates) are set. Given our emphasis on capturing systemic amplification risk in a stress

³⁷We do not model risks associated with derivatives contracts, largely due to the complexity involved in modelling margin calls and changes in derivatives exposures without granular data on these derivatives contracts. Because we omit derivatives models, our model captures neither liquidity flows associated with margin calls nor the impact on banks' solvency of deteriorating counterparty creditworthiness. Recent work at the Bank of England has found that risks to non-bank financial institutions from derivatives margin calls are currently low (see ?), suggesting that including this missing channel would not significantly impact on our results. By excluding derivatives, we may also miss losses (or gains) from derivatives positions hedging exposures. We also do not explicitly model the availability and cost of long-term unsecured funding. While price increases and restrictions in the availability of long-term funding likely add to the pressures on banks, the impact during the relatively short timescales we focus on (and within which the other mechanisms we focus on in the model that operate) that is expected to play out would be limited.

scenario, we take an approach consistent with the relevant literature (see e.g. [Caccioli et al. \(2013\)](#), [Kok & Montagna \(2013\)](#), [Gai et al. \(2011\)](#)) and do not model contract formation dynamics in the interbank and repo markets; the only markets we model are those for tradable assets. By not allowing for the possibility that institutions that face funding shocks may acquire new funding, our model could overestimate contagion risks associated with such shocks.³⁸

Tradable Assets and Markets The *value* of an institution’s tradable assets T_i is given by

$$T_i = \sum_{a \in \mathcal{A}} T_{ia} = \sum_{a \in \mathcal{A}} \sum_{m=1}^{M_a} T_{iam} = \sum_{a \in \mathcal{A}} \sum_{m=1}^{M_a} s_{iam} p_{am}, \quad (3.6)$$

where T_{ia} is the value of the tradable assets of institution $i \in \mathcal{F}$ of type $a \in \mathcal{A}$, which could be government bonds a_1 , corporate bonds a_2 , equities a_3 , or other tradable assets a_4 . T_{iam} denotes the value of tradable asset m of type $a \in \mathcal{A}$ held by institution $i \in \mathcal{F}$, where $m = 1, \dots, M^a$ and M^a denotes the number of different types of assets of type a . A tradable asset m could be an Italian 10-year government bond, for example. The price of a tradable asset m of type $a \in \mathcal{A}$ is given by p_{am} and the units held by institution $i \in \mathcal{F}$ is given by s_{iam} .

In line with [Cont & Schaanning \(2017\)](#), [Greenwood et al. \(2015\)](#), we do not model the *counterparty* (i.e. the issuer) and the *cash-flows* (e.g. dividends) associated to tradable assets, but focus on the *interconnections* formed by institutions $i \in \mathcal{F}$ that hold an asset m of type $a \in \mathcal{A}$ in common, to enable overlapping portfolio contagion (see Section 3.5.2). The overlapping portfolio network is reconstructed using the random network method employed in [Kok & Montagna \(2013\)](#) (see Appendix A.1.2).

Following the literature on contagion via asset sales (see e.g. [Caccioli et al. \(2015, 2014\)](#)), we take a simplified reduced-form approach to modelling the price impact of asset sales.³⁹ Empirical research, such as by [Bouchaud \(2010\)](#), suggests that the price impact function is concave and is linear for small volumes of sales.

For simplicity, and given that the volume of sales at each time in our model is limited, we assume that the price impact is linear (in line with [Greenwood et al. \(2015\)](#)). Given this approach, the price at time t of asset m of type $a \in \mathcal{A}$, p_{am}^t is given by

$$p_{am}^t = p_{am}^{t_0} \max\{1 - \beta_{am} f_{am}^t; 0\}, \quad (3.7)$$

³⁸We can assess the importance of this assumption by exploiting the flexibility of the framework to ‘turn off’ funding contagion, which would be analogous to assuming that institutions can frictionlessly source new funding in case they face withdrawals.

³⁹A more advanced study of price dynamics could be facilitated by modelling the limit order book in exchange-traded markets ([Paulin et al. \(2018\)](#)). Even though we do not model limit order books here, our generic framework supports such modelling.

and is capped so that it never falls below zero. In equation 3.7, f_{am}^t denotes the cumulative fraction of net asset sales of asset m of type $a \in \mathcal{A}$ (relative to the market capitalisation) up to time t and β_{am}^t is the asset’s price impact parameter. For instance, $\beta_{am} = 0, 1, 2$ means that the price of tradable asset m of type $a \in \mathcal{A}$ falls by $x = 0, 5, 10\%$ if 5% of the total market capitalisation has been sold.

Interbank Contracts In line with [Amini et al. \(2013\)](#), [Hałaj & Kok \(2013\)](#), the value of the interbank assets I_i is given by the sum of the notional exposures to bank counterparties that have not yet defaulted $I_{ij} \mathbb{1}_{\{j \notin \mathcal{D}\}}$, where \mathcal{D} is the set of defaulted banks. If a *counterparty* $j \in \mathcal{B}$ has defaulted, the bank receives $(1 - LGD_j)I_{ij}$ of cash C_i , where LGD_j is the loss given default (LGD). Interbank liabilities \tilde{I}_i are given by the sum of the notional borrowing from other banks I_{ji} , so $\tilde{I}_i = \sum_{j \in \mathcal{B}} I_{ji}$. We reconstruct the interbank network using bank balance sheet data and the reconstruction method proposed by [Hałaj & Kok \(2013\)](#) and employed by [Kok & Montagna \(2013\)](#). This method iteratively picks a random pair of banks and assigns a random number from a uniform distribution between 0% and 20%, which determines what percentage of the bank’s residual interbank assets is deposited in the other bank’s remaining interbank liabilities (truncated if larger). The exposure between any pair of banks is capped at 20%, a choice that is motivated by the Basel III large exposure limits ([BCBS \(2013\)](#)).⁴⁰

Following [Kok & Montagna \(2013\)](#), we assume that interbank contracts are overnight contracts (i.e. mature every day) and are automatically rolled over, unless they are explicitly *not* rolled over, in which case the affected bank receives a *liquidity* shock. An interbank contract may, for example, be withdrawn if a counterparty has liquidity needs ([Aymanns et al. \(2018\)](#)) or engages in liquidity hoarding ([Acharya & Skeie \(2011\)](#), [Heider et al. \(2009\)](#)). We only consider funding reductions to raise cash to meet contractual obligations or regulatory constraints in the model, but it is also possible to incorporate liquidity hoarding (see e.g. [Wiersema et al. \(2019\)](#) for a way to do so).

Repos and Reverse Repos Under a repurchase agreement, an institution j will sell a tradable asset m of type $a \in \mathcal{A}$ to an institution i at a time t and repurchase the security

⁴⁰We note that our results are averaged across N realised interbank and security network reconstructions, and error bars show the standard deviation of the outcome as a result of these different reconstructed networks. Of course, these error bars only reflect the standard deviation that is realised with a particular network reconstruction method. If outcomes were not only averaged across realised networks using a particular reconstruction method, but also across multiple reconstruction methods, reflecting that the ‘functional form’ of the network may be unknown, then the standard deviation could be greater. The network structure matters: it is shown by several studies to significantly impact the probability and extent of contagion (see e.g. [Gai & Kapadia \(2010\)](#)). See [Anand et al. \(2017\)](#) for a horserace of different network reconstruction methods, and see [Gandy & Veraart \(2017\)](#) for an alternative, Bayesian approach to reconstructing the interbank network.

at a time $T > t$ at pre-specified price. In effect, in this transaction institution j provides a loan secured by assets (collateral) to a *counterparty* i . If institution i defaults during the lifetime of the contract, bank j is legally entitled to take the received collateral and may (fire) sell it to recover as much of the notional R_{ji} (or more) as possible. To ensure that enough cash can be recovered upon the sale of the collateral, collateral m of type $a \in \mathcal{A}$ typically receives a haircut h_{am} .

We assume that each individual repo contract is collateralised by one type of non-cash collateral s_{ijam}^e (where s_{ijam}^e denotes that specific asset m of type a of institution i is placed as collateral to institution j and hence remains for accounting purposes as an encumbered ‘e’ asset on i ’s balance sheet) – one of the tradable assets T_i – and that cash collateral C_{ij}^e of institution i can be used to supplement this if necessary.⁴¹ We impose the restriction that when an institution has used an amount of a particular asset to secure repo funding, that asset is no longer available to the institution to liquidate until that repo contract is terminated. As with interbank contracts, we assume that repo contracts are overnight contracts that automatically roll over unless one of the counterparties explicitly opts not to do so.

Whenever the price p_{am} of the asset collateral s_{ijam}^e falls, the value of the collateral after the haircut may no longer be sufficient to fully collateralise the repo loan R_{ji} . In such cases, the institution receives a margin call M_{ji} to restore full collateralisation, which it must meet either with more of the underlying asset or cash collateral. If it has insufficient of either, it liquidates other asset types to obtain cash. The margin call M_{ji}^t from institution $j \in \mathcal{F}$ to institution $i \in \mathcal{F}$ is given by

$$M_{ji}^t := R_{ji}^t - (1 - h_{am}^t) s_{ijam}^{e,t-1} p_{am}^t + C_{ij}^{e,t-1} \quad (3.8)$$

$$= \begin{cases} > 0 & i \in \mathcal{F} \text{ must pledge } M_{ji}^t \text{ value of extra ‘haircutted collateral’ to } j \in \mathcal{F}; \\ = 0 & \text{no margin call;} \\ < 0 & j \in \mathcal{F} \text{ must return } |M_{ji}^t| \text{ value of ‘haircutted collateral’ to } i \in \mathcal{F}. \end{cases}$$

The margin call obliges institution $i \in \mathcal{F}$ to place extra ‘unencumbered’ (u) asset collateral s_{iam}^u or cash collateral C_i^u ⁴² to make equality $R_{ji} = (1 - h_{am}) s_{ijam}^e p_{am} + C_{ij}^e$ hold again. An institution i can only meet a margin call with existing items on its balance sheet if it

⁴¹The superscript ‘e’ signifies that the posted collateral stays for accounting purposes on the balance sheet of institution i , but is an ‘encumbered asset’ in the sense that it is no longer available to the institution to sell while the repo contract is extant.

⁴²We note that tradable assets s_{iam} be further broken down in that part which is unencumbered s_{iam}^u (can be liquidated as no counterparty has a claim on it) and the sum of the encumbered s_{ijam}^e collateral posted to each counterparty $j \in \mathcal{F}$. That is, $s_{iam} = s_{iam}^u + \sum_{j \in \mathcal{F}} s_{ijam}^e$. Likewise, cash can be split in its unencumbered and encumbered part: $C_i = C_i^u + \sum_{j \in \mathcal{F}} C_{ij}^e$.

has sufficient unencumbered assets of the type s_{iam}^u already placed in the repo contract R_{ji} , or if it has sufficient unencumbered liquid instruments C_i^u . Else, it needs to liquidate unencumbered assets of other asset types (e.g. firesale tradable assets $s_{iam}^u > 0$) to raise sufficient cash C_i^u that can be placed as cash collateral. Since an institution i could have multiple repo contracts R_{ji} it may also face multiple margin calls at every time step t , which it meets sequentially. The total *value* of the reverse repos of institution $i \in \mathcal{F}$ is given by $R_i = \sum_{i \in \mathcal{F}} R_{ij}$ and its total repo value is given by $\tilde{R}_i = \sum_{i \in \mathcal{F}} R_{ji}$.

It is common that an institution $i \in \mathcal{F}$ is allowed to re-hypothecate collateral received as part of its reverse repo R_{ij} position by placing it in its own repo contract R_{ki} , for a $j, k \in \mathcal{F}$. When an institution has offsetting reverse repo R_{ij} and repo contracts R_{ki} , its position is called match book. In such case, the margin call associated with a reverse repo contract is opposite to the margin call associated with a repo contract (i.e. $M_{ij} = -M_{ki}$). As a consequence, the institution can just pass on the collateral it received in the reverse repo contract R_{ij} to the repo contract R_{ki} , or the other way around. Hence, the institution is not exposed to liquidity risk unless delays in the delivery of collateral occur (Gorton & Muir (2016)), which we do not capture. In our model (see Section 3.5.1), we assume that each bank $i \in \mathcal{B}$ in its role as an intermediary is largely matched book (as their reverse repo R_i and \tilde{R}_i given by data largely offset in size, see Appendix A.1.2), so are little exposed to margin call risk, whereas each hedge fund $j \in \mathcal{H}$ is not matched book and thus exposed to margin calls M_{ji} .

In line with Bookstaber, Paddrik & Tivnan (2014), we assume that a bank provides reverse repo funding to hedge funds (see details in Appendix A.1.2), and a bank itself receives repo funding from an external financier that is not explicitly modelled. We set the haircuts h_{am}^t (see Section 3.5.2) for government bonds a_1 , corporate bonds a_2 , equities a_3 and other tradable assets a_4 respectively equal to $h_{a_1,m}^t = 2\%$, $h_{a_2,m}^t = 4\%$, $h_{a_3,m}^t = 15\%$, and $h_{a_4,m}^t = 4\%$ (such as mortgage-backed securities), in line with (BCBS (2004)) $\forall t$, for $m = 1, \dots, M^a$. Cash collateral does not receive a haircut. We could but do not consider how haircuts may change (e.g. increase) over time t in periods of distress. The potentially sharp increase in haircuts in financial crises has been empirically examined by Gorton & Metrick (2009) and shown by Brunnermeier & Pedersen (2009) to be an additional driver of margin calls-induced liquidations.

Contagion Mechanisms The contracts discussed above act as edges in a financial network, and are the channels through which contagious shocks can be passed on or amplified. We will discuss four, sometimes interacting, types of contagion that we model explicitly.

Exposure Loss Contagion Exposure loss contagion is a form of node amplification (see Section 3.4.2). It occurs when liquidation following the default of bank $j \in \mathcal{B}$ leads to further contagion, which are induced by the exposure losses incurred by each of its interbank contract I_{ij} counterparties $i \in \mathcal{B}$. By default, we set the loss given default for all banks equal to one hundred percent (i.e. $LGD_i = 100\%$, $\forall i \in \mathcal{B}$), since Cont et al. (2010) argue that over the short time period typically considered by a system-wide stress test counterparties of a defaulted bank are unlikely to have a positive recovery. In traditional models of exposure loss contagion, such as Amini et al. (2013), exposure losses can cause a bank to default, thereby potentially setting in motion a chain of further defaults. In our model, exposure losses may also weaken a bank such that it needs to de-lever to become less vulnerable. Hence, in our model exposure losses may not merely foster default-domino effects through *post-default* contagion, but also spawn *pre-default* contagion in the form of, for example, overlapping portfolio contagion or funding contagion. Exposure losses can also contribute directly to pre-default contagion, because banks mark down the value of their interbank exposures as their counterparties' solvency deteriorates, leading to further deterioration in solvency as modeled by Bardoscia et al. (2017) - but we do not capture that direct channel in our implementation.

Overlapping Portfolio Contagion Overlapping portfolio contagion, another form of node amplification, occurs when the sale of a tradable asset causes the asset price to drop, leading in turn to a downward valuation of marked-to-market assets of other institutions that hold the sold asset as well. Where institutions wish to retain their leverage around some target, this may prompt delevering via asset sales and further price falls (Duarte & Eisenbach (2015)).

Traditionally, many models have generally motivated delevering by using a binding leverage constraint λ_i (see e.g. Cont & Schaanning (2017), Greenwood et al. (2015)), though a forthcoming paper Coen et al. (2019) models asset sales in the face of risk-weighted capital ratio and liquidity coverage ratios. Our model includes these constraints and also allows for asset sales to be triggered by contractual obligations (such as the obligation to pay back a loan or meet a margin call). Importantly, in our model banks have options other than selling assets in order to delever; they can also reduce interbank funding I_i or reverse-repo funding R_i exposures, for example. Therefore, institutions affected by overlapping portfolio contagion will not necessarily transmit and amplify marked-to-market shocks even if they are forced to act, but may instead trigger funding contagion or collateral contagion. Where marked-to-market shocks are sufficient to cause institutional default and liquidation, this can also trigger exposure loss contagion.

Collateral Contagion As noted, when falls in asset prices lead to margin calls on repo contracts that cannot be met with available cash or collateral, institutions will be forced to liquidate other assets in order to raise cash collateral instead. We call this ‘collateral contagion’, which is a form of edge amplification 3.4.2. Hence, overlapping portfolio contagion and collateral contagion can reinforce each other, similar to the margin-price spiral Brunnermeier & Pedersen (2009) has identified.

Funding Contagion Funding contagion, which is a form of node amplification, occurs when a funding shock provokes the institution to raise liquidity by withdrawing funding from its counterparties. Our model allows banks to cease rolling over funding to each other, either to raise cash to meet contractual obligations or to deleverage. However, in the face of funding shocks banks are not limited to taking such actions, and could for example also raise cash by selling securities.

3.5.3 Behaviour

As noted before, when financial institutions act in our model their option set may be limited by various constraints. However, within that option set there may still be ample choice. The way in which institutions make that choice is based on behavioural assumptions. In this section, we discuss these two elements of behaviour. Specifically, we first discuss how, at a conceptual level, banks in our model attempt to stay away from their binding regulatory constraints. Subsequently, we outline how we operationalise this conceptual approach in our stress testing model. We then discuss how banks act in situations when multiple constraints bind simultaneously, and also discuss the ‘pecking orders’ banks use to decide how to meet their constraints. Finally, we turn to asset managers and hedge funds, and describe how we model their behaviour.⁴³

We again stress that our generic framework is flexible by design, so that we could for example also consider the impact of different or more heterogeneous behavioural assumptions. In the stress testing model institutions of the same type (bank, asset manager, hedge fund) use the same set of behavioural rules, but they still act heterogeneously at any point in time because of their institution-specific combination of binding constraints and balance sheet properties.

⁴³Our model does not currently consider (1) a number of behavioural options available to banks (e.g. dividend cuts), (2) strategic interactions, or (3) endogenous intervention by central banks (e.g. lender of last resort facilities). It also does not allow for buying behaviour, which might act to dampen, for example, the price impacts of fire sales. However, our generic framework and model implementation does allow us to experiment - by including (some of) these behavioural options as assumptions in the model - and to evaluate how important they are to model outcomes.

Behaviour of Banks: Minima, Buffers and Targets We assume that banks choose to maintain a ‘management buffer’ at all times. Once banks fall below this buffer value, they respond to get back to some target (this is in line with the approach taken in e.g. [Bookstaber, Paddrik & Tivnan \(2014\)](#), [Cont & Schaanning \(2017\)](#)). Such behaviour is consistent with the empirical findings of [Adrian & Shin \(2010\)](#) and also has intuitive appeal: institutions have to monitor and adjust their balance sheets to comply with regulatory standards and are likely to take some buffer space to prevent them dipping below regulatory requirements too quickly.

Under this setup, banks act when at least one of the following conditions holds: (1) $\lambda^M \leq \lambda_i < \lambda_i^B \leq \lambda_i^T$; (2) $\rho^M \leq \rho_i < \rho_i^B \leq \rho_i^T$; and/or (3) $\Lambda_i < \Lambda_i^B$ (where Λ_i^B may be above, below, or equal to Λ^S). The superscript B and T denote the buffer value and target value of the constraint, which are bank-specific.

A bank determines the level of its management buffer and target at least in part based on its assessment of the usability of its regulatory buffer (see discussion in Sections [3.5.1](#) and [3.7](#)). For example, if banks consider their regulatory buffers to be fully usable, they may set their management buffer and target at a level that is lower than the regulatory buffer level, enabling them to use the buffer to absorb shocks. On the other hand, when banks behave as if buffers are not usable (in other words, if they behave as if buffers are requirements), their management buffer and target may exceed the regulatory requirements, to make sure that some shocks can be absorbed before these regulatory requirements are triggered.

Usability of Regulatory Buffers and Targets To assess the impact of regulatory buffer usability on systemic risk (see Section [3.6.5](#)), we introduce a parameter u that determines what fraction of each regulatory buffer (risk-based capital buffer, leverage buffer, and liquid asset buffer) a bank is willing to use. So if we set u to 25%, then a bank will seek to prevent its capital ratio falling below 75% of their regulatory buffer standard by taking actions to rebuild their capital ratio towards a desired target value. Alternatively, if we set u to 0%, regulatory buffers are not considered to be usable at all, and banks will take actions to avoid dipping below them. Given a usability u of buffers, the buffer value at which institutions act to return to target for each constraint is given by

$$\rho^B = \rho^M + (1 - u)(y^\rho \rho_i^{CB}), \quad (3.9)$$

$$\lambda^B = \lambda^M + (1 - u)(y^\lambda \rho_i^{CB}), \quad (3.10)$$

$$\Lambda^B = (1 - u)\Lambda^S, \quad (3.11)$$

where y^λ and y^ρ tell the size of the regulatory leverage buffer and regulatory risk-weighted buffer relative to the Basel III standard (e.g. $y^\rho = 2$ means regulators have doubled the risk-weighted buffer standard applicable to each bank relative to the Basel III standard). The default parameters for usability (u), size of the regulatory buffer standards (y^ρ , y^λ), and the target values (ρ_i^T , λ_i^T and Λ_i^T) are shown in Table 3.2.

Banks' Actions When Facing Multiple Constraints If necessary, banks can choose from a set of actions to rebuild capital and liquidity ratios and meet repayment obligations. The action that is most effective will likely depend on the constraint that binds. Our model first sets out the set of actions available to banks to meet each constraint independently, which are compared against the pecking order we impose on these actions. This process yields multiple pecking orders that banks use to meet their constraints.

We consider the following pecking orders:

Margin calls - An institution i meets a margin call by first attempting to post more cash or unencumbered ('u') assets s_{iam}^u of the type m, a underlying the repo contract. If that is not possible, it will raise cash by liquidating other types of unencumbered assets (including interbank contracts and reverse repo lending) in proportion to its current holdings. We will now describe this mathematically.

We start with explaining the case where the margin call is positive (i.e. $M_{ji}^t > 0$). As explained in Section 3.5.2, a margin call may be met with the same type of collateral that is already placed as part of the contract, or, if that is not sufficient with cash collateral. In this light, an institution $i \in \mathcal{F}$ will meet a positive margin call $M_{ji}^t > 0$ by pledging $s_{ijam}^{e,t,E}$ extra (E) units of collateral m of type $a \in \mathcal{A}$ to institution $j \in \mathcal{F}$ at time t . That is, $s_{ijam}^{e,t,E}$ is given by

$$s_{ijam}^{e,t,E} = \min\left\{\frac{M_{ji}^t}{(1 - h_{am}^t)p_{am}^t}; s_{iam}^{u,t-1}\right\} \mathbb{1}_{\{M_{ji}^t > 0\}}, \quad (3.12)$$

where we note that the units $s_{ijam}^{e,t,E}$ pledged can never exceed the units of unencumbered collateral m of type $a \in \mathcal{A}$ that institution $i \in \mathcal{F}$ has of type a, m , $s_{iam}^{u,t-1}$. If the units of pledged collateral $s_{ijam}^{e,t,E}$ are not sufficient to fully meet the margin call, then institution $i \in \mathcal{F}$ has to pledge $C_{ij}^{e,t,E}$ extra cash collateral, given by

$$C_{ij}^{e,t,E} = \min\{\max\{M_{ji}^t - s_{ijam}^{e,t,E}(1 - h_{am}^t)p_{am}^t; 0\}; C_i^{u,t-1}\} \mathbb{1}_{\{M_{ji}^t > 0\}}, \quad (3.13)$$

where we note that $C_{ij}^{e,t,E}$ can never exceed the amount of unencumbered cash $C_i^{u,t-1}$ that institution $i \in \mathcal{F}$ has. If at this point institution $i \in \mathcal{F}$ has still not fully satisfied its margin call M_{ji}^t , then it has to resort to liquidating assets (see also e.g. [Gai et al. \(2011\)](#)). The amount of assets institution $i \in \mathcal{F}$ has to liquidate l_i^t to meet the remainder of the margin call is given by

$$l_i^t = \max\{M_{ji}^t - s_{ijam}^{e,t,E}(1 - h_{am}^t)p_{am}^t - C_{ij}^{e,t,E}, 0\} \mathbb{1}_{\{M_{ji}^t > 0\}}.^{44} \quad (3.14)$$

It liquidates assets according to the ‘margin call pecking order’ described in Section ??.

If the amount of cash that institution $i \in \mathcal{F}$ raises from liquidating assets is still not sufficient to honour its margin call M_{ji}^t , then it defaults. In such case, the reverse repo party $j \in \mathcal{F}$ is contractually allowed to permanently keep all the collateral ($s_{ijam}^{e,t}$ and $C_{ij}^{e,t}$) in the repurchase agreement R_{ji}^t (see Section 3.5.2). We assume that institution $i \in \mathcal{F}$ will (fire) sell the non-cash collateral (i.e. $s_{ijam}^{e,t}$) to eliminate any exposure to the collateral ([Shleifer & Vishny \(2011\)](#)), which raises cash.

If, on the other hand, the margin call is negative (i.e. $M_{ji}^t < 0$), then the reverse repo party $j \in \mathcal{F}$ must return part of the collateral it has received from the repo party $i \in \mathcal{F}$. It must return some collateral, because the repo contract R_{ji}^t is now overcollateralised given the haircuts h_{am}^t that apply and given the current price of the collateral p_{am}^t . We assume that the reverse repo party $j \in \mathcal{F}$ first returns (R) $C_{ji}^{e,t,R}$ amount of cash collateral it received from the repo party $i \in \mathcal{F}$, given by

$$C_{ji}^{e,t,R} = \min\{C_{ij}^{e,t-1}; |M_{ji}^t|\} \mathbb{1}_{\{M_{ji}^t < 0\}}. \quad (3.15)$$

Subsequently, if that is not enough, the reverse repo party $j \in \mathcal{F}$ returns $s_{ijam}^{e,t,R}$ units of non-cash collateral received from the repo party $i \in \mathcal{F}$, given by

$$s_{ijam}^{e,t,R} = \frac{\max\{|M_{ji}^t| - C_{ij}^{e,t,R}, 0\}}{1 - h_{am}^t} \mathbb{1}_{\{M_{ji}^t < 0\}}. \quad (3.16)$$

Repaying maturing liabilities - A bank initially meets payment obligations with cash. If it has insufficient cash, it will raise cash by liquidating assets in the following order: (1) interbank contracts I_i ; (2) reverse repos R_i ; (3) unencumbered tradable assets T_i^u (starting with the tradable assets that have the least price impact).

⁴⁴Institution $i \in \mathcal{F}$ may also slightly liquidate more assets than l_i^t to take any potential liquidation cost into account.

Defending the risk-weighted capital ratio - Banks strengthen their risk-weighted capital ratio ρ_i by liquidating assets with the highest risk-weight first, in order to raise cash with a zero risk weight. We will now describe this mathematically.

A bank returns to a target for the risk-weighted capital ratio ρ_i^T whenever its risk-weighted capital ratio ρ_i (defined in equation 3.1) falls below its buffer ρ_i^B and has not failed yet (i.e. $\rho_i \geq \rho^M$). A bank $i \in \mathcal{B}$ returns to its target ratio ρ_i^T by reducing non-zero risk-weight assets A_{ip} (for $\omega_p \neq 0$, for $p \in \mathcal{P}$). As noted above, we assume that the bank returns to target by reducing the most high risk-weight assets first, as this is the most effective way to quickly get back to the capital ratio target ρ_i^T .⁴⁵ Given the risk weights that apply, the order to reduce non-zero risk-weighted assets A_{ip} is given by: (1) unencumbered corporate bonds $T_{a_2}^u$; (2) unencumbered other tradable assets $T_{a_4}^u$; (3) unencumbered equities $T_{a_3}^u$; (4) interbank assets I_i ; (5) reverse repo R_i .

The iterative method employed by bank $i \in \mathcal{B}$ to aim to reach its target ρ_i^T is as follows. It liquidates \hat{r}_{ip_4} amount of asset type A_{ip_4} . It can never reduce more assets than the unencumbered assets $A_{ip_4}^u$ it has of this type. That is, \hat{r}_{ip_4} is given by $\hat{r}_{ip_4} = \min\{r_{ip_4}, A_{ip_4}^u\}$, where r_{ip_4} is given by

$$r_{ip_4} = \frac{1}{\omega_{p_4}} \left[\sum_{p \in \mathcal{P}} \omega_p A_{ip} - \frac{\tilde{E}_i}{\rho_i^T} \right], \quad (3.17)$$

and follows from

$$\rho_i^T = \frac{\tilde{E}_i}{\omega_{p_4} (A_{ip_4} - r_{ip_4}) + \sum_{p \in \mathcal{P} \setminus p_4} \omega_p A_{ip}}. \quad (3.18)$$

If $\hat{r}_{ip_4} < r_{ip_4}$ then bank $i \in \mathcal{B}$ did not have enough unencumbered assets $A_{ip_4}^u$ of type p_4 to reach its target ρ_i^T . Hence, it will next reduce \hat{r}_{ip_6} amount of the next asset in the pecking order A_{ip_6} . Where \hat{r}_{ip_6} is again given by $\hat{r}_{ip_6} = \min\{r_{ip_6}, A_{ip_6}\}$ and r_{ip_6} is given by

$$r_{ip_6} = \frac{1}{\omega_{p_6}} \left[\sum_{p \in \mathcal{P}} \omega_p A_{ip} - \sum_{p=p_4} \omega_p A_{ip}^u - \frac{\tilde{E}_i}{\rho_i^T} \right] \mathbb{1}_{\{\hat{r}_{ip_4} < r_{ip_4}\}}. \quad (3.19)$$

We observe that the amount of (unencumbered) assets that have been designated to be liquidated in the previous round of the iterative procedure have been reduced from the sum. We continue this iteration for as many times as its needed, by extending this logic,

⁴⁵We note however that banks may implement optimisation strategies to minimise liquidation losses that may result in them selling more liquid assets in preference to less liquid assets, as in [Coen et al. \(2019\)](#).

to reach the target ρ_i^T up to the last non-zero risk weight that can be reduced by at most $\hat{r}_{ip_8} = \min\{r_{ip_8}, A_{ip_8}\}$, where r_{ip_8} is given by

$$r_{ip_8} = \frac{1}{\omega_{p_8}} \left[\sum_{p \in \mathcal{P}} \omega_p A_{ip} - \sum_{p=p_4, p_6, p_5, p_7} \omega_p A_{ip}^u - \frac{\tilde{E}_i}{\rho_i^T} \right] \mathbb{1}_{\{\hat{r}_{ip_x} < r_{ip_x}, \text{ for } x = 4, 6, 5, 7\}}. \quad (3.20)$$

In case the following condition is true the bank $i \in \mathcal{B}$ cannot fully reach its target, even in the absence of liquidation cost

$$\frac{\tilde{E}_i}{\sum_{p \in \mathcal{P}} \omega_p A_{ip} - \sum_{p=p_4, p_6, p_5, p_7} \omega_p A_{ip}^u} < \rho_i^T. \quad (3.21)$$

Defending the leverage ratio - Banks first delever by using cash to proportionally pay back liabilities L_i .⁴⁶ Where this is insufficient, we assume that banks liquidate assets in order of increasing liquidation costs. Therefore, the pecking order for liquidation is the same as the one used when meeting payment obligations. We will now describe this mathematically.

A bank $i \in \mathcal{B}$ returns to its leverage target λ_i^T whenever its leverage ratio λ_i (defined in equation 3.2 as the bank's CET1 equity \tilde{E}_i ⁴⁷ over its asset exposure \hat{A}_i ⁴⁸) falls below its buffer value λ_i^B and it has not defaulted yet (i.e. $\lambda_i \geq \lambda^M$). A bank $i \in \mathcal{B}$ returns to its leverage target λ_i^T by delevering d_i amount, rather than issuing new equity to uplift the leverage ratio λ_i as issuing equity is typically not feasible in times of distress (Greenwood et al. (2015)). The delevering amount d_i is given by

$$d_i = \left[\hat{A}_i \frac{1}{\lambda_i^T} - \tilde{E}_i \right] \mathbb{1}_{\{\lambda^M < \lambda_i \leq \lambda_i^B\}}, \quad (3.22)$$

⁴⁶We note that in the UK, central bank reserves do not contribute to the leverage ratio, so this option strictly speaking would not be available.

⁴⁷In the stress test we would like to capture how asset losses and liability changes effect the value of the CET1 equity \tilde{E}_i . To be able to do this, we approximate the CET1 equity of a bank i at time t as $\tilde{E}_i^t \approx E_i^t - \Delta_i^{t0}$, where Δ_i^{t0} is given by the difference between book equity E_i and CET1 equity \tilde{E}_i at time zero. That is, $\Delta_i^{t0} := E_i^{t0} - \tilde{E}_i^{t0}$. This approximation is reasonable. The CET1 equity \tilde{E}_i of a bank strongly relates to the book equity of a bank $E_i := A_i - L_i$, but may not be equal to it. With this approximation, we basically assume that the difference between the equity E_i and the CET1 equity \tilde{E}_i is constant over time.

⁴⁸As we do not have data to determine how the leverage exposure \hat{A}_i changes as a function of asset value changes A_i , we approximate the leverage exposure \hat{A}_i^t at time t as the asset value A_i^t at time t minus some fixed adjustment $\hat{\Delta}_i^{t0}$ determined at time zero. That is, $\hat{A}_i^t \approx A_i^t - \hat{\Delta}_i^{t0}$. We compute $\hat{\Delta}_i^{t0}$ at time zero (i.e before we shock the system) as the difference between the total assets A_i^{t0} and leverage exposure \hat{A}_i^{t0} at time zero (i.e. $\hat{\Delta}_i^{t0} := A_i^{t0} - \hat{A}_i^{t0}$) and keep it constant throughout the stress test. The leverage exposure at time zero \hat{A}_i^{t0} is given by data as $\hat{A}_i^{t0} = \frac{\tilde{E}_i^{t0}}{\lambda_i^{t0}}$.

which follows from

$$\lambda_i^T = \frac{(A_i - d_i) - (L_i - d_i) - \Delta_i^{t_0}}{A_i - d_i} \mathbb{1}_{\{\lambda^M < \lambda_i \leq \lambda_i^B\}} = \frac{\tilde{E}_i^t}{A_i - d_i} \mathbb{1}_{\{\lambda^M < \lambda_i \leq \lambda_i^B\}}. \quad (3.23)$$

A bank $i \in \mathcal{B}$ delevers d_i amount by liquidating d_i amount of assets.

Defending the Liquidity Coverage Ratio - we focus on the numerator of the LCR, and assume that banks boost their LCR Λ_i by liquidating non-HQLA assets and raising cash to increase their holdings of high quality liquid assets, starting by liquidating the least costly assets. We will now describe this mathematically.

A bank $i \in \mathcal{B}$ returns to a LCR target Λ_i^T if its LCR Λ_i (defined in equation 3.5) falls below its buffer Λ_i^B . We assume it does so by reducing non-HQLA (i.e. non- Q_i) assets to generate cash C_i^u , which counts towards its HQLA Q_i (i.e. the denominator of the LCR Λ_i), rather than reducing net outflows Θ_i (Θ_i is the denominator of the LCR Λ_i). The amount of non-HQLA assets q_i that a bank $i \in \mathcal{B}$ will liquidate to return to its target is given by

$$q_i = \{\Lambda_i^T \Theta_i - Q_i\} \mathbb{1}_{\{\Lambda_i < \Lambda_i^B\}}, \quad (3.24)$$

which follows from

$$\Lambda_i^T = \frac{Q_i + q_i}{\Theta_i} \mathbb{1}_{\{\Lambda_i < \Lambda_i^B\}}. \quad (3.25)$$

A bank $i \in \mathcal{B}$ liquidates q_i amount of non-HQLA assets.

Multiple constraints that bind simultaneously - If multiple constraints bind simultaneously, we assume that banks prioritise meeting these constraints as follows:

1. Meet contractual obligations (i.e. repayment obligations and margin calls);
2. Improve the risk-weighted capital ratio ρ_i ;
3. Improve the leverage ratio λ_i by paying back liabilities with cash, liquidating further assets if necessary
4. Boost the liquidity coverage ratio Λ_i .

We motivate this order with reference to the observation that contractual constraints are commonly more strictly enforced (which would lead to default) than regulatory constraints (FED (2010), BIS (2014), Brown & Dinç (2011)). When both the leverage ratio and the risk-weighted capital ratio bind, we assume that banks will first take action to improve the risk-weighted capital ratio ρ_i before acting to alleviate the leverage constraint λ_i . We justify this assumption on the basis that re-building the risk-weighted capital ratio ρ_i (by liquidating non-zero risk-weighted assets) raises cash, which the bank can subsequently use to delever should this be necessary. Actions taken with the primary aim of reducing the leverage ratio, however, may have no impact on the risk-weighted capital ratio (for example if zero risk-weighted assets are liquidated and the cash used to pay off liabilities).

In our model, banks address their LCR Λ_i last for a number of reasons. First, the LCR Λ_i is the regulatory constraint that should be minimal (at least in theory) because it is a buffer rather than a regulatory minimum requirement. Second, taking actions in our model to improve the risk-weighted capital ratio ρ_i will in general also boost the LCR Λ_i : if a bank liquidates assets with a non-zero risk weight this will raise cash C_i^u which, insofar it is not used to delever, will increase the numerator of the LCR Λ_i (the HQLA Q_i , see equation 3.5). This reasoning does not always hold. For example, if a bank decides to reallocate investments from assets with a high risk-weight to assets with a low risk-weight that does not count as a HQLA Q_i , improving the risk-weighted capital ratio may not increase a bank's LCR Λ_i . When the actions banks take to improve their risk-weighted capital ratio and to delever have the net effect to push up the LCR Λ_i , they need to take fewer additional actions to return the LCR to its target ratio Λ_i^T . Moreover, further action may not be necessary at all if these actions push up the LCR Λ_i from $\Lambda_i < \Lambda_i^B$ to $\Lambda_i \geq \Lambda_i^B$.

Bank Failure Contagion models often assume that banks are liquidated on default (see e.g. Kok & Montagna (2013), Caccioli et al. (2013)). But resolution frameworks in most jurisdictions have been undergoing rapid changes since the 2007-2008 crisis to enable the orderly resolution of banks (Armour et al. (2016)). We consider two edge cases for what happens when banks ‘fail’, which we term *disorderly liquidation* and *contagion-free resolution*.⁴⁹

In the edge case of *disorderly liquidation* all banks that fail are rapidly liquidated: tradable assets are fire-sold and short-term secured and unsecured loans are withdrawn (in line with Kok & Montagna (2013)), and unsecured creditors take losses (see Section 3.5.2 and 3.5.2) while secured creditors take title to repo collateral. In the other edge case,

⁴⁹We define failure as a breach of minimum capital requirements or illiquidity.

contagion-free resolution, every defaulted bank is resolved without any contagion: the bank simply becomes inactive.⁵⁰ In reality, the consequences of bank failure would be between these two extremes (see e.g. BoE (2017a), Klimek et al. (2015), Hüser et al. (2017), Chennells & Wingfield (2015)). We do not study the impact of specific resolution regimes here. What matters for our purposes is that our qualitative findings hold for both edge cases, suggesting they apply across a broad range of outcomes.

Asset manager behaviour Each asset manager $i \in \mathcal{M}$ has the obligation to pay back shares at their prevailing net asset value when they are redeemed by investors. To do so, an asset manager first use cash. If this is insufficient to meet redemptions, it liquidates tradable assets in a ‘vertical slice’ (i.e. proportional to their asset holdings).

The NAV of an asset manager $i \in \mathcal{M}$ (see Section 3.5.1 for a balance sheet description) is given by

$$\eta_i = \frac{A_i - L_i}{\sigma_i} = \frac{E_i}{\sigma_i}. \quad (3.26)$$

The performance of an asset manager $i \in \mathcal{M}$ in terms of its NAV η_i^t at time t can be measured relative to a reference time point, which we take to be the beginning of the stress test t_0 . We can define the relative loss χ_i^t at time t of the NAV η_i of the representative asset manager $i \in \mathcal{M}$ as

$$\chi_i^t = \frac{\eta_i^{t_0} - \eta_i^t}{\eta_i^t}. \quad (3.27)$$

In line with the empirical evidence of Coval & Stafford (2007), we simply assume that the asset manager investors redeem shares proportional to the relative loss of their NAV χ_i^t .⁵¹ That is, the cumulative fraction of the original number of asset managers shares $\sigma_i^{t_0}$ that is withdrawn up to time t , f_i^t is given by

$$f_i^t = \chi_i^t. \quad (3.28)$$

The asset manager has the obligation to pay back the shares that are redeemed at their prevailing NAV η_i . If the asset manager does not have enough cash C_i to do so it must liquidate tradable asset T_i . This can give rise to a contagious fire sales.

⁵⁰This implementation of *contagion-free resolution* does not reflect our assessment of current resolution regimes. We simply want to capture the edge case where resolution is contagion-free, in order to study the impact on systemic risk.

⁵¹We model the redemptions pressures in a simplistic way to be able to make a point that asset managers can affect banking sector stability (see Section 3.6.3), but given our stress test framework we could easily implement asset managers in a more involved way (e.g. along the lines of Baranova et al. (2017)).

Hedge fund behaviour The behaviour of a hedge fund $i \in \mathcal{H}$ is driven (1) the obligation to meet margin calls (see the specification of how this is done earlier in this section). (2) the obligation to repay a withdrawn repo agreement, and (3) an internal-risk limit to remain below a leverage bound (if it exceeds that limit, it would be forced to delever). In each of these cases, the hedge fund employs a pecking order to determine its response. To meet a margin call M_{ji} , a hedge fund acts in the same way that banks do (see Section 3.5.3): it first pledges more unencumbered collateral of the type already placed, then places unencumbered cash, and finally proportionally liquidates unencumbered assets. To repay under a repo contract R_{ji} , a hedge fund proportionally liquidates unencumbered assets. Similar to banks, a hedge fund delevers whenever its leverage ratio $\lambda_i = \frac{E_i}{A_i}$ ⁵² falls underneath its buffer value λ_i^B to return to its leverage target λ_i^T .⁵³ As explained in Bookstaber, Paddrik & Tivnan (2014), each hedge fund faces an implicit minimum leverage ratio λ_i^M (which is specific to that hedge fund) implied by the haircuts it faces on its collateral. When a hedge fund liquidates unencumbered assets to raise cash in order to pay back liabilities, it does so proportionally.

3.6 Policy experiments and Results

For the same initial shock, we compare the outputs from our system-wide stress testing model for the European financial system to those from a microprudential stress test. As a microprudential baseline, we use the European Banking Authority (EBA) 2018 stress test results.⁵⁴ The EBA stress test was conducted with static balance sheets and did not model second-round effects that arose as a consequence of banks' responses (Ebner (2018)), despite the fact that surveys involving participants in previous EBA stress tests have suggested that these effects could be sizable (Brinkhoff et al. (2018)).⁵⁵ The process works as follows: following an initial shock, the EBA stress test calculates the initial impact on each individual bank and provides a microprudential output. Using this output as a baseline, we use the input from the EBA's microprudential stress test as an *input* for our system-wide stress testing model. On the basis of this input, our system-wide model then accounts for potential contagion mechanisms that might amplify this initial shock.

⁵²Note the leverage ratio λ_i of hedge funds has a different definition from the leverage ratio of banks (see equation 3.2).

⁵³The default parameters for the excess buffer above the hedge fund's minimum (i.e. $\lambda_i^B - \lambda_i^M$) and the excess target above the hedge fund's buffer (i.e. $\lambda_i^T - \lambda_i^B$) are given in Table 3.2.

⁵⁴See the 2018 EBA microprudential stress test outputs here: <https://eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing/2018/results>.

⁵⁵Participants are allowed to cut dividends in response to the impact of the stress under certain conditions, and restrictions on distributions associated with entering regulatory capital buffers are also included. Cost-cutting is, however, constrained.

To evaluate the relevance of the system-wide stress test, we compare its output to the microprudential output produced by the EBA’s stress test.

Our results suggest that the inclusion of contagious dynamics among banks leads to a starkly different (and, generally, more severe) risk outlook on banks’ resilience. In addition, we show that the efficacy of the resolution process for failing banks is an important driver of systemic risk amplification. We then move from this macroprudential but bank-centered analysis to a truly system-wide perspective by including (representative) non-bank financial institutions, and show that this leads to further - albeit limited⁵⁶ - shock amplification. We also evaluate how different contagion mechanisms can reinforce each other such that their combination has a greater impact than the sum of its parts. Finally, we show how the usability and size of regulatory buffers are crucial determinants of the level of systemic amplification risk.

3.6.1 Default Parameters and Visualisation

Unless otherwise stated, all experiments use the same default parameters, which are given in Table 3.2.⁵⁷ Our experiments and their results focus on the magnitude of systemic risk amplification as a function of (1) the severity of the initial shock, (2) the calibration of our price impact parameters, and (3) whether or not banks are liquidated in an orderly way when they fail.

We present the results in a consistent format. On the *y-axis*, we represent a *commonly used systemic risk measure*, \mathbb{E} - an approach that is similar to that used in Gai & Kapadia (2010), Gai et al. (2011), Paulin et al. (2018). This measure is defined by the average fraction of bank defaults in a systemic event (‘the average extent of a systemic event’), with a ‘systemic event’ being defined as a situation where at least 5% of the banking system defaults (see Appendix A.1.3 for a precise definition of \mathbb{E}). In line with Paulin et al. (2018), this average is computed across $N=100$ simulation runs, where in each run the reconstructed interbank and tradable networks are randomly redrawn (see Section 3.5.2 and 3.5.2 for the reconstruction methods).⁵⁸ Given the uncertainty about the network structure - as noted, we do not have access to the data required to calibrate the model to the network structure - the randomness in the system (for each simulation run) stems solely from the randomness in the reconstructed networks (as in Gai & Kapadia (2010)); per x-y-axis point all else is kept constant.

⁵⁶This may be explained by the lack of granular modelling of the non-banks due to data limitations.

⁵⁷Each timestep in this model can be thought of as representing a timeframe of about a day to a few days. Asset sales and actions to stop rolling over repo and interbank contracts can thus be taken and completed within each timestep. Simulations generally converge after a handful of timesteps, with the longest of those presented here taking around 20 timesteps to converge.

⁵⁸Error bars are omitted since they are very small for $N = 100$ and tend to zero for an increasingly large N .

For each experiment, we show the systemic risk assessment \mathbb{E} resulting from both the system-wide stress test (coloured lines) and the microprudential stress test (grey-coloured lines⁵⁹). Since the microprudential stress test does not capture contagion defaults, the grey-coloured lines could be seen as the average fraction of initial bank defaults in a systemic event \mathbb{E} , while the coloured lines represent the total default fraction as defined above. Their difference represents the average fraction of contagion defaults in a systemic event (‘the average extent of contagion’). For simplicity, we will frequently refer to \mathbb{E} as displaying systemic risk (or, conversely, financial stability) or the (initial) fraction of bank defaults.

To highlight the sensitivity of financial stability to the *severity of the initial adverse scenario* and *market liquidity*, we vary the magnitude of the initial shock x on the x -axis. We do so by applying a scalar of between 0 and 2 to the losses from the 2018 EBA stress test, for which $x = 1$. In addition, we vary the price impact by between 0% and 10% if 5% of the market capitalisation of the asset has been sold (see section 3.5.2 for details). Of course, as is common for models of contagion dynamics,⁶⁰ the *magnitude* of systemic losses generated in these experiments are sensitive to a number of assumptions and parameters. Accordingly, and in line with use-cases of models using similar techniques (see e.g. ??), our (current) system-wide stress testing model is not designed to provide highly precise *quantitative* predictions, but instead provides qualitative findings. Importantly, our qualitative findings are robust to varying assumptions and parameters.

⁵⁹The coloured markers on top of a grey line indicate the coloured line, which the grey line is associated to. If the microprudential stress test outcome is the same for the different coloured lines, then the line is just displayed in grey.

⁶⁰Cont & Schaanning (2017), for example, demonstrate the significant sensitivity of systemic outcomes to price impact parameters in their model of price-mediated contagion via asset sales.

Table 3.2: Default settings for the figures in the result Section 4.7.

Parameter Category	Default settings	Brief description and motivation
Initial shock	$x = 1$	Severity of initial shock of the risk-weighted capital and leverage ratio relative to the 2018 European Banking Authority (EBA) adverse scenario. Hence, $x = 1$ means that adverse scenario of the system-wide stress test matches that of the 2018 EBA microprudential stress test.
Institutions	Banks turned on. Hedge funds & asset managers turned off.	This choice is motivated by data quality. Since our initialisation of asset managers and hedge funds is rough (based on <i>ECB Statistical Warehouse Data</i>), we by default exclude them from our model ('turn them off').
Contracts and contagion channels	Overlapping portfolio contagion, funding contagion, exposure loss contagion & collateral contagion turned on.	We include ('turn on') all relevant contagion channels, because modelling a subset of contagion channels may lead to an underestimation of systemic risk (see e.g. Kok & Montagna (2013) , Caccioli et al. (2013)).
Constraints	$\rho^M = 4.5\%$, $y^\rho \rho_i^{CB}$ $\lambda^M = 3\%$, $y^\lambda \lambda_i^{CB}$ $\Lambda^S = 100\%$ where $y^\rho = y^\lambda = 1$	The regulatory capital requirements, and capital and liquidity buffer standards are set in line with Basel III. The buffer standards are set at $y^\rho = y^\lambda = 1$ times the Basel III standard (i.e. equal to the Basel III standard).
	$\Delta_i^{\rho,t_0} = \Delta_i^{\rho,data}$ $\Delta_i^{\lambda,t_0} = \Delta_i^{\lambda,data}$	We assume that if regulatory capital buffer sizes are adjusted relative to the Basel III standard, banks alter their capital ratios by an equal percentage in order to comply with the new regulatory standard.
Market	Asset price fall is $x = 5\%$ if 5% of the market capitalisation has been sold.	This is in line with a standard assumption in the literature, see e.g. Schnabel & Shin (2004) , Cifuentes et al. (2005) , Gai & Kapadia (2010) , and Caccioli et al. (2014) .
Behaviour	$\rho_i^B = \rho_i^M + (1 - u)(y\rho_i^{CB})$, $\lambda_i^B = \lambda_i^M + (1 - u)(y\lambda_i^{CB})$ $\Lambda_i^B = (1 - u)\Lambda_i^S$ where $u = 50\%$	Banks act to return to target whenever they have exhausted $u = 50\%$ of their regulatory capital or liquidity buffers.
	$\rho_i^T = \rho_i^B + 1\%$, $\lambda_i^T = \lambda_i^B + 1\%$, $\Lambda_i^T = \Lambda_i^B + 5\%$	The target value is 1% above the capital buffers, and 5% above the liquidity buffer.

3.6.2 From Micro to Macro: A Macroprudential Overlay for the EBA 2018 Stress Test

We first study how the systemic-risk assessment of system-wide and microprudential stress tests differ, for different levels of severity of the initial, adverse scenario (as given by the scaled 2018 EBA scenario). Following the impact of the initial shock on their balance sheets, banks take actions to return to their targets. These actions, though individually rational, collectively create contagion, amplifying the initial shock.

Figure 3.3 (default parameters) and 3.4 (triple leverage buffers λ_i^{CB}) show the as-

assessment of systemic risk by the EBA’s microprudential stress test (grey lines) and our macroprudential stress test (orange lines) as a function of the initial shock. Three key findings emerge: 1) microprudential stress tests alone are insufficient to assess financial stability; 2) whether banks fail in a disorderly or managed way has a significant impact on financial stability; and 3) when we treat the leverage ratio at the time of the EBA stress test as a binding constraint, it produces greater financial instability than when we consider the risk-based capital ratio alone.⁶¹




Figure 3.3: This figure shows systemic risk \mathbb{E} as a function of the scaled impact of the 2018 EBA scenario. The coloured lines show the system-wide stress test outcome and the grey lines show the (scaled 2018 EBA) microprudential stress test outcome. These results illustrate that, for a given microprudential stress test outcome, the financial system can be stable or unstable depending on its shock-amplifying tendency.

First, our results confirm the intuition that *for a given microprudential stress test outcome, the stability of the financial system depends on the system’s shock-amplifying tendency*,⁶² which implies that microprudential stress tests alone are insufficient to assess financial stability. For example, comparing the orange lines in the left and right panels of Figure 3.3 shows that the system can be very unstable ($\mathbb{E} \approx 0.9$) or quite stable ($\mathbb{E} \approx 0.1$) for a microprudential stress test outcome of $\mathbb{E} = 0$ given an initial shock of

⁶¹A number of banks have leverage ratios close to – or below – their Basel III minimum requirements following the initial shock. This in part reflects the fact that the leverage ratio minimum requirement was not in force in most of the EU at the time of the stress test, though banks were required to disclose their leverage ratios. This stands in contrast to the situation in the UK, where major UK banks and buildings societies have been subject to a minimum Tier 1 leverage ratio requirement of 3% and an additional countercyclical leverage buffer (CCLB) for several years (Bank of England (2015)), and where banks remained comfortably above their leverage ratio hurdle rates in the 2018 stress test after management actions had been considered (Bank of England (2018)). When we remove the leverage ratio constraint in our model, contagion reduces substantially, but the qualitative findings remain robust.

⁶²The shock-amplifying tendency of the system depends, among others, on the ‘usability’ (see Section 3.6.4) and size (see Section 3.6.6) of regulatory buffers as well as the resolution regime (this section).

$x = 0.5$, demonstrating that the microprudential stress test outcome provides incomplete insight into the system’s stability. Comparing two panels of Figure 3.3 also demonstrates that *systemic risk is much lower in the case of ‘contagion-free resolution’ than in the case of ‘disorderly liquidation’* – reaffirming the importance of bank resolution to promoting financial stability. However, even if resolution is contagion-free, amplification of the initial stress scenario may still occur due to the actions institutions take – for example, to avoid default (e.g. by delevering when the usable part of the capital buffer has been exhausted; see Section 3.5.3 on the banks’ behaviour). In the right plot, the excess systemic risk (orange line) above the initial impact (grey line) is solely generated by such ‘pre-default contagion’, as resolution is assumed to be ‘contagion-free’ (see Section 3.5.3).

These results also confirm that financial stability may be highly non-linear in the impact of the initial shock, with the onset and sharpness of the turn towards instability depending on the system’s shock-amplifying tendency. This can be seen in the figure, which in the case of ‘disorderly liquidation’ shows sharp increases in the systemic risk measure as the severity of the initial shock increases. Comparing the ‘disorderly liquidation’ plots of Figure 3.3 and Figure 3.4 (where leverage ratio buffers are tripled) shows that the system becomes more shock absorbing when leverage buffers are increased. This is not surprising, because increasing the leverage buffer not only delays the onset of the non-linear jump towards instability in the case of ‘disorderly liquidation’ (from $x = 0.2$ in Figure 3.3 to $x = 1.2$ in Figure 3.4), but also makes the non-linearity less pronounced. We note that in the case of ‘contagion-free resolution’, we do not observe such sharp non-linearities.⁶³

Second, it is clear from Figure 3.3 that *microprudential stress tests may significantly overestimate financial stability*, particularly in cases where banks fail in a disorderly manner or where the macroeconomic shock is particularly severe. This finding is consistent with the findings of Sarin & Summers (2016), who show that market measures suggest banks would not be resilient to a scenario such as that applied by the 2018 Federal Reserve (FED) stress test even though the FED microprudential stress test outcome deems banks to be robust in such a case.

Microprudential stress tests generally expose banks to severe scenarios calibrated to previous crises, and so implicitly aim to include the impacts of higher order contagion

⁶³In the case of contagion-free resolution, it is possible for systemic risk to be lower if the initial shock is larger; for example, the fraction of bank defaults for an initial shock of $x = 0.8$ is smaller than that for $x = 0.6$ in Figure 3.4. If banks in our in our system-wide stress test model default after the initial shock, they take no actions and cause no further amplification, whereas if they survive but are constrained they will amplify the shock. In reality, the initial shock would not be instantaneous, so those banks that default due to the shock could still take actions to try to avoid this outcome, thus amplifying losses beyond those captured here.

effects. The non-linear nature of such effects, however, means that simply setting a severe scenario does not guarantee that the full financial stability implications of contagion dynamics will be captured – not least because the shock-amplifying tendency of a financial system markedly changes over time (e.g. due to shifts in the resolution regime, risk perception in markets). For example, in the edge case of ‘disorderly liquidation’ in Figure 3.3, even setting a very severe microprudential stress scenario (e.g. of $x = 1$ given by the 2018 EBA stress test, which implicitly seeks to include higher-order effects) would not capture the degree of instability we observe when systemic amplification mechanisms are included, even for mild initial shocks (of $x = 0.25$). At approximately $\mathbb{E} \approx 40\%$, as given by difference of the orange line $\mathbb{E} \approx 40\%$ (at $x = 0.25$) and the grey line $\mathbb{E} \approx 10\%$ (at $x = 1$), the underestimation of systemic risk in the microprudential stress test is significant.

Third, our results suggest that *the risk of financial instability in the European banking system is driven by the leverage ratio constraint, which binds more than the risk-weighted or LCR constraint*. Figure 3.4 illustrates that increasing leverage ratio buffers significantly reduces systemic risk. We further illustrate the relative importance of the leverage ratio in Figures A.2 and A.3 in Appendix A.3, which show that if we impose only the risk-weighted capital ratio constraint the system remains stable for much larger regions of the initial shock than when we impose only the leverage ratio constraint. This result is a function of the fact that the banks in our system are on average closer to breaching their leverage ratio constraints than their risk-weighted capital buffer constraints, both before and after the initial shock (see summary statistics in Table A.3 in Appendix A.3).

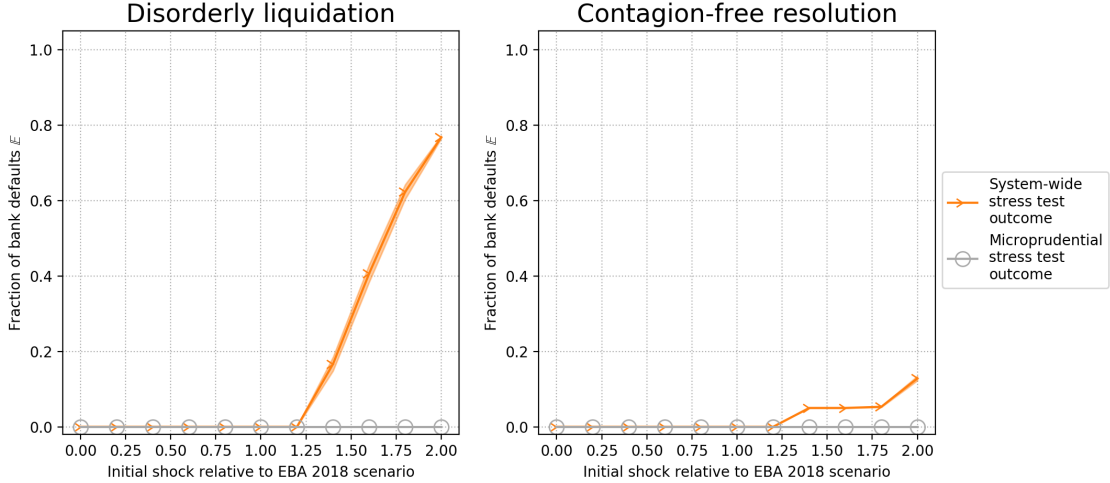


Figure 3.4: This figure has the same set-up as Figure 3.3, except here we have tripled (i.e. $y^p = 3$ in equation 3.9) the combined risk-weighted capital buffer ρ_i^{CB} . Tripling the buffer reduces the shock-amplifying tendency of the financial system and delays the non-linear divergence of the system-wide stress test outcome from the microprudential stress test outcome. It also almost completely eliminates systemic risk \mathbb{E} in the case of ‘contagion-free resolution’.

3.6.3 Contagious Feedback Loops Between Banks and Non-Banks

Our baseline model only includes banks and their interactions (see Table 3.2). In the second policy experiment, we add asset managers and hedge funds to our financial system to show that *expanding the types of financial institutions included in the stress test changes the expected financial stability outcomes*. In Figure 3.5, we show how systemic risk outcomes change when we add non-banks, first separately and then together. We show results for the ‘contagion-free resolution’ and ‘disorderly liquidation’ cases. For the latter, we additionally show results for set-ups in which banks’ capital buffers are doubled. Including hedge funds allows us to assess the risks posed by margin calls and the withdrawal of funding from their bank counterparties in terms of prompting asset sales that further depress prices and amplify banking system losses. On the other hand, including asset managers captures the risk that the price impacts of banks’ asset sales affect the performance of asset managers’. That, in turn, could prompt shareholder redemptions which may force asset managers to sell into the market to meet them, further reinforcing fire sale dynamics. To the best of our knowledge, we are the first to include the above-mentioned contagious feedback loops among banks, hedge funds, and asset managers.

We find that *including hedge funds and asset managers in our system-wide stress test increases systemic risk* modestly. Accordingly, excluding hedge funds and asset managers from a banking system stress test may lead regulators to underestimate systemic

risk (or overestimate the resilience of banks), particularly if banks' actions under stress are likely to affect those institutions and the markets they operate in. The relatively modest magnitude of the effect of including these institutions likely reflects the fact that we only include EU-based asset managers and hedge funds, not those based offshore. As a consequence, hedge funds in particular only hold a small share of total assets in our model.⁶⁴ However, at the same time non-bank financial institutions (such as hedge funds) might be willing and able to buy when banks are forced to sell, potentially mitigating the price impacts resulting from asset sales. Therefore, a more comprehensive inclusion of non-banks could also *support* banking-sector and market stability. This type of countercyclical behaviour by institutions such as hedge funds has been observed, for example by [Czech & Roberts-Sklar \(2017\)](#) who also note that while asset managers often behave countercyclically, this can reverse in times of stress. The Bank of England's recent paper on system-wide stress simulation ([Aikman et al. \(2019\)](#)) finds that funding constraints for hedge funds can lead them to sell assets, which is in line with the result produced by our model. So whether non-bank financial institutions act pro- or countercyclically and thus amplify or dampen stress depends largely on the nature of the stress they face. This scenario-dependency underscores the need to include non-banks in system-wide stress tests in a way that captures their exposures to different types of valuation and liquidity shocks in order to obtain a holistic picture of systemic stability.

⁶⁴Hedge funds included in the *ECB Statistical Warehouse* data hold approximately 2.7% of the total assets in the banking sector included in our model. The leverage of our hedge funds is also modest, reducing the risk that they will need to undertake significant asset sales even in the face of material funding outflows. We expect that banking sector stability is more heavily influenced by hedge funds if their size or leverage increases, and note that we do not account for the distribution of leverage between hedge funds. In this context, we stress the importance of initiatives to try to measure fund leverage, for example see <https://www.iosco.org/news/pdf/IOSCONEWS515.pdf>. The aggregate asset value of asset managers meanwhile is much more significant at approximately 57.2% of the aggregate asset value of the banking sector, hence the impact of their inclusion on systemic risk is larger.

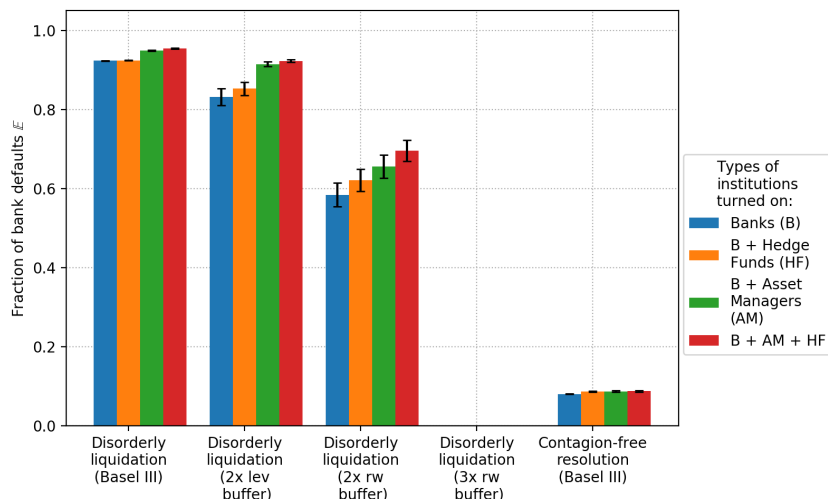


Figure 3.5: This figure shows the average fraction of bank defaults in a systemic event \mathbb{E} for stress tests including different constellations of institutions (i.e. banks, banks and asset managers (AMs), banks and hedge funds (HFs), banks and AMs and HFs) and for different regulatory regimes (i.e. for ‘disorderly liquidation’ or ‘contagion-free resolution’, for Basel III settings or for Basel III settings except that the leverage buffer λ_i^{CB} or risk-weighted capital buffer ρ_i^{CB} is doubled or tripled). Stability of the banking sector is (negatively) affected by non-banks (i.e. AMs and HFs); exclusions of these institutions from (banking) stress tests is thus likely to lead to an overestimation of resilience.

3.6.4 Amplification of Contagion Mechanisms

In the third policy experiment we show that *some combinations of contagion mechanisms are mutually amplifying*: the impact of the combination of contagion channels is greater than the sum of their impacts when considered individually. As discussed in Section 3.5.2, our stress test model includes four contagion mechanisms, (1) overlapping portfolio contagion (O), (2) exposure loss contagion (E), (3) funding contagion (F), and (4) collateral contagion (C). We use the flexibility the generic framework provides by the explicit modelling of contractual features and counterparty relationships to exclude (‘turn off’) each of these channels, by (1) setting the price impact equal to zero; (2) setting LGD equal to zero; (3) redirecting interbank contracts and repo contracts to external nodes that are always able to repay, and (4) removing the margin call obligation from repo contracts. Then we assess the impacts of various combinations of these channels in Figure 3.6 where, for instance, the label *O&E* means that only overlapping portfolio contagion and exposure loss contagion are included (‘turned on’).

Figure 3.6a shows that contagion mechanisms are mutually amplifying.⁶⁵ For example, if we assume a price impact of 5%, we find that systemic risk due to overlapping portfolio contagion alone is small (at around $\mathbb{E} \approx 5\%$), systemic risk due to exposure

⁶⁵To illustrate the relevant dynamics most clearly, we use the ‘disorderly liquidation’ case and increase banks’ capital buffers for this experiment. Our qualitative conclusions are robust to different parameter settings.

loss contagion is moderate in size (i.e. around $\mathbb{E} \approx 25\%$), and instability due to funding contagion and collateral contagion is absent (i.e $\mathbb{E} = 0\%$). However, the systemic risk of these four contagion mechanisms considered together is substantial (around $\mathbb{E} \approx 85\%$). Figure 3.6b, which shows a direct measure of this amplification, illustrates this finding. Focusing on the 5% price impact point in Figure 3.6b (the middle set of bars), we observe that the ratio of the systemic risk \mathbb{E} caused by the joint set of contagion mechanisms *over* the systemic risk produced by the sum of the individual contagion mechanisms could be as large as approximately three when all contagion mechanisms are considered.⁶⁶ Based on these findings, it is clear that modelling contagion mechanisms in isolation may lead to an underestimation of systemic risk as large as 300%. As far as we are aware, we are the first to show that overlapping portfolio contagion, exposure loss contagion, funding contagion and collateral contagion are mutually amplifying.

Our results also show that *the degree of amplification of systemic risk is heterogeneous for different sets of jointly-considered contagion mechanisms, and varies with the liquidity of markets*. This is illustrated by the different heights of the bars in Figure 3.6b. By comparing the height of the bars for the different price impacts in Figure 3.6b and A.4, we observe that the degree of amplification is heavily dependent on the *price impact*. For instance, the amplification is much smaller for a 0% price impact than at the 5% price-impact point.⁶⁷ This result clearly shows that market illiquidity can act as a powerful amplifier of other contagion mechanisms. As far as we are aware, we are also the first to highlight that the degree of amplification is heterogeneous for different sets of contagion mechanisms and in market illiquidity.

⁶⁶The same results are shown in absolute terms in Figure A.4 in Appendix A.3, which shows that the contagion mechanisms that amplify each other most in relative terms may not be the same contagion mechanisms that amplify each other most in absolute terms.

⁶⁷In Figure 3.6b, the amplification smaller than one is an artefact of the finite size of the European system, which prevents the systemic risk \mathbb{E} produced by the joint set of contagion mechanisms to exceed that of the sum of the parts. Since the individual contagion channels already cause almost maximal instability (for illustrative purposes, consider $\mathbb{E} = 0.7$ for each of the four contagion mechanisms), summing their systemic risk is going to be higher ($0.7 + 0.7 + 0.7 + 0.7 = 2.8$) than the systemic risk caused by the joint set of contagion mechanisms (which, for illustrative purposes, we set equal to $\mathbb{E} = 0.95$).

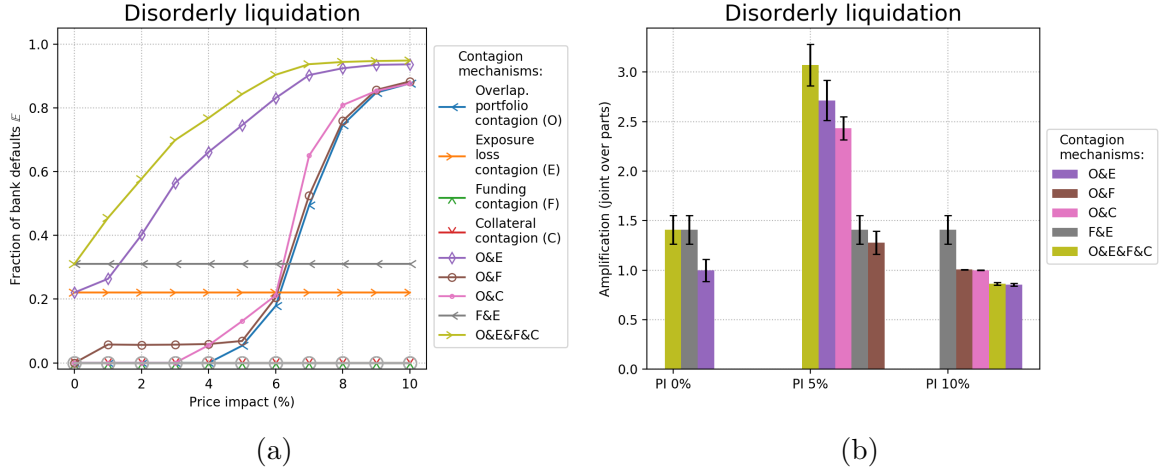


Figure 3.6: This figure shows the amplification among contagion mechanisms (overlapping portfolio contagion, exposure loss contagion, funding contagion and collateral contagion) for the case of ‘disorderly liquidation’ where the leverage buffer λ_i^{CB} is made two-and-a-half times larger (i.e. $y^\lambda = 2.5$). For instance, ‘O & E’ means that overlapping portfolio contagion and exposure loss contagion are included (‘turned on’) and the other contagion mechanisms are excluded (‘turned off’). Shock amplification is heterogeneous among different sets of contagion mechanisms and in the market liquidity. Plot 3.6a shows systemic risk \mathbb{E} as a function of the price impact for various combinations of contagion mechanisms. Plot 3.6b also elucidates Plot 3.6a by showing the *amplification* among sets of contagion mechanisms for different price impacts (PI). Amplification is computed as the systemic risk of the joint set of contagion mechanisms \mathbb{E} over the sum of the systemic risk \mathbb{E} of the individual contagion mechanisms. Amplification greater than one means that the considered contagion mechanisms are mutually amplifying. If the amplification is equal to one, then contagion mechanisms do not amplify each other.

3.6.5 ‘Usability’ of Buffers and Contagion

Pre-default contagion is in large part a function of institutional behaviour, which is why we examine how different behaviour in the face of constraints affects systemic contagion. In particular, we show in Figure 3.7 that *the more ‘usable’ banks perceive their buffers to be, the lower the risk that they will take actions (pre-default) that cause systemic amplification*. Figure 3.7 illustrates this point for the case of ‘contagion-free resolution’, where the usability of the risk-weighted capital buffer ρ_i^{CB} , the leverage buffer λ_i^{CB} , and the LCR Λ^S are varied in quantiles from $u = 0\%$ to $u = 100\%$.⁶⁸ To the best of our knowledge, while Basel III and some academics (e.g. Goodhart et al. (2008), Goodhart (2013)) have qualitatively underscored the importance of usable buffers for financial stability, we are the first to quantitatively demonstrate it in a system-wide stress test setting.

⁶⁸This result also holds for the case of ‘disorderly liquidation’, see Figure A.6 in Appendix A.3. In our experiments, the usability of the liquid asset buffer is not important for systemic risk (see Figure A.7 in Appendix A.3). This is because the 2018 EBA stress test scenario considers a solvency shock rather than a liquidity shock. Moreover, most deleveraging options banks have improve their liquidity position.




Figure 3.7: This figure shows systemic risk \mathbb{E} for the case of ‘contagion-free resolution’ as a function of the price impact (left plot) or as a function of the scaled 2018 EBA scenario (right plot) for different ‘levels of usability’ of capital buffers. If banks consider $u = 25\%$ of their regulatory buffers to be usable, they will act, if $u = 25\%$ of their regulatory leverage buffer λ_i^{CB} , risk-weighted capital buffer ρ_i^{CB} , or LCR Λ^S buffer is exhausted, to avoid a further depletion of the regulatory buffer (see Section 3.5.3 for implementation details). Resilience increases in the ‘usability’ of regulatory buffers. This result holds irrespective of the market liquidity or the stress scenario.

3.6.6 Calibration of Buffers with System-Wide Stress Tests

In the final policy experiment, we show that *the size of regulatory buffer required to limit systemic risk may be underestimated if system-wide dynamics are not taken into account*. Figure 3.8 shows systemic risk \mathbb{E} for different buffer sizes and for both bank failure edge cases as a function of the initial shock. The top row shows how stability changes if regulators double or quadruple the regulatory risk-weighted capital buffer relative to the Basel III standard, and the bottom row shows the same for the regulatory leverage buffer.⁶⁹

⁶⁹We assume that banks continue to maintain the same management buffer over their regulatory buffers as in our default calibration, see the default settings in Table 3.2.

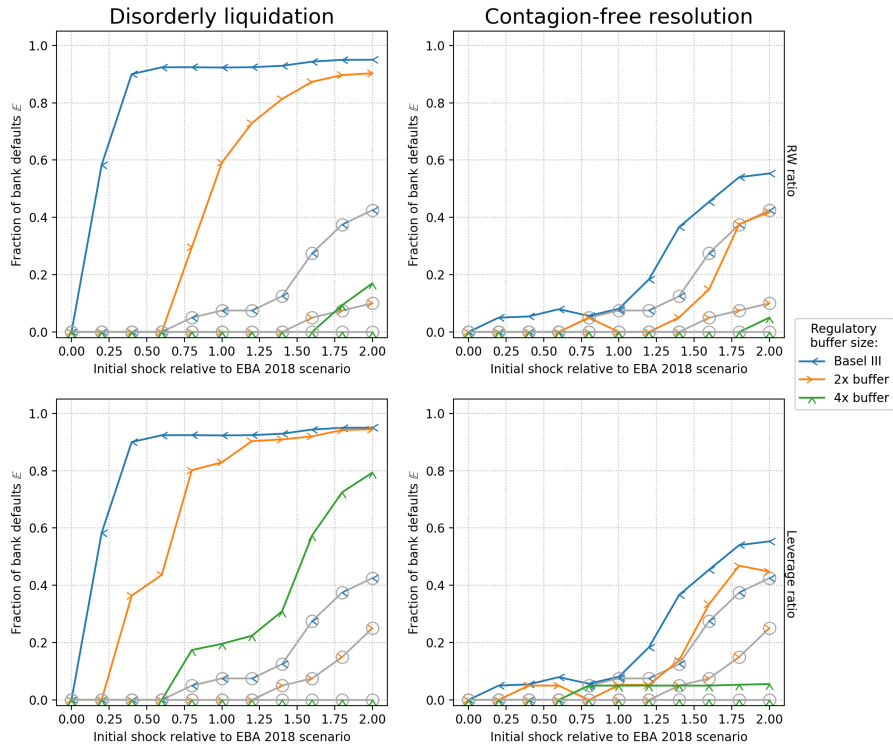


Figure 3.8: This figure shows systemic risk \mathbb{E} for both the case of ‘disorderly liquidation’ and ‘contagion-free resolution’ as a function of the (scaled) 2018 EBA scenario for different regulatory buffer sizes. The top row shows the effect of doubling and quadrupling the regulatory risk-weighted capital buffer ρ_i^{CB} compared to the Basel III standard. The leverage ratio λ_i and the LCR Λ_i are also included (‘turned on’), but kept equal to their Basel standard. The bottom row shows the same as the top row, except now the regulatory leverage buffer λ_i^{CB} is doubled or quadrupled relative to the Basel III standard. Enlarging the capital buffers markedly enhances financial stability. This suggests that regulators relying on microprudential stress tests (grey-coloured lines) rather than system-wide stress tests (coloured lines) to calibrate regulatory buffers will overestimate resilience.

Figure 3.8 shows that, as the regulatory buffer size increases (whether through an increase in the risk-weighted ratio or the leverage ratio), systemic risk drops for any initial shock size regardless of the resolution edge case. We obtain a similar result for any level of price impact (see Figure A.9 in Appendix A.3). Significantly smaller capital buffers are needed to achieve the same level of financial stability in regimes where banks fail via a ‘contagion-free resolution’ than when they undergo a ‘disorderly liquidation’.

Crucially, when we take contagion dynamics into account the buffers required to contain systemic risk are significantly higher. This result suggests that relying solely on microprudential stress tests to calibrate buffers risks overestimating resilience. Figure A.9 in Appendix A.3 further illustrates this findings, and shows that the microprudential stress tests estimates that the Basel III buffers can effectively mitigate systemic risk to

a level under $\mathbb{E} = 10\%$ in case of disorderly liquidation, whereas a when system-wide dynamics are taken into account buffers need to be more than doubled to achieve the same outcome. This result is consistent with the results of our policy experiment on buffer usability: banks that have *more sizable, usable buffers* can absorb more shocks without having to engage in destabilising actions. Increasing regulatory buffers tends to bring down contagion defaults more than it reduces the initial defaults (see Figure A.9 in Appendix A.3). This points to the special function that regulatory capital buffers perform in containing contagion and in reducing the inherent shock amplifying tendency of the financial system. As far as we are aware, we are the first to demonstrate the importance of using system-wide stress tests to calibrate buffers to avoid underestimating the buffer size that is needed to maintain stability.

Figure A.9 illustrates this risk. Imagine that regulators believe the initial shock size will not exceed the 2018 EBA shock ($x \in [0, 1]$), and that they wish to curtail systemic risk below $\mathbb{E} = 10\%$ for a price impact in interval $[0\%, 10\%]$, in a regime where banks are ‘disorderly liquidated’. In this scenario, the microprudential stress test would find that the Basel III buffers are sufficient (the grey-blue ‘Basel III’ line at $\mathbb{E} \approx 5\%$ in the top-left panel of Figure A.9). However, when system-wide dynamics are taken into account, regulators would find that they need to more than double the risk-weighted capital buffers to achieve the same result (the green ‘3x buffer’ line at $\mathbb{E} = 0\%$ in the top-left panel of Figure A.9 is the first line to fall underneath $\mathbb{E} = 10\%$).

3.7 Policy Implications and Conclusion

In a highly connected financial system, seemingly local shocks can be propagated and amplified to take on systemic importance. However widely recognised, this reality is not yet reflected in the design of financial stress tests, which are not yet system-wide in scope and do not coherently combine multiple interacting contagion and amplification mechanisms as well as the behavioural responses of heterogeneous financial institutions to shocks. We have outlined a generic framework for the development of system-wide financial stress tests with multiple interacting contagion and amplification channels and heterogeneous financial institutions. We have shown that this framework – thanks to the way in which it conceptualises financial systems, its advanced simulation engine, and its software library (the ‘Economic Simulation Library’, or ‘ESL’) – can flexibly implement stress tests ranging from simple toy models to large-scale, data-driven models with a high degree of verisimilitude.

We used this framework to implement a system-wide stress test for the European financial system that incorporates amplification risks associated with default contagion,

price-mediated contagion via asset sales, funding contagion, and liquidity stress via margin calls. When comparing our findings to the European Banking Authority’s stress test from 2018, we found that our system-wide approach reveals hidden weaknesses in the resilience of the financial system. This raises the concern that current stress test results may provide false comfort to regulators, markets, and the public at large. Our findings have at least three important implications for policymakers.

1. System-wide stress tests are necessary complements to microprudential stress tests to assess systemic risk. Our findings support and add to the growing body of evidence suggesting that capturing endogenous shock amplifications in stress test is critical to assess financial stability. Our finding that a positive microprudential stress tests outcome does not guarantee resilience, and that this problem cannot simply be resolved by increasing the severity of the stress scenario as a proxy for amplification dynamics, providing grounds for concern. Because our findings demonstrate that financial stability is critically defined by amplification dynamics, microprudential stress tests can be meaningfully complemented by macroprudential overlays (e.g. [Dees & Henry \(2017\)](#), [Fique \(2017\)](#)) and by considering contagion dynamics in regulatory exercises (e.g. [BoE \(2017b\)](#)).

2. The usability of capital is key to systemic resilience. Our findings suggest that (perceived) restrictions on the usability of capital can increase systemic risk. Perception is hard to regulate, and there are other legitimate considerations that necessarily inform the design of regulatory buffers (e.g. incentives to behave opportunistically that may call for restrictions to dividend payments when buffers are depleted, see [Armour et al. \(2016\)](#)), and our findings do not speak to how this result might best be achieved. They do, however, call attention to the sharp rise in pre-default contagion that arises when banks take action to avoid using their buffer capacity - actions that are individually rational but collectively destabilising. This motivates a careful consideration on the part of regulators when setting stress test hurdle rates or penalties for using buffer capacity – a point also recognised by [BoE \(2013\)](#).

3. The calibration of capital buffers should explicitly take into account system-wide dynamics. Our results show that failing to account for system-wide amplification risks may cause regulators to set capital buffers at too low a level. The current practice of using microprudential stress tests to calibrate capital requirements can therefore be meaningfully complemented by the use of system-wide stress tests. Regulators have the power to do so, for example when calibrating capital requirements under Pillar II of the

Basel supervisory framework (BCBS (2009) or, in the UK, as part of the calibration of the countercyclical capital buffer. Indeed, the recent incorporation of feedback and amplification effects in regulatory stress tests is a step in this direction (see e.g. BoE (2017b)).⁷⁰

The models used to calibrate capital requirements could explicitly and systematically reflect the role of bank resolution in mitigating systemic risk, such that smaller capital buffers are required if resolution is likely to be effective. Incorporating the benefits of bank resolution in system-wide stress tests that account for heterogeneity would mirror the work done by Brooke et al. (2017), who consider the benefits of an effective bank resolution regime in reducing the optimal level of capital requirements (and buffers) at a macro level. The results of system-wide stress tests would also provide a richer dataset to inform the calibration of existing regulatory capital surcharges for systemically important financial institutions (see e.g. Enriques et al. (2019)), and could even inform the calibration of a newly created top-up buffer that is explicitly designed to account for systemic risk – which, following Greenwood et al. (2017), is re-calibrated yearly to account for time-varying idiosyncratic and systemic risk.

In this paper, we take a first step in providing regulators with a generic framework that can help to them implement system-wide stress tests, and our results highlight why doing so is important. But developing this framework, and the system-wide stress test models that it hosts, remains a work in progress that will require further research and investment in capacity, software and data. Our findings highlight that further study of, for example, heterogeneous behaviour in the face of constraints, bank resolution, and non-bank behaviour will be critical to understanding contagion and amplification. Now that our generic framework enables regulators to build and use large-scale, data-driven models, the importance of data availability at a granular level⁷¹ grows further – particularly given the importance of calibration. And, finally, although our generic framework marks an important step forwards, it is itself incomplete. Integration of derivatives markets, for example, presents a key modelling challenge.⁷²

⁷⁰The outcome of such an exercise may not be an across-the-board increase in capital requirements; the effects may instead be heterogeneous, with some institutions that are more central to the functioning of the financial system being subjected to stricter requirements. This is in line with the concept of ‘network-sensitive regulation’, proposed by Enriques et al. (2019).

⁷¹More data (in the format described in Section 3.4.1) is needed to model interconnections at the contract-level, which gives important information about the pathways of contagion within a financial system. Since the 2007-2008 financial crisis, regulators have vastly enhanced data collection – for example on interbank contracts, security holdings, repurchase agreements and derivative markets (Abad et al. (2016)) – but especially for the non-banking sector, more (and better quality) data is remains a challenge.

⁷²So far derivatives markets have only, partially, been stress tested on a stand-alone basis (see e.g. Bardoscia et al. (2018), Paddrik & Young (2017), Paddrik et al. (2016)). Similarly, the roles of hedging and countercyclical behaviour, for example when fundamentalist investors buy in a fire-sale and dampen the shock, are key to consider.

Chapter 4

Systemic Implications of the Bail-In Design

4.1 Summary

The 2007-2008 financial crisis forced governments to choose between the unattractive alternatives of either bailing out a systemically important bank (SIBs) or having it fail in a system-wide disruptive manner. Bail-in has been put forward as the primary tool to resolve a failing bank, which would end the too-big-to-fail problem by letting stakeholders shoulder the losses, while minimising the calamitous impact of a bank's failure on the economy and the financial system ([WBG \(2017\)](#)). Though the aptness of bail-in has been evinced in relatively minor idiosyncratic bank failures, its efficacy in maintaining stability in cases of large bank failures and episodes of system-wide crises remains to be practically tested.

This paper investigates the stability implications of the bail-in design, for all these cases. We develop a multi-layered network model of the European financial system that captures the prevailing endogenous-amplification mechanisms: exposure loss contagion, overlapping portfolio contagion, funding contagion, bail-inable debt revaluations, and bail-inable debt runs.

We apply this stress test model to data provided by *S&P Global Market Intelligence*, the *ECB Statistical Warehouse*, the 2018 *European Banking Authority* (EBA) stress test results, allowing us to initialise balance sheets of European banks and non-banks, as well as decompose banks' liabilities in seniority classes. In line with [Hałaj & Kok \(2013\)](#), [Kok & Montagna \(2013\)](#), we reconstruct the bank debt and common asset holding networks, which interconnect these institutions. The loss absorbing requirements, which set a minimum on the amount of long-term loss-absorbing debt that banks should issue (and cannot be cross-held by banks), further inform the calibration of the maturity profile and non-bank holdings of bail-inable debt.

Our results reveal that financial stability hinges on the bail-in design, which consists of a set of ‘primary’ and ‘secondary’ bail-in parameters. This consists of the failure threshold, recapitalisation target, debt-to-equity conversion rate, loss absorption requirements, debt exclusions and bail-in-design uncertainty. Our results demonstrate that an early bail-in, strong recapitalisation and fair distribution of equity compensation by means of debt-to-equity conversion rates tends to enhance financial resilience. While a late bail-in, weak recapitalisation and unfair compensation undermines stability. We further show that excluding run-prone, short-term debt from the application of the bail-in tool, increasing the requirements on loss absorbing debt and providing investors with certainty about the bail-in design lowers the risk of contagion. While subjecting short-term debt to the application of the bail-in tool, having weaker requirements on the amount of loss absorbing debt that banks should hold and more uncertainty regarding the bail-in design, which prevents pricing risks, are a recipe for bail-in-induced instability.

Next, we find that bail-in usually works when relatively small European systemically important banks (SIBs) idiosyncratically fail regardless of the elected bail-in parameters – consistent with empirics. However, we discover that bailing-in banks in a system-wide crisis may heftily exacerbate financial fragility when the bail-in parameters are not well-tuned; and observe that ill-designed bail-ins may induce widespread contagion if large European D-SIBs idiosyncratically fail. Strikingly, we observe a phase shift from an unstable to stable system, if resolution authorities do choose appropriate bail-in parameters. Instability remains curbed even if systemic effects cause multiple banks to be bailed-in amid pervasive distress. Our evidence fortunately suggests that the pivot for stability is in the hands of policymakers; and that well-designed bail-ins could *credibly* be administered, *even* in system-wide crises. It also suggests, however, that the current policy parameters might be in the regime of instability.¹

Authors of Paper - Charles Goodhart, J. Doyne Farmer & Alissa M. Kleinnijenhuis (first author). The paper will available here: [Goodhart et al. \(2020\)](#).

4.2 Introduction

The failure of investment bank Lehman Brothers was perhaps the defining event of the 2007-2008 financial crisis, bringing the financial system and the real economy to the brink of the abyss ([Bernanke \(2017\)](#)). Incumbent governments were forced to choose between

¹We thank the participants at (internal) seminars at the International Monetary Fund, Federal Reserve Board of Governors, European Commission, European Central Bank, Institute for New Economic Thinking at the Oxford Martin School, Mathematical Institute at the University of Oxford, National Institute of Economic Research in London and MIT Sloan School of Management for their invaluable feedback.

the unattractive alternatives of either bailing-out a systemically important bank (SIB) or having to disorderly fail it in a manner that may have threatened financial stability (French et al. (2010), Bernanke (2010)). Ordinary bankruptcy procedures were entirely inadequate for dealing with SIB failures (Bernanke (2017)). The Squam Lake Report, an influential report offering guidance on the reform of financial regulation, recommended that authorities be given the necessary powers to affect an orderly resolution for systemically important institutions (French et al. (2010)). This recommendation has been adopted. The Bank Recovery and Resolution Directive (BRRD) establishes a common approach to resolution in the European Union (EU). Title II, the Orderly Liquidation provision of the Dodd-Frank Act, provides a process to quickly and efficiently liquidate a large, complex financial company that is close to failing.

Bail-in has been widely hailed as the primary tool to resolve a failing systemically important bank (e.g. FSB (2013), Chennells & Wingfield (2015), BoE (2017a)). It would end the ‘too-big-to-fail’ problem by letting stakeholders shoulder the losses – addressing moral hazard – while minimising the calamitous impact of a bank’s failure on the economy and the financial system (FSB (2013)). Bail-in is a statutory power by the resolution authority that enables to absorb losses and recapitalise a failing bank by incurring losses on its creditors – while potentially compensating them with a stake in the bank’s equity. In an ‘open’ bail-in, the goal is to revitalise the bank, whereas in a ‘closed’ bail-in the aim is to orderly shut it down (Avgouleas & Goodhart (2015)).

The efficacy of bail-in to conserve financial stability in plausible crisis scenarios remains an open question. While bail-in has proved successful in dealing with the failure of relatively minor SIBs, its aptness in handling cases of large bank failures and in steering through episodes of system-wide crises without jeopardising stability remains to be practically tested; a major financial crisis in the period after the Great Recession has yet to occur. Bail-ins on small SIBs have been victoriously carried out without compromising stability in Cyprus, the Netherlands, Denmark and the United Kingdom, for example (WBG (2017)).² Leading experts remain sceptical as to whether bail-ins can be safely carried out in periods of high financial turmoil without further exacerbating distress. Ben Bernanke notes that: ‘controversies remain over how effective even a Title II resolution would be in the context of a significant financial crisis.’ (Bernanke (2017)). Avgouleas & Goodhart (2015) argue that ‘the bail-in approach may, indeed, be much superior to bailouts in the case of idiosyncratic failure. In other cases, the bail-in process may entail

²Though, it must be noted that this was not without complementary bail-out funds. Partially so, because imposing losses on real-economy creditors was politically inexpedient. This concern, however, is outside the scope of our paper.

important risks'. The question of how the bail-in design might modify the system-wide outcome for better or for worse thus continues to be unanswered.

In this paper we comprehensively investigate the systemic implications of the bail-in design. For this, we extend a multi-layered network model of the European financial system developed by [Farmer et al. \(2020\)](#). In addition, we calibrate this model using *S & P Global Market Intelligence* data and the 2018 European Banking Authority (EBA) disclosures of the stress test results. The calibrated model captures the prevailing contagion mechanisms that could endogenously amplify shocks emanating from bail-ins: exposure loss contagion, overlapping portfolio contagion, funding contagion, bail-inable debt revaluations, and bail-inable debt runs.

In sum, our results suggest that financial stability *hinges* on the bail-in design. Our first main result reveals how the value of a set of bank-specific, 'primary' bail-in parameters – which include the failure threshold, the recapitalisation target, and the debt-to-equity conversion rate – sweepingly alters systemic risk. We find that when banks are bailed-in using a 'good' primary bail-in design consisting of an early bail-in, strong recapitalisation and fair distribution of debt-to-equity compensations, systemic risk is kept in check. Whereas when banks are bailed-in late, weakly recapitalised and unevenly compensated by equity claims, bail-ins can induce instability and exacerbate existing financial distress.

Our second main result underscores the supporting role that the more structural, 'secondary' bail-in parameters play in further enhancing or hindering stability. It shows that the stability wedge between 'good' and 'bad' primary parameters is further widened by the choice of secondary parameters. We advance to untangle the stability contribution of each of the secondary bail-in parameters: the debt exclusions from bail-in, the loss absorption requirements and the uncertainty in the bail-in design. We find that excluding run-prone, short-term debt from the application of the bail-in tool, increasing the loss absorbing requirements relative to the *status quo*, and decreasing uncertainty in the bail-in design regulators will apply in a prospective bail-in significantly reduces systemic risk; and do so to an extent that the primary parameters alone were not able to achieve, highlighting their essential role in supporting a credible bail-in regime.

Our final main result attests that it is imperative to take multiple contagion mechanisms and non-bank holdings of bail-inable debt into account when evaluating the system-wide implications of the bail-in design. Merely considering the exposure loss contagion that could ensue from bail-ins – as [Hüser et al. \(2017\)](#) have done – would falsely suggest that

the EU financial system becomes resilient to severe shocks when bail-ins are introduced to deal with failures of SIBs. Instead, if, rightly, four more prevailing contagion mechanisms are considered, then stability may be substantially damaged for severe enough shocks – more gravely so with a poor bail-in design. This result also shows that excluding bail-inable debt holdings of non-banks from the systemic analysis is a substantial mistake. Given that the loss absorbing requirements prohibit eligible bail-inable debt to be the cross-held by systemically important banks, a significant fraction of bail-inable debt is held by the non-banking sector. If exposure losses and mark-to-market losses amass to bail-inable debt holdings of leveraged non-banks, then these non-banks may be impelled to delever to avoid margin calls. Forced asset liquidations by non-banks negatively feed-back onto the banking system via overlapping portfolio contagion. If, on the other hand, these losses mount on non-leveraged non-banks then the real economy suffers. Destabilising spirals between the banking and non-banking system are amplified by a poor bail-in design.

Overall, our results show that bail-in usually works for idiosyncratic failures of smaller European SIBs regardless of the elected ‘primary’ parameters – which is consistent with previous experience (WBG (2017)). For all other cases, however, the primary bail-in parameters *do* crucially matter. We find that bank bail-ins may heftily exacerbate financial fragility in financial crisis episodes if ‘primary’ parameters are ‘badly’ chosen; and also observe that ill-designed bail-ins may induce widespread contagion if larger European SIBs idiosyncratically fail. Strikingly, we witness a phase shift from an unstable to stable system, if resolution authorities do choose ‘good’ bail-in parameters. Instability remains curbed even if systemic effects cause multiple banks to be bailed-in amid pervasive distress. Our evidence fortunately suggests that the pivot for stability is in the hands of policymakers; and that well-designed bail-ins could *credibly* be administered, *even* in system-wide crises. It also suggests, however, that the current policy parameters might be in the regime of instability.

Our contribution adds to the nascent network literature on the systemic effects of bail-in. Klimek et al. (2015) employ an agent-based network model to evaluate the economic and financial ramifications of bail-in. They compare its performance against other resolution mechanisms. Hüser et al. (2017) evaluate the systemic implications of bail-in in the EU, drawing on a calibrated multi-layered network model to bank debt and equity cross-holdings. Bernard et al. (2017) investigate the incentives for banks to contribute to a voluntary bail-in arise from their exposure to credit and price-mediated contagion. These papers neither investigate the systemic impact of the bail-in design, nor include

the prevailing contagion mechanisms and non-banks in this analysis. Instead, they take the bail-in design as is and merely explore the repercussions of exposure loss contagion (Klimek et al. (2015), Hüser et al. (2017) and Bernard et al. (2017)) and overlapping portfolio contagion (Bernard et al. (2017)). By ignoring a set of prevailing interacting contagion mechanisms, they risk underestimating the systemic footprint of the bail-in design. Though bail-in has been designed with systemic considerations in mind,³ it is not enough to assert its suitability on a system-wide scale. As Aymanns et al. (2016) have shown for the case of the Basel II leverage requirements, well-intended microprudential regulation may undermine financial resilience when systemic feedbacks are taken into account. This makes the investigation of the stability implications of the bail-in design in a networked financial system a necessary gap to fill.

The organisation of the rest of the paper is as follows. Section 4.3 introduces resolution and the bail-in tool. Section 4.4 describes the data we use to initialise our model. Section 4.5 develops our system-wide stress testing model, building forth upon Farmer et al. (2020). It consists of two parts. Section 4.5.1 focusses on modelling the design of the bail-in mechanism. Section 4.5.2 models the systemic impact of the bail-in design by capturing multiple interacting contagion mechanisms in the financial system. As part of this effort, it develops a novel method for revaluating bail-inable debt (in Section 4.5.2.2) and poses three ‘what-if’ scenarios that could lead to bail-inable debt collapses (In Section 4.5.2.3). Section 4.7 portrays our results, which focus on investigating the systemic impact of the bail-in design in cases of idiosyncratic bank failures and in cases of system-wide crises. Section 4.8 links our contribution to the literature. Finally, Section 4.9 discusses the policy implications of our findings.

4.3 Introduction to Bail-In

In this section we will introduce the bail-in tool to resolve a failing bank. A first aim of this section is to provide a general introduction to the bail-in mechanisms. A second goal is to introduce the ‘primary’ and ‘secondary’ bail-in design parameters. Our model choices will be based on this.

The Bank Recovery and Resolution Directive (BRRD), adopted in 2014, establishes a common approach within the European Union (EU) to the recovery and resolution of banks and investment firms. EU national governments transpose the BRRD into their laws. The BRRD provides resolution authorities with new powers to effectively resolve

³See: Directive 2014/59/EU of the European Parliament and of the Council.

a bank. Although some attention will be paid to banking regulations outside the EU, we review the BRRD below because the bail-in model to be developed will be applied in this chapter to data from the *ECB Statistical Warehouse* and the 2018 *European Banking Authority* (EBA) stress test results.

4.3.1 Failure Threshold: Conditions to Resolve a Bank

A bank will be resolved if the following three conditions are fulfilled:⁴

1. A bank is deemed to failing or likely to fail (FLTF);
2. Resolving the bank is in the public interest;
3. Failure of the bank cannot be prevented otherwise (e.g. early intervention measures, conversion of claims outside resolution in accordance with Article 59(1)a, a private sector solution).

While condition 3 is clear, condition 1 and 2 benefit from further explanation. A bank is deemed to be FLTF (condition 1) if at least one of the following circumstances apply:⁵⁶

1. The bank infringes or is likely to infringe upon its requirements for continued authorisation (e.g. minimum capital requirements, because it has incurred or is likely to incur large asset losses);
2. The bank is expected to insolvent in the near future;
3. The institution is or is likely to be illiquid soon.

Resolving a bank is in the public interest (condition 2) if a bank resolution meets one or more of the following key resolution objectives:

1. To ensure the continuity of critical functions (e.g. payment services on behalf of customers, taking deposits from and extending loans to households and small business, clearing and settling financial transactions; providing custody services (BoE (2017a)));
2. To avoid financial instability and specifically prevent contagion;
3. To avoid reliance on public funds (via bail-outs);

⁴See: Article 33(1) of the BRRD.

⁵See: Article 33(4) of the BRRD.

⁶See the specification of the FLTF trigger in our model in Section 4.5.1.3.

and if winding up the institution under normal insolvency proceedings would not meet the resolution objectives to the same extent.⁷

In essence, the failure threshold of a bank that is eligible to be resolved is thus given by the point at which the bank is ‘deemed failing-or-likely-to-fail’. Going forward, we will refer to this cut-off point as the ‘failing-likely-to-fail ratio’, or simply, the failure threshold. *The failure threshold constitutes the first ‘primary’ bail-in parameter.*

To comply with the resolution objectives stated above, regulators (e.g. [BoE \(2017a\)](#)) typically resolve systemically important banks (SIBs) by applying a resolution tool, such as the bail-in tool, and liquidate small banks via the regular insolvency procedure. Contrary to SIBs, the smallest firms usually do not pose a threat to financial stability or have sufficiently critical functions to justify the use of resolution tools. *In line with the resolution approach, we will resolve SIBs and liquidate non-SIBs in our model.*⁸

4.3.2 The Prime Resolution Tool is the Bail-In Tool

If the three resolution conditions – previously discussed – are met then the resolution authority may resolve a bank applying one or a combination of the following resolution tools⁹: (1) the sale of business tool; (2) the bridge institution tool; (3) the asset separation tool;¹⁰; and (4) the bail-in tool.

While in principle all four resolution tools are available to the resolution authority. In practise, a resolution authority is likely to use the bail-in tool to resolve SIBs. The reason is that splitting up large and complex firms, so that critical functions can be preserved while other parts may be wound down, may not be feasible in a timely manner. In light of this, bail-in is typically the only tool that meets the resolution objectives ([Chennells & Wingfield \(2015\)](#)). *Since bail-in is the prime resolution tool ([Hüser et al. \(2017\)](#)), we will assume that all SIBs in our model are resolved via the bail-in tool.*

4.3.3 The Bail-In Tool

Bail-in is a *statutory* power by the resolution authority to recapitalise a bank that meets the resolution conditions to the extent necessary to restore its *short-term* viability. *The bail-in tools allows to absorb losses and recapitalise a resolved bank by writing down and/or*

⁷See: Article 33(5) of the BRRD.

⁸Since, as we will learn, all the banks in our model are SIBs, we resolve banks rather than disorderly liquidate them.

⁹See: Article 37(3)(4) of the BRRD.

¹⁰Exceptionally, the asset separation tool may only be used together with another tool, see: Article 37(5) of the BRRD.

converting the bank's bail-inable claims to equity.

In case the bail-in tool is used to recapitalise the original bank in order to restore the bank to viability, one refers to the process as an ‘open bail-in’. In case the bail-in tool is wielded to recapitalise the bridge bank, we speak of a ‘closed bail-in’. In a closed bail-in, the loss-absorbing liabilities will remain in the original legal entity, which will undergo the regular insolvency procedure, while the critical activities are transferred to the newly created bridge entity (Chennells & Wingfield (2015)). The BRRD, which applies to institutions in the European Union¹¹, allows for both open and closed bail-ins.¹², while Title II of the Orderly Liquidation provision of the Dodd-Frank Act, which applies in the United States, provides mainly for a closed bail-in process (Avgouleas & Goodhart (2015)).

4.3.3.1 Recapitalisation target

Short-term viability is aimed to be restored by recapitalising the institution by a sufficient amount such that it complies with the conditions for authorisation and sustains or regains market confidence.¹³ In other words, *a bank will be recapitalised to a certain ‘recapitalisation target’ – our second ‘primary’ bail-in parameter – compliant with these guidelines.*

Notably, bail-in thus focusses on addressing solvency rather than liquidity to make a bank viable again. Nonetheless, bail-in may indirectly address liquidity issues: a better recapitalised bank typically has lower funding costs and lesser issues with retaining access to market funding (Burrows et al. (2012)).

4.3.3.2 Debt-to-Equity Conversion Rates

The recapitalisation of the bank is brought about by writing down and/or converting to equity the bank's bail-inable claims (as noted in Section 4.3.2). It takes place in two sequential phases (see details in Section 4.5.1):

1. The loss absorption phase (which we will also refer to as ‘phase *a*’);
2. The recapitalisation part (which we will also refer to as ‘phase *b*’).

In line with Hüser et al. (2017) and in accordance with technical documents on the bail-in process (SRB (2017)), *we strictly split the loss absorption and recapitalisation phase of the bail-in mechanism when specifying our model.*

In the loss absorption phase, a bank's bail-inable claims may be partially or fully

¹¹See: Article 1(1) of the BRRD.

¹²See: Article 46(2) of the BRRD.

¹³See: Article 43(2)a of the BRRD.

written down – and potentially converted to an equity claim – to lift an institution out of insolvency.¹⁴ A failed bank may be found to be insolvent, because of incurred losses of because of previously unrecognised losses. In the recapitalisation phase, the bank’s bail-inable claims receive further haircuts, which are typically converted to equity to raise the capital base further above zero to a recapitalisation target.

Article 50 of the BRRD sets out the principles for the deb-to-equity conversion ratio, which regulators should use when determining how to set the conversion rate. *The debt-to-equity conversion rate, our third ‘primary’ bail-in parameter, determines how many shares a creditor in priority class k of the bailed-in bank receives per unit of haircut applied to its principle amount of bail-inable claim.* Article 50 tasks the EBA with the duty to provide guidelines regarding how creditors may be appropriately compensated by means of the conversion rate. This guideline provides two principles (P) regarding how the conversion rate in each priority class k should be set (EBA (2017c)):

1. *P1- No Creditor Worse Off (NCWO)*: Resolution authorities should seek to ensure that no creditor or shareholder is expected to incur greater net losses than it would have incurred in winding up the bank under normal insolvency proceedings.¹⁵ This condition is in place as is a safeguard for creditors. The principle seeks to avoid having to use the resolution financing arrangements. Under Article 75 of the BRRD claimants are entitled to the difference between the expected loss under resolution and insolvency proceedings if this is positive.¹⁶
2. *P2 - The Preservation of the Hierarchy of Claims*: According to Article 50 of the BRRD, regulators may apply differential conversion rates under certain circumstances. For example, differential rates may be applied to ensure that no creditor is worse off, or to ensure that the creditor is appropriately compensated for the haircut it has incurred. Whenever differential conversion rates are applied, the conversion rate to more senior liabilities under the applicable insolvency law must be higher than the conversion rate to more junior liabilities under the applicable insolvency law. Claims in the same priority class must be treated in an equitable manner. The resolution authority should not apply differential conversion rates, if it is not in the interest of the resolution objectives (see Section 4.3.1 and Article 34 of the BRRD (EBA (2017c))).

¹⁴See: Article 46(1)a of the BRRD.

¹⁵See also: Article 74(2) of the BRRD.

¹⁶Additionally, resolution authorities should assess whether appropriate regard has been had to the right of property under the EU charter of Fundamental rights (EBA (2017c)).

4.3.3.3 Loss Absorbing Requirements

The bail-in tool can only be effective in absorbing losses and recapitalising a bank to a desired capital ratio if the bank has sufficient bail-inable debt. In order to ensure that banks have sufficient bail-inable debt at the start of the bail-in, or in other words, have enough ‘loss absorbing capacity’, the Financial Stability Board (FSB) and the Bank Recovery and Resolution Directive (BRRD) have established minimum requirements for banks’ loss absorbing capacity. The FSB has established a common minimum total loss absorbing capacity (TLAC) that applies to all globally systemically important banks (G-SIBs) taking effect on 1 January 2019 (FSB (2015b)). The BRRD has put in place a minimum requirement for own funds and eligible liabilities (MREL),¹⁷ having taken effect since 1 January 2016. The MREL does not only apply to European GSIBs, but to each bank under the remit of the BRRD. Other than the TLAC standard, it is determined on a case-by-case basis. *These loss absorption requirements constitute our first set of ‘secondary’ bail-in parameters.*

Not all debt instruments are eligible to count towards the loss absorption requirements. While differences exist in the debt eligibility criteria of TLAC and MREL, they have two important rules in common. First, eligible debt must have a time to maturity of at least one year. Second, eligible debt cannot be cross-held by banks. An extensive comparison among the debt eligibility criteria of TLAC, MREL and bail-inable debt can be found in Appendix B.5.1. (The eligibility criteria of bail-inable debt are also separately discussed next.)

4.3.3.4 Debt Exclusions from Bail-In

Instruments that are eligible to receive haircuts in the loss absorption and recapitalisation phase of bail-in are part of what we refer to as the ‘hierarchy of bail-inable claims’. The hierarchy of bail-inable claims respects the hierarchy of claims prescribed in regular insolvency proceedings, except for the fact that certain debts are excluded. Similar to the treatment of losses in the insolvency hierarchy, bail-inable debt contracts in a lower priority class k (or in other words, in a lower seniority class) are subject to losses sooner than debt contracts k_x (where $k_x > k$) in a higher priority class. In keeping with the insolvency hierarchy, the resolution authority presiding over a bail-in should write-down or convert instruments in the following order:¹⁸ (1) CET1 capital; (2) AT1 capital; (3) T2 capital (4) subordinated debt; and (5) other eligible liabilities. Instruments in the same priority class k (or in other words, in the seniority class) must be treated in an equitable manner by distributing losses proportionally among bail-inable liabilities in that same

¹⁷See: Article 45 of the BRRD.

¹⁸Article 48(1) of the BRRD.

priority class k .¹⁹

Even though all bailed-in debt in a priority class is treated in an equitable manner, all debt (i.e. including non-bail-inable debt) may not be. Relevant exclusions of debt from the bail-inable debt hierarchy – which are included in the regular insolvency hierarchy – are:²⁰ (1) covered deposits (including those covered by the deposit guarantee scheme (DSG)); (2) liabilities to institutions with a (time to) maturity less than seven days; and (3) secured liabilities up to their collateral value (including covered bonds). *Our second ‘secondary’ parameter is the set of debt exclusions from bail-in.*

In addition to fixed exclusions, exclusions from bail-in may also be applied on a discretionary, *ad hoc* basis. Resolution authorities may decide to (partially) exclude certain liabilities, such as derivatives, from the application of write-down or conversion powers when at least one of the following conditions is met:²¹ (a) the exclusion is strictly necessary and proportionate to avoid contagion; (b) the exclusion is necessary and proportionate to achieve continuity of critical functions and core business lines; (c) bailing in that liability is not possible within the short-time frame; and (d) the application of the bail-in tool to those liabilities would cause a destruction in value such that the losses borne by other creditors would be higher than if those liabilities were excluded from bail-in. When exclusions are applied, breaching the No Creditor Worse Off (NCWO) condition, becomes more likely (as we will show).

4.3.3.5 Uncertainty in the ‘Primary’ Bail-In Design

The determination of the necessary loss absorption and recapitalisation amount hinges on the valuation of assets and liabilities at the time of the application of the bail-in tool. This valuation is to be executed by a designated independent valuer.²²²³ The valuation of a bank’s balance sheet is a complicated task that is inherently inconclusive. First, balance sheets of all large institutions tend to be complex and opaque. Further, contracts on the balance sheets of banks typically have multiple (equally valid) methods of valuation with volatile valuation outcomes. Finally, balance sheets often contain unrecognised losses.

In part due to the difficulty to value balance sheets, some resolution authorities indicate (e.g. [BoE \(2017a\)](#)) that bringing phase *a* and *b* to conclusion may take months. However, in line with [Hüser et al. \(2017\)](#), *we assume that a bail-in will be completed in*

¹⁹Article 48(2) of the BBRD.

²⁰See: Article 44(1) of the BRRD.

²¹See: Article 44(2) of the BRRD.

²²See: Article 36(1) of the BRRD.

²³More comprehensively, the purposes of valuation are to inform: (a) whether a firm is FLTF; (b) the scope liabilities subject to bail-in; (c) the amount of loss absorption and recapitalisation needed; (d) the restructuring plan; (e) the market value estimation; (f) the recovery each investor would have had in insolvency to deal with the No Creditor Worse Off (NCWO) condition ([BoE \(2017a\)](#)).

one time step for the following reasons. First, in principle phase *a* and *b* of the bail-in mechanisms *could* be completed in a day. Take contingent convertibles (CoCos), the contractual analogue of bail-in, as an example.²⁴ Upon issuance, it is clear what the CoCo trigger and conversion rate is. Hence, when the CoCo is triggered the debt write-down and/or equity conversion is immediate.

Second, we believe that bail-in *should* be completed in a short-time period. Large uncertainty tends to take over financial markets when the terms of the bail-in, and thus the losses that will be suffered, are only announced after months, thereby aggravating financial instability. Bail-in could learn from CoCos in the following regard: the terms of the bail-in must be clear a priori to price risk (i.e. bail-inable debt), and bail-in should be materialised quickly. *This brings us to our third ‘secondary’ bail-in parameter, (un)certainty in the primary bail-in design: are the three primary parameters (i.e. the failure threshold, the recapitalisation target, and debt-to-equity conversion rates) that a resolution authority intends to apply in a prospective bank bail-in knowable and known in advance?*

4.3.4 Following a Bail-In: Restructuring

Once both parts of the bail-in mechanism are completed, a bank is restructured to address the cause of its failure and restore its *long-term* viability. As part of this process, the bank is required to submit a business reorganisation plan within a month after the application of the bail-in tool.²⁵ This business reorganisation plan stipulates the timing and approach to, for example: (a) withdraw from loss-making activities; (b) restructure existing activities to make these competitive; and/or (c) sell assets or of business lines.²⁶ The resolution authority will assess the likelihood of the plan to restore long-term viability, and may require amendments.²⁷ *We will not model this, as our model focusses on financial crisis dynamics in the short-term: one month (or in other words, 30 timesteps) – where institutions balance their books and fulfil payment obligations once a day.*

²⁴A statutory bail-in mechanism differs from contractual write-off or conversion features, such as contingent convertibles (CoCos). While both involve creditor-financed recapitalisations, CoCos are private financial contracts with principal and scheduled coupon payments that can be automatically converted into equity or written down when a predetermined trigger event occurs, whereas bail-in is a statutory power that enables resolution authorities to eliminate or dilute existing shareholders, and to write-down or convert claims (Rutledge et al. (2012)).

²⁵See: Article 52(1) of the BRRD.

²⁶See: Article 52(6) of the BRRD.

²⁷See: Article 52(1) of the BRRD.

4.4 Data

We apply our system-wide stress test model to data provided by *S&P Global Market Intelligence*, the *ECB Statistical Warehouse* and the 2018 *European Banking Authority* (EBA) stress test results, allowing us to initialise balance sheets of European banks and non-banks, as well as decompose banks' liabilities in seniority classes.

We include the 48 banks in our model that participated in the 2018 EBA stress test.²⁸ We initialise their balance sheets using *S&P Global Market Intelligence* data. The balance sheet of each bank is depicted in Figure 4.1. The precise data-codes we used from *S&P Global Market Intelligence* to initialise the balance sheet items can be found in Table B.1 in Appendix B.1.1. The *S&P Global Market Intelligence* also allows us to map the liabilities to seniority classes of debt, which is necessary to estimate each bank's bail-inable debt. Table B.2 in Appendix B.1.2, and the accompanying formulas, describe exactly how this is done. Figure 4.1 illustrates the composition of each bank's bail-inable debt, with in its left column a breakdown of the bank's assets \mathcal{A} and in its right column a breakdown of the bank's liabilities L_i .

C_i , Cash	D_i , Deposits
Y_i , External Assets	\tilde{I}_i , Interbank Liabilities
T_i , Tradable Assets	\tilde{R}_i , Repos
I_i , Interbank Assets	\tilde{O}_i , Other Liabilities
R_i , Reverse Repos	
O_i , Other Assets	E_i , Equity

Figure 4.1: Stylised balance sheet of a bank $i \in \mathcal{B}$, where \mathcal{B} denotes the set of banks. The bank's assets A_i are given by the sum of its cash C_i , external assets Y_i , tradable assets T_i , interbank assets I_i , reverse repos R_i , and other assets O_i . The bank's liabilities L_i are given by the sum of its deposits D_i , interbank liabilities \tilde{I}_i , repos \tilde{R}_i , and other liabilities \tilde{O}_i . The bank's book equity E_i is given by the difference of its assets A_i and liabilities L_i . The bank's CET1 equity \tilde{E}_i is approximated by equation B.2. The value of the bank's own funds F_i , its total capital instruments, is given by the sum of its CET1 equity \tilde{E}_i , AT1 equity \tilde{E}_i^{AT1} and T2 equity \tilde{E}_i^{T2} (see equation 4.6).

²⁸See the list of participating banks here: <https://eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing/2018/results>.

We use the estimates of the *ECB Statistical Warehouse* to construct a representative balance sheet of all non-banks in the European financial system. We next, proportionally, split χ percent of the balance sheet of the representative non-bank in a ‘leveraged non-bank’ and $(1 - \chi)$ percent in a ‘non-leveraged non-bank’ (and remove the representative non-bank from the system), to reflect that we do not have reliable estimates of the relative size of leveraged non-banks.²⁹ To express that we do not precisely know the initial leverage of the leveraged non-bank we define a parameter λ^{t0} .³⁰ The stylised balance sheet of the leveraged and non-leveraged non-bank are exhibited in Figure 4.2.

C_i , Cash	L_i , Liabilities
T_i , Tradable Assets	
B_i , Bail-Inable Debt Holdings	
O_i , Other Assets	E_i , Equity

Figure 4.2: Stylised balance sheet of a non-bank $i \in \mathcal{N}$, where \mathcal{N} denotes the set of non-banks. Each non-bank’s assets A_i are given by the sum of its cash C_i , tradable assets T_i , bail-inable debt holdings issued by banks B_i , and other assets O_i . The non-banks have a generic liability L_i , which is equal to zero in the case of the non-leveraged non-bank. The equity E_i of each non-bank is given by the difference of its assets A_i and liabilities L_i . Its initial leverage is defined by $\lambda^{t0} := \frac{E_i^{t0}}{A_i^{t0}}$, and altered by varying L_i^{t0} .

The loss absorbing requirements inform the calibration of the maturity profile and non-bank holdings of bail-inable debt. Since the loss absorbing requirements allow debt to count towards the measure only if: (1) the time to maturity of the debt contract is greater than a year; and (2) the debt contract is not cross-held by banks, we know all the debt below the requirements will be held by non-banks. Accordingly, we link all the bail-inable debt below each bank’s requirement to non-banks (χ to the leveraged non-bank and $(1 - \chi)$ to the non-leveraged non-bank); and we interconnect all the bail-inable debt above each bank’s requirement using the debt reconstruction method employed by [Halaj](#)

²⁹In the result section we will vary parameter χ to investigate how banking sector stability is affected by the size of the leveraged non-banks who hold bail-inable debt.

³⁰In the result section we will vary parameter λ^{t0} to investigate how banking sector stability is affected by the degree of leverage in the leveraged non-banking system.

& Kok (2013), Kok & Montagna (2013). Each bank’s non-bail-inable debt is also interconnected using this method.³¹ Based on our knowledge of the amount of tradable assets that each institution (bank or non-bank) has on its balance sheet, we randomly reconstruct common asset holding networks using Kok & Montagna (2013). To acknowledge that our results use reconstructed networks, we average our findings across N simulation runs – in which in each simulation run the networks is randomly redrawn – and show their standard deviation.

We next discuss our model.

4.5 Model

In this section we develop our system-wide stress testing model, which builds forth upon Farmer et al. (2020). This consists of two parts.

The aim of Section 4.5.1 is to arrive at a mathematical model of design characteristics of bail-in mechanisms – consisting of ‘primary’ (i.e. failure threshold, recapitalisation target, conversion rates) and ‘secondary’ (debt exclusions from bail-in, loss absorption requirements, primary bail-in design ‘uncertainty’) parameters – which has informally been introduced in Section 4.3.

The pursuit of Section 4.5.2 is to model multiple interacting contagion mechanisms in a financial system, where bail-in is the preferred method for dealing with bank failures. We will seek to explain how we jointly model five prevailing contagion mechanisms: bail-in-induced exposure loss contagion, overlapping portfolio contagion, funding contagion, bail-inable debt revaluations and halts on roll-overs of bail-inable debt (referred to as bail-inable debt ‘runs’). Additionally we will describe how contagious amplifications can arise due to both the bail-in design of a bank failure, and institutions’ behaviour in response to exogeneous or endogeneous shocks (generated by the bail-in design).

The definition of the notation used in this Section is provided in Table B.3 in Appendix B.2.

4.5.1 Modelling the Design of the Bail-In Mechanism

We approach our endeavour to model the design of the bail-in mechanism by chronicling the steps to execute a bail-in – thereby touching upon the role of each constituent part: a bank’s pile of bail-inable debt, the loss absorption requirements, haircuts in the

³¹Any excess of assets or liabilities could be seen as foreign (i.e. outside the EU) investments or foreign funding, and is linked to a passive ‘external node’.

loss absorption phase, the recapitalisation target, haircuts in the recapitalisation phase, the debt-to-equity conversion rates (fair & unfair conversion rates), and the dilution of existing equity holders.

4.5.1.1 Hierarchy of Bail-inable Claims: Debt Exclusions (First ‘Secondary’ Bail-In Parameter)

We introduced the debt inclusions in and exclusions from bail-in in Section 4.3.3.4. We now will mathematically formalise this.

The total value of bank i ’s bail-inable claims \hat{B}_i is given by the sum of bail-inable debt in each priority class (see Figure 4.3); that is,

$$\hat{B}_i := \sum_{k=1}^5 B_i^k, \quad (4.1)$$

where B_i^k denotes the value of bail-inable debt in priority class k of bank $i \in \mathcal{B}$. The mapping between a bank’s liabilities and bail-inable debt composition can be found in Section 4.4 on the model initialisation.

Since in a bail-in haircuts may be applied to priority classes $k = 2, \dots, 5$, but may not be applied to priority class $k = 1$ (which automatically re-values following asset value A_i changes), it will be useful to also define the bail-inable debt in the set of priority classes $\mathcal{K} = [k_2, \dots, k_5]$ excluding priority class k_1

$$B_i := \sum_{k \in \mathcal{K}} B_i^k < L_i \quad (4.2)$$

as we often deal with this quantity.

We note that the total value of a bank’s bail-inable debt B_i is by definition less than or equal to its amount of liabilities L_i . The reason is that some debt is excluded from the application of the bail-in tool – as we first threw light on in Section 4.3.3.4.

We further define B_{ji}^{km} to be the bail-inable contract of bank i held by institution j in priority class k with time to maturity m .

Figure 4.4 portrays our estimation of the composition of each bank’s debt \hat{B}_i based on the model initialisation to data described in Section 4.4. Figure 4.5 summerises Figure 4.4 into an aggregate.

Bail-In-Able Debt	Interbank and Other Liabilities (Excluding Short-Term Maturities)	k_5
	Interbank and Other Liabilities (Excluding Short-Term Maturities)	k_4
	T2 Capital	k_3
	AT1 Capital	k_2
	CET1 Capital	k_1

Figure 4.3: Shows the composition of bail-inable debt comprised of multiple priority classes. The bail-inable debt consists of debt or capital instruments in multiple priority classes k_x , for $x = 1, \dots, 5$. Given the stylised bank balance sheets we consider, the debt or capital instruments in each priority class are given by the following: (k_1) CET1 capital \tilde{E}_i ; (k_2) AT1 capital \tilde{E}_i^{AT1} ; (k_3) T2 capital \tilde{E}_i^{T2} ; (k_4) subordinated interbank and other liabilities (excluding contracts with a time to maturity less than 7 days); and (k_5) senior interbank and other liabilities (excluding contracts with a time to maturity less than 7 days).




Figure 4.4: Shows the size of the priority classes for European banks that partook in the 2018 EBA stress test. The various priority classes consist of the following instruments in our model: k_1 consists of CET1 equity \tilde{E}_i ; k_2 consists of AT1 capital \tilde{E}_i^A ; k_3 consists of T2 capital \tilde{E}_i^{T2} ; k_4 consists of interbank contracts \tilde{I}_i and other liabilities \tilde{O}_i ; k_5 consists of the same type of instruments as k_4 ; k_6 consists of deposits D_i ; and k_7 consists of secured contracts, among which repurchase agreements (repos) \tilde{R}_i . Notably, our estimation of bail-innable debts show that European banks have little debt in priority classes k_1 , k_2 , k_3 , k_4 and k_7 , while having lots of debt in priority classes k_5 and k_6 . Since priority classes k_1 to k_4 are small, it is likely that these priority classes suffer large losses in a bail-in, while creditors in priority class k_5 are probably less likely to suffer losses.




Figure 4.5: Shows the aggregate debt composition of each bank. We note that this result displays the same as Figure 4.5, but now aggregates the debt across banks per priority class and displays their relative size in percentages.

4.5.1.2 Loss Absorbing Requirements (Second ‘Secondary’ Bail-In Parameter)

To ensure that banks hold sufficient bail-inable debt \hat{B}_i (see equation 4.1) they face loss absorbing requirements. To that end, as we explained in the introductory Section 4.3.3.3, each European bank is subject to the MREL loss absorbing requirement. Some – those that are G-SIBs – are also subject to the TLAC requirement.³² Our goal in this Section is to spell out the TLAC requirement first and the MREL requirement next, so that we may use these to infer the percentage of each bank’s debt that is held by non-banks (as we eluded upon in the data Section 4.4). We are, to the best of our knowledge, the first to express TLAC and MREL requirements in precise formulas.

TLAC G-SIBs are required to hold sufficient TLAC-eligible instruments Z_i such that it exceeds the minimum risk-weighted TLAC requirement $T^{\rho, M}$ and the minimum leverage

³²The list of banks that are designated as G-SIBs in 2018 can be found here: <https://www.fsb.org/2018/11/fsb-publishes-2018-g-sib-list/>.

exposure requirement $T^{\lambda,M}$.³³

In addition, the CET1 equity \tilde{E}_i that is used to meet the regulatory risk-weighted capital buffers ρ_i^{CB} or the regulatory leverage buffer λ_i^{CB} may not also be used to meet the TLAC standard (FSB (2015b))(see Farmer et al. (2020) for a precise description of these buffer variables) to ensure that buffers remain usable. A buffer is considered to be usable by a bank, if the bank is willing to absorb liquidity or valuations shocks using the buffer. That is, at all times a G-SIBs loss absorbing capacity Z_i should be such that both

$$T_i^\rho = \frac{Z_i}{\Omega_i} \geq T^{\rho,M} + \rho_i^{CB} \quad (4.3)$$

and

$$T_i^\lambda = \frac{Z_i}{\hat{A}_i} \geq T^{\lambda,M} + \lambda_i^{CB} \quad (4.4)$$

hold (FSB (2015b)). In other words, the amount of TLAC-eligible instruments Z_i should at all times satisfy

$$Z_i \geq \max\{\Omega_i(T^{\rho,M} + \rho_i^{CB}), \hat{A}_i(T^{\lambda,M} + \lambda_i^{CB})\}, \quad (4.5)$$

where Ω_i denotes bank i 's the risk-weighted assets and \hat{A}_i signifies bank i 's the asset exposure (Farmer et al. (2020)).

From 1 January 2019 the minima are set to $T^{\rho,M} = 16\%$ and $T^{\lambda,M} = 6\%$. From 1 January 2022 these will be $T^{\rho,M} = 18\%$ and $T^{\lambda,M} = 8\%$.³⁴ A breach, or likely breach, of minimum TLAC is ordinarily treated by resolution authorities as seriously as a breach, or likely breach, of minimum regulatory capital requirements. Regulators take the view that a breach of the TLAC should be addressed immediately, to ensure that sufficient loss-absorbing capacity is available in resolution (FSB (2015b)).

MREL MREL is determined on a case-by-case basis for European banks. The MREL for a bank is determined based on the judgement of the regulator, which is guided by the criteria set out in Article 45(6) of the BRRD. These include a consideration of a bank's: (a) resolvability; (b) loss absorption needs; (c) recapitalisation needs; (d) size; (e) business model; (f) funding model; (g) risk profile and (h) systemic impact. When determining the MREL requirement, the regulator will also avoid MREL shortfalls due

³³As noted in Section 4.3.3.3, the instruments which are eligible to count towards the TLAC measure are discussed in Appendix B.5.1.

³⁴In our opinion it would be clearer if authorities would directly include in the TLAC requirement that the risk-weighted and leverage-based buffer standards must remain usable. Rather than specifying that CET1 equity that is used to meet regulatory buffers is excluded from TLAC. When including the buffers that must remain usable, the updated TLAC requirements as of 1 January 2019 become $T^{\rho,M} = 16\% + \rho_i^{CB}$ and $T^{\lambda,M} = 6\% + \lambda_i^{CB}$.

to eligible bail-inable liabilities B_i (defined in equation 4.1) that it anticipates must be excluded from bail-in on a discretionary basis, due to contagion risks or no-creditor-worse-off (NCWO) considerations for example.³⁵

Even though MREL is determined on a case-by-case basis, a default MREL requirement exists (SRB (2017), EBA (2016)). We will describe this now. A bank must hold sufficiently many MREL eligible instruments M_i so that it can absorb losses and be recapitalised to a level where it sustains market confidence.³⁶ Hence, the minimum MREL amount M_i^M is the sum of the MREL amount required for loss absorption M_i^L and the MREL amount required for recapitalisation M_i^R (SRB (2017)).

MREL Baseline Loss Absorption Amount The baseline loss absorption amount M_i is (subject to some safeguards) set consistently with the sum of the bank’s pillar one capital requirements ρ_i^{P1} , pillar two capital requirements ρ_i^{P2} (the latter is set by the microprudential stress tests, among other determining factors (Farmer et al. (2020))), and risk-weighted buffer standard ρ_i^{CB} .³⁷ The pillar one capital requirements ρ_i^{P1} states that a bank should have at least 8% own funds F_i , defined as

$$F_i = \tilde{E}_i + \tilde{E}_i^{AT1} + \tilde{E}_i^{T2}, \quad (4.6)$$

relative to its risk-weighted assets Ω_i .³⁸ The reason for this default choice for the loss absorption amount is that the pillar I, II and buffer standards already reflect the view of the regulator regarding the unexpected losses an institution should be able to absorb (EBA (2015)). This baseline loss absorption amount M_i has two safeguards. The first safeguard is that the loss absorption amount can never be lower than buffer standards computed using the Basel I standardised approach BI_i (the so-called ‘Basel I floor’³⁹).⁴⁰ The second safeguard is that the loss absorption amount can never be less than the bank’s leverage requirement λ_i^M (and the bank’s leverage buffer λ_i^{CB} ⁴¹).

³⁵See: Article 44(3) of the BRRD.

³⁶As noted in Section 4.3.3.3, the instruments which are eligible to count towards the MREL measure are discussed in Appendix B.5.1.

³⁷See: Article 1 of Directive 2014/59/EU.

³⁸More specifically, the pillar one capital requirement ρ_i^{P1} states that a bank should have at least: (a) 4.5% CET1 equity \tilde{E}_i relative to its risk-weighted assets Ω_i ; (b) 6% CET1 equity \tilde{E}_i and AT1 equity \tilde{E}_i^{AT1} relative to its risk-weighted assets Ω_i ; and (c) 8% own funds F_i (i.e. CET1 equity \tilde{E}_i , AT1 equity \tilde{E}_i^{AT1} and T2 equity \tilde{E}_i^{T2}) relative to its risk-weighted assets Ω_i .

³⁹See: Article 500 of Regulation (EU) No 575/2013

⁴⁰See: Article 1 of Directive 2014/59/EU.

⁴¹Article 1(e) of Directive 2014/59/EU does not specify whether ‘any applicable leverage requirement’ includes a leverage buffer standard λ_i^{CB} proposed by Basel III (FSB (2017)) or solely refers to the minimum leverage requirement λ_i^M . This leverage ratio buffer λ_i^{CB} will apply to G-SIBs from 1 January 2018 onwards and will be set at 50% of a G-SIB’s higher-loss absorbency risk-weighted requirements (see Farmer et al. (2020)). In our opinion this leverage buffer should be included for consistency with the risk-weighted loss absorbing requirements, where risk-weighted buffers are also taken into account.

Article 45 of the BRRD stipulates that the denominator of the MREL requirement should be equal to the sum of the bank's own funds F_i and liabilities L_i . The EBA has proposed to alter the denominator of the MREL to the bank's risk exposure amount Ω_i , to make it consistent with the TLAC requirement (EBA (2015)).

Thus, the baseline MREL loss absorption amount under the BRRD approach is given by⁴²⁴³

$$M_i^L := \max\{(\rho_i^{P1} + \rho_i^{P2} + \rho_i^{CB})\Omega_i, BI_i, (\lambda^M + \lambda_i^{CB})\hat{A}_i\} \frac{1}{L_i + F_i}, \quad (4.7)$$

where the pillar I requirement ρ_i^{P1} (i.e. the sum of the amount of CET1 equity, AT1 capital, and T2 capital a bank should hold relative to its risk exposure Ω_i) is set to 8%, the pillar II requirement is set based on the stress test result among others, and the combined buffer standard ρ_i^{CB} is set based on the Basel III standards. Further, \hat{A}_i denotes the leverage exposure.

MREL Baseline Recapitalisation Amount The baseline MREL recapitalisation amount M_i^R differs depending on whether the institution is a G-SIB that will likely be bailed in, a small firm that will likely be liquidated, or medium-sized bank that may have less strong recapitalisation needs than a G-SIB (EBA (2015)). If a bank is expected to be liquidated, then recapitalisation is not necessary. Hence, in that case the recapitalisation amount M_i^R is set to zero (i.e. $M_i^R = 0$). In the case where the bank is likely to be resolved via a bail-in, the recapitalisation amount is set such that (subject to the previously mentioned safeguards) the bank meets the conditions for continued authorisation (i.e. the pillar I and pillar II requirement ρ_i^{P1}, ρ_i^{P2}) and sustains market confidence (i.e. it has buffers on top of its requirements). Market confidence is believed to be regained by recapitalising a bank beyond its requirements by an amount equal to the combined buffer standards ρ_i^{CB} (EBA (2015)). Hence, for G-SIBs the baseline recapitalisation amount M_i^R equals the baseline loss absorption amount M_i^L , that is, $M_i^R = M_i^L$. Medium-sized firms typically only need a fraction $r_i \in [0, 1]$ of the recapitalisation amount that a G-SIB needs. Hence, for such firms $M_i^R = r_i M_i^L$. Thus, the baseline MREL requirement M_i^M equals

⁴²Since we do not model the distinction between Basel III and Basel I risk weights, and neither use a microprudential stress test to set the pillar II requirement ρ_i^{P2} , the formula for the baseline loss absorption amount simplifies to $M_i^L = \max\{(\rho_i^{P1} + \rho_i^{CB})\Omega_i, (\lambda^M + \lambda_i^{CB})\hat{A}_i\}$.

⁴³Under the EBA's approach, using the risk exposure Ω_i as the denominator, we would have $M_i^L = \max\{(\rho_i^{P1} + \rho_i^{P2} + \rho_i^{CB})\Omega_i, BI_i, (\lambda^M + \lambda_i^{CB})\hat{A}_i\} \frac{1}{\Omega_i}$.

$$M_i^M := M_i^L + M_i^R = \begin{cases} M_i^L, & \text{if } i \text{ will be liquidated;} \\ 2M_i^L, & \text{if } i \text{ is a G-SIB that will be bailed in;} \\ (1 + r_i)M_i^L, & \text{if } i \text{ will be resolved with lesser recapitalisation needs.} \end{cases} \quad (4.8)$$

where M_i^L is defined in equation 4.7.⁴⁴

As is the case for the TLAC loss absorbing requirement, CET1 equity \tilde{E}_i that is counted towards the regulatory risk-weighted buffer ρ_i^{CB} and regulatory leverage buffer λ_i^{CB} cannot be counted towards MREL (EBA (2016)). This means that a bank must always make sure to have an amount of eligible MREL instruments M_i such that


$$M_i \geq M_i^M (F_i + L_i) + \max\{\rho_i^{CB} \Omega_i, \lambda_i^{CB} \hat{A}_i\} \quad (4.9)$$

holds.⁴⁵ Similar to a TLAC breach, a breach of MREL is treated as seriously as a breach of the minimum capital requirements (EBA (2015)).


Figure 4.6b shows how much bail-inable debt \hat{B}_i each bank participating in the 2018 EBA stress test holds relative to its MREL, and if applicable, TLAC requirement. We use the most binding of a bank's loss absorbing constraints to determine the percentage of its bail-inable that that is held by non-banks. We bring to the fore that we may overestimate the degree to which a bank exceeds its requirements, because we use the eligible debt in bail-in \hat{B}_i as a numerator, which counts eligible debt with a time to maturity greater than 7 days. Rather than using the TLAC T_i or MREL M_i eligible instruments, which only include debt with a time to maturity greater than 1 year.

⁴⁴EBA (2016) indicates that banks typically must at least hold an amount of MREL instruments M_i equal to 8% of liabilities and own funds $L_i + F_i$. This condition could be added to equation 4.8.

⁴⁵In our opinion it would be clearer if authorities would directly include the risk-weighted and leverage-based buffer standards that must remain usable, rather than specifying that CET1 equity that is used to meet regulatory buffers is excluded from MREL. When including the buffers that must remain usable, the updated minimum loss absorbing amount M_i^L as defined in equation 4.7 becomes $M_i^L := \max\{(\rho_i^{P1} + \rho_i^{P2} + 2 * \rho_i^{CB})\Omega_i, BI_i, (\lambda^M + 2 * \lambda_i^{CB})\hat{A}_i\} \frac{1}{L_i + F_i}$. The updated minimum loss absorbing amount M_i^L will then also feed into the an updated minimum recapitalisation amount M_i^R via equation 4.8.



(a)



(b)

Figure 4.6: Shows the amount of bail-inable debt that banks who participated in the 2018 EBA stress test have relative to their TLAC and MREL standards. The bars plot the size of a bank’s bail-inable debt \hat{B}_i (see equation 4.1) relative to its applicable loss absorption minima. These are relative to: (a) risk-weighted assets Ω_i (blue); (b) asset exposure \hat{A}_i (orange); and (c) own funds F_i (see equation 4.6) and liabilities L_i (green). The diamond squares denote the bank’s regulatory minimum for each ratio. Only banks that are G-SIBs face TLAC requirements.

4.5.1.3 Failure Threshold (First ‘Primary’ Bail-In Parameter)

While we have introduced the conditions under which a bank is bailed-in in Section 4.3.1, we have not yet formalised the failure threshold. We will do this here.

Section 4.3.1 specified that a SIB eligible for resolution should be bailed-in if it is deemed to be FLTF. Determining when a bank is FLTF is not an exact science, and regulatory guidelines to determine this point differ across jurisdictions. For instance, a bank that is resolved under Title II of the Orderly Liquidation provision of the Dodd-Frank Act is presumably bailed in when it is (expected to be) insolvent (Goodhart & Avgouleas (2016)), whereas a bank that is resolved under the BRRD is likely bailed in when it breaches its capital requirements, in line with the Basel III recommendations. Further, regulators also have discretion to determine the FLTF point on a case-by-case basis.

Although the FLTF point is not set in stone, as a modeller you have to make a determination for the FLTF. In line with Hüser et al. (2017), we assume a bank is bailed in the first time $t \in \mathbb{N}$ (where \mathbb{N} is the set of natural numbers) when its risk-weighted capital ratio ρ_i , given by its equity \tilde{E}_i over its risk-weighted assets Ω_i , that is

$$\rho_i = \frac{\tilde{E}_i}{\Omega_i}, \quad (4.10)$$

falls below its FLTF trigger ρ_i^F .⁴⁶ The bail-in time τ_i of bank i is given by

$$\tau_i = \inf\{t : \rho_i^t < \rho_i^F\}. \quad (4.11)$$

The bail-in time τ_i signifies the state of the balance sheet at the start of bail-in, whereas the sub-time steps τ_i^a and τ_i^b denote the state of the balance sheet right after phase a and b are completed, which, as explained in Section 4.3.3, is assumed to be so in the same time step τ_i .

In preparation for Section 4.5.2.2, which discusses the valuation of bail-inable debt, we also need to define the bail-in time τ_i^{nm} in Monte Carlo run n (out of N runs in total) applicable to a bail-inable contract issued by bank i with a time to maturity m . This is given by

$$\tau_i^{m,n} = \inf\{s \leq t + m : \rho_i^{s,n} < \rho_i^F\}. \quad (4.12)$$

⁴⁶In Section 4.7, we will evaluate how the choice of the FLTF trigger ρ_i^F affects financial stability.

⁴⁷Alternatively, we could have also defined the bail-in time τ_i as the first time the bank falls below its trigger for at least one of its capital ratios: its risk-weighted capital ratio ρ_i or its leverage ratio λ_i , where the leverage ratio is defined as the bank’s CET1 equity \tilde{E}_i over its asset exposure \tilde{A}_i . That is $\tau_i := \inf\{t \in \mathbb{Z} : \rho_i^t < \rho_i^F \text{ or } \lambda_i^t < \lambda_i^F\}$. For simplicity, we do not assume use this bail-in trigger.

Equation 4.12 shows that a bail-in in run n will only matter for the contract's payoff whenever bail-in takes place before the contract's maturity T (since $t + m = t + T - t = T$).⁴⁸

For simplicity, we will omit the i -subscript or m -superscript whenever the bank or contract in question is clear.

4.5.1.4 Loss Absorption and Recapitalisation through Haircuts

Once a bank has met the FLTF condition (see equation 4.11), it will be bailed in. The bail-in starts with applying haircuts to absorb losses and to recapitalise the bank to a recapitalisation target – as we will discuss next.

To absorb losses in phase a and to recapitalise the bank in phase b , haircuts are applied to the bank's bail-inable debt B_i . The total haircut applied to bank i 's bail-inable debt B_i^k in priority class k is given by h_i^k . This in turn is given by the sum of the haircut applied in the loss absorption phase $h_i^{a,k}$ and the haircuts applied in the recapitalisation phase $h_i^{b,k}$. That is, the total haircut h_i^k in priority class k is given by

$$h_i^k := h_i^{k\tau_a} + h_i^{k\tau_b}. \quad (4.13)$$

Extra sub- or superscripts are added to equation 4.13 to denote contract-specific haircuts. For instance, $h_{ji}^{km\tau_a}$ denotes the haircut applied in phase a to a bail-inable debt contract $B_{ji}^{km\tau}$ issued by bank i and held by institution j in priority class k with time to maturity m .

We proceed to discuss the necessary loss absorption and recapitalisation amount to revitalise the bank and the associated haircuts to affect this.

Loss Absorption Amount Phase a applies whenever the bank is insolvent (i.e. $\tilde{E}_i^\tau < 0 \implies \rho_i^\tau < 0$) at the start of bail-in τ . In that case, the CET1 equity of the bank \tilde{E}_i^τ , in other words priority class k_1 , is not sufficient to absorb the asset losses necessitating haircuts to higher priority classes $k \in \mathcal{K}$ to fully absorb the losses. The total loss \hat{l}_i^τ that *needs* to be absorbed is given by the degree of insolvency, that is

$$\hat{l}_i^\tau := |\tilde{E}_i^\tau| \mathbb{1}_{\{\rho_i^\tau < 0\}} = \begin{cases} 0, & \text{if } \rho_i^\tau \geq 0 \text{ (no phase } a\text{);} \\ > 0, & \text{if } \rho_i^\tau < 0 \text{ (phase } a \text{ applies).} \end{cases} \quad 49 \quad (4.14)$$

⁴⁸If a bail-in in run n does not occur at or prior to the contract's maturity T , then we set the bail-in time $\tau_i^{m,n}$ to infinity by convention.

⁴⁹Obviously, in reality the CET1 equity \tilde{E}_i cannot become negative. However, by our definition of the CET1 equity \tilde{E}_i (see equation B.2), an asset value A_i less than $L_i - \Delta_i^{t_0}$ implies a negative CET1 equity signifying that asset losses cannot be absorbed by CET1 equity \tilde{E}_i .

The total loss that *can* be absorbed by the bank's bail-inable debt B_i (see equation 4.2) is given by

$$l_i^\tau := \min\{B_i^\tau, \hat{l}_i^\tau\} \quad (4.15)$$

The loss absorbing requirements (see Section 4.5.1.2) aim to ensure that a bank has sufficient bail-inable debt B_i to avoid situations where the losses cannot be fully absorbed (i.e. $B_i^\tau < \hat{l}_i^\tau$). If the bail-inable debt B_i turns out to be insufficient to absorb losses in phase a , then the bank remains insolvent. In such cases, regulators need to decide whether to save the bank or let it fail. In our experiments (see Section 4.7) we either liquidate the bank or recapitalise the bank via a bail-out (see details in Appendix B.3.1). Also, perhaps needless to say, we terminate bail-in, so that phase b is not applied. In summary, for a given loss \hat{l}_i^τ that *needs* to be absorbed, there are two cases

$$\hat{l}_i^\tau = \begin{cases} > B_i^\tau, & \implies \text{the bank remains insolvent and will be liquidated or bailed out;} \\ \leq B_i^\tau, & \implies \text{the bank has sufficient bail-inable debt to absorb losses.} \end{cases} \quad (4.16)$$

In any case, after phase a , bank i 's CET1 equity \tilde{E}_i^τ updates to

$$\tilde{E}_i^{\tau a} = \begin{cases} \tilde{E}_i^\tau \geq 0, & \text{if } \rho_i^\tau > 0 \text{ (CET1 equity unchanged when phase } a \text{ is not necessary);} \\ 0, & \text{if } \hat{l}_i^\tau \leq B_i^\tau \text{ (successful phase } a \text{ eliminates insolvency);} \\ < 0, & \text{if } \hat{l}_i^\tau > B_i^\tau \text{ (unsuccessful phase } a \text{ does not eliminate insolvency).} \end{cases} \quad (4.17)$$

From equation 4.26 we clearly observe that a successful phase a means that the insolvency is eliminated but nothing more: the CET1 equity \tilde{E}_i is not lifted above 0 as this is the task of phase b (see Section 4.5.1.4).

We proceed to formalise how the *feasible* loss absorption l_i^τ (see equation 4.15) is achieved through applying haircuts.

Haircuts in Loss Absorption Phase Loss absorption l_i^τ is achieved by applying haircuts $h_i^{k\tau a}$ according to the hierarchy of bail-inable claims and proportional within each priority class k . Let us specify this procedure. Let k_b be the highest priority class in the hierarchy of bail-inable claims at time τ that the loss l_i^τ fully exceeds, that is, k_b is defined as

$$k_b := \max\{x \in \mathcal{K} : l_i^\tau \geq \sum_{k=2}^x B_i^{k\tau}\}. \quad (4.18)$$

By implication, k_{b+1} is the first priority class that the loss l_i^τ does not fully exceed and $k \geq k_{b+2}$ are the priority classes which are unaffected by phase a . Using equation 4.18, the total haircuts $h_i^{k\tau_a}$ that befall priority classes $k \in \mathcal{K}$ in phase a are thus given by

$$h_i^{k\tau_a} = \begin{cases} B_i^{k\tau}, & \text{for } k = 2, \dots, k_b; \\ l_i^\tau - \sum_{k=2}^{k_b} B_i^{k\tau}, & \text{for } k = k_{b+1}; \\ 0, & \text{for } k \geq k_{b+2}. \end{cases} \quad (4.19)$$

Due to the application of the haircuts $h_i^{k\tau_a}$, the bail-inable debt in priority classes $k \in \mathcal{K}$ updates after the completion of phase a to

$$B_i^{k\tau_a} = \begin{cases} B_i^{k\tau} - h_i^{k\tau_a} = 0, & \text{for } k = 2, \dots, k_b; \\ B_i^{k\tau} - h_i^{k\tau_a} > 0, & \text{for } k = k_{b+1}; \\ B_i^{k,\tau}, & \text{for } k \geq k_{b+2}. \end{cases} \quad (4.20)$$

Given the proportional distribution of the total haircuts $h_i^{k\tau_a}$ over individual contracts $B_{ji}^{km\tau}$ in a priority class $k \in \mathcal{K}$ (discussed in Section 4.5.1.1), the contract-specific haircuts $h_{ji}^{km\tau_a}$ in phase a are given by

$$h_{ji}^{km\tau_a} = \begin{cases} h_i^{k\tau_a} \frac{B_{ji}^{km\tau}}{B_i^{k\tau}} = B_{ji}^{km\tau}, & \text{for } k = 2, \dots, k_b; \\ h_i^{k\tau_a} \frac{B_{ji}^{km\tau}}{B_i^{k\tau}}, & \text{for } k = k_{b+1}; \\ 0, & \text{for } k \geq k_{b+2}. \end{cases} \quad (4.21)$$

As a consequence of the haircuts $h_{ji}^{km\tau_a}$, the value of bail-inable debt contracts $B_{ji}^{km\tau_a}$ in priority classes $k \in \mathcal{K}$ update in phase a to

$$B_{ji}^{km\tau_a} = \begin{cases} 0, & \text{for } k = 2, \dots, k_b; \\ B_{ji}^{km\tau} - h_{ji}^{km\tau_a}, & \text{for } k = k_{b+1}; \\ B_{ji}^{km\tau}, & \text{for } k \geq k_{b+2}. \end{cases} \quad (4.22)$$

Recapitalisation Target (Second ‘Primary’ Bail-In Parameter) After the losses have been absorbed (see Section 4.5.1.4) by applying haircuts in phase a (see Section 4.5.1.4), the bank must be recapitalised to a target capital ratio ρ_i^T at which the bank should be viable again. As explained in Section 4.3.3, regulators believe a bank to be viable again when a bank’s capital ratio ρ_i is comparable to that of its peers or, at a minimum, satisfies its authorisation conditions (e.g. its minimum capital requirements).

Given a recapitalisation target ρ_i^T , the desired target recapitalisation amount \hat{b}_i is implied by

$$\rho_i^T = \frac{\tilde{E}_i^{\tau_a} + \hat{b}_i}{\Omega_i^{\tau_a}}, \quad (4.23)$$

where $\Omega_i^{\tau_a} = \Omega_i^{\tau}$, since the risk-weighted assets Ω_i are not affected by bail-in. Equation 4.23 informs how much CET1 equity (i.e. \hat{b}_i) must be added (or in other words, how many bail-inable debt must receive haircuts) to the bank's phase- a equity $\tilde{E}_i^{\tau_a}$ to hit the recapitalisation target ρ_i^T . By rewriting equation 4.23, we find that the *necessary* target recapitalisation amount is given by

$$\hat{b}_i = \rho_i^T \Omega_i^{\tau_a} - \tilde{E}_i^{\tau_a} = (\rho_i^T - \rho_i^{\tau_a}) \Omega_i^{\tau_a}. \quad (4.24)$$

Equation 4.24 shows that the necessary recapitalisation target amount \hat{b}_i is larger if the recapitalisation target ρ_i^T is set higher relative to the post-phase- a capital ratio $\rho_i^{\tau_a}$.

The resolution authority is unable to fully recapitalise the bank if its bail-inable debt $B_i^{\tau_a}$ after the completion of phase a is insufficient. The *feasible* recapitalisation amount b_i is thus capped by $B_i^{\tau_a}$ as shown below

$$b_i = \min\{\hat{b}_i, B_i^{\tau_a}\}. \quad (4.25)$$

If the phase- a bail-inable debt $B_i^{\tau_a}$ is insufficient to fully recapitalise the bank ($\hat{b}_i > B_i^{\tau_a}$), then the bank's phase- b capital ratio will be below its target ($\rho_i^{\tau_b} < \rho_i^T$). In such case, the bank could either satisfy its authorisation conditions ($\rho^{T2,M} \leq \rho_i^{\tau_b} < \rho_i^T$) or not ($\rho_i^{\tau_b} < \rho^{T2,M} \leq \rho_i^T$). If it does not, the regulator must take further action, for instance, to either liquidate the bank or to further recapitalise the bank via a bail out. If the bail-in target is not reached but the bank does satisfy its authorisation conditions, then further regulatory action to fully recapitalise the bank is optional (see Appendix B.3.1). To avoid a scenario where the bank has insufficient bail-inable debt to fully recapitalise, total loss absorbing requirements have been introduced (see Section 4.5.1.2). The total loss absorbing requirements thus both serve to ensure that loss absorption (as explained in 4.5.1.4) and recapitalisation can be successfully completed. In summary,

$$\rho_i^{\tau_b} = \begin{cases} \rho_i^T, & \text{if } \hat{b}_i \leq B_i^{\tau_a} \text{ (phase } b \text{ successful);} \\ < \rho_i^T, & \text{if } \hat{b}_i > B_i^{\tau_a} \text{ (phase } b \text{ not or partially successful).} \end{cases} \quad (4.26)$$

Irrespective of whether phase b is fully successful, the bank's phase- b CET1 equity $\tilde{E}_i^{\tau_b}$ is given by

$$\tilde{E}_i^{\tau_b} = \tilde{E}_i^{\tau_a} + b_i. \quad (4.27)$$

Haircuts in the Recapitalisation Phase Let k_r be the first priority class in the hierarchy of bail-inable claims at time τ_a which the recapitalisation amount b_i exceeds, that is

$$k_r := \max\{x \in \mathcal{K} : b_i \geq \sum_{k=2}^x B_i^{k\tau_a}\}. \quad (4.28)$$

Then the total haircuts $h_i^{k\tau_b}$ that befall priority classes $k \in \mathcal{K}$ in phase b are given by

$$h_i^{k\tau_b} = \begin{cases} B_i^{k\tau_a}, & \text{for } k = 2, \dots, k_r; \\ b_i - \sum_{k=2}^{k_r} B_i^{k\tau_a}, & \text{for } k = k_{r+1}; \\ 0, & \text{for } k \geq k_{r+2}. \end{cases} \quad (4.29)$$

Due to the application of the haircuts $h_i^{k\tau_b}$, the bail-inable debt in priority classes $k \in \mathcal{K}$ updates after the completion of phase b to

$$B_i^{k\tau_b} = \begin{cases} B_i^{k\tau_a} - h_i^{k\tau_b} = 0, & \text{for } k = 2, \dots, k_r; \\ B_i^{k\tau_a} - h_i^{k\tau_b} > 0, & \text{for } k = k_{r+1}; \\ B_i^{k\tau_a}, & \text{for } k \geq k_{r+2}. \end{cases} \quad (4.30)$$

The contract-specific haircuts $h_{ji}^{km\tau_b}$ for priority classes $k \in \mathcal{K}$ in phase b are given by

$$h_{ji}^{km\tau_b} = \begin{cases} h_i^{k\tau_b} \frac{B_{ji}^{km\tau_a}}{B_i^{k\tau_a}} = B_{ji}^{km\tau_a}, & \text{for } k = 2, \dots, k_r; \\ h_i^{k_{r+1}\tau_b} \frac{B_{ji}^{k_{r+1}m\tau_a}}{B_i^{k_{r+1}\tau_a}}, & \text{for } k = k_{r+1}; \\ 0, & \text{for } k \geq k_{r+2}. \end{cases} \quad (4.31)$$

Due to the application of the contract-specific haircuts $h_{ji}^{km\tau_b}$, the value of generic bail-inable debt contracts in priority classes $k \in [2, 5]$ update after the completion of phase b to

$$B_{ji}^{km\tau_b} = \begin{cases} 0, & \text{for } k = 2, \dots, k_r; \\ B_{ji}^{km\tau_a} - h_{ji}^{km\tau_b}, & \text{for } k = k_{r+1}; \\ B_{ji}^{km\tau_a}, & \text{for } k \geq k_{r+2}. \end{cases} \quad (4.32)$$

We now proceed to discuss the conversion rates offered as a compensation for the haircuts incurred in phase a and phase b .

Debt-to-Equity Conversion Rates (Third ‘Primary’ Set of Bail-In Parameters) While we introduced the debt-to-equity conversion rates – and the Principles for setting these – in Section 4.3.3.2, we have not yet modelled it. We proceed by formally defining the conversion rates. Then we unfold what the implications of two Principles (NCWO and Preservation of Hierarchy of Claims) are in terms of setting ‘fair’ or ‘unfair’ conversion rates.

It turns out that it is much easier to satisfy both Principles when one splits the total conversion rate Δ_i^k in: (1) the conversion rate that applies to haircuts h^{τ_a} incurred in the loss absorption phase Δ_{ia}^k ; and (2) the conversion rate that applies to haircuts h^{τ_b} incurred in the recapitalisation phase Δ_{ib}^k . The total conversion rate Δ_i^k relates to the conversion rate in the loss absorption phase Δ_{ia}^k and the recapitalisation phase Δ_{ib}^k as follows

$$\Delta_i^k = \frac{\Delta_{ia}^k h_{ji}^{km\tau_a} + \Delta_{ib}^k h_{ji}^{km\tau_b}}{h_{ji}^{km}}, \quad (4.33)$$

where we recall from equation 4.13 that the total contract-specific haircut h_{ji}^{km} is given by the sum of the phase-*a* haircut $h_{ji}^{km\tau_a}$ and phase-*b* haircut $h_{ji}^{km\tau_b}$. Equation 4.33 tells that the total conversion rate Δ_i^k is given by the number of shares a creditor receives per unit haircut applied to the principle amount of a bail-inable claim B_{ji}^{km} in priority class *k*, which, is in accordance with the definition of the debt-to-equity conversion rate Δ_i^k given at the start of this section.

It will also be useful to split the conversion rate in the ‘fair’ conversion rate $\tilde{\Delta}^k$ and the ‘unfair’ conversion rate $\hat{\Delta}_i^k$ for both phase *a* and *b* (to be specified further in the next sections). That is,

$$\Delta_{ia}^k = \tilde{\Delta}_a + \hat{\Delta}_{ia}^k; \quad (4.34)$$

$$\Delta_{ib}^k = \tilde{\Delta}_b + \hat{\Delta}_{ib}^k. \quad (4.35)$$

When conversion rates are applied in phase *a* and *b*, a haircutted creditor in priority class *k* receives a share $\epsilon_{ji}^{km\tau_b}$ of bank *i*’s CET1 equity $\tilde{E}_i^{\tau_b}$. So shares $\epsilon_{ji}^{km\tau_b}$ and conversion rates are linked by the following formula

$$\epsilon_{ji}^{km\tau_b} = \frac{\Delta_{ia}^k h_{ji}^{km\tau_a} + \Delta_{ib}^k h_{ji}^{km\tau_b}}{\eta_i^{\tau_b}}, \quad (4.36)$$

where $\eta_i^{\tau_b}$ denotes the number of outstanding shares of bank *i*’s CET1 equity \tilde{E}_i after phase *b*. The phase-*b* number of CET1 equity shares $\eta_i^{\tau_b}$ of bank *i* is given by the sum of the existing CET1 equity shares after the application of the haircuts in the loss absorption phase $\eta_i^{\tau_a}$ (see equation 4.70) and the newly created shares in phase *a* and *b*, that is

$$\eta_i^{\tau_b} = \eta_i^{\tau_a} + \sum_{k=2}^{k_b} \Delta_{ia}^k B_i^{k\tau} + \Delta_{ia}^{k_b+1} (l_i^{\tau} - \sum_{k=2}^{k_b} B_i^{k\tau}) + \sum_{k=2}^{k_r} \Delta_{ib}^k B_i^{k\tau_a} + \Delta_{ib}^{k_r+1} (b_i - \sum_{k=2}^{k_r} B_i^{k\tau_a}), \quad (4.37)$$

where we recall that k_b and k_r are defined in equation 4.18 and 4.28.

As the phase-specific conversion rate has been split in the fair and unfair part (see

equation 4.35) we can split up the received share $\epsilon_{ji}^{km\tau_b}$ in the fair $\tilde{\epsilon}_{ji}^{km\tau_b}$ and unfair $\hat{\epsilon}_{ji}^{km\tau_b}$ part. That is,

$$\epsilon_{ji}^{km\tau_b} = \tilde{\epsilon}_{ji}^{km\tau_b} + \hat{\epsilon}_{ji}^{km\tau_b}. \quad (4.38)$$

The value of j 's claim E_{ji}^{kmt} is given by its share $\epsilon_{ji}^{km\tau_b}$ of bank i 's CET1 equity \tilde{E}_i^t , that is

$$E_{ji}^{kmt} = \epsilon_{ji}^{km\tau_b} \tilde{E}_i^t, \quad (4.39)$$

where at time τ_b we have $t = \tau_b$. From equation 4.39 it is clear that j 's claim E_{ji}^{kmt} of a fixed share $\epsilon_{ji}^{km\tau_b}$ of equity changes over time t as bank i 's CET1 equity \tilde{E}_i^t changes. This is relevant, because it means that converted debt to equity can revalue downwards in the time period after the bail-in has been completed, even if no (significant) net losses were suffered due to the conversion itself.

We next discuss the two conversion rate principles imply for setting debt-to-equity rates. We consider two extremes that may satisfy these principles: ‘fair’ and ‘unfair’ conversion rates. To be best of our knowledge, we are the first to formally (using formulas) work out the implications of these principles.

Fair Conversion Rates The phase-specific conversion rates are set fairly when

$$\text{Fair Conversion Rate in Phase a : } \Delta_{ia}^k = \tilde{\Delta}_a = 0 \text{ for } k \in \mathcal{K}; \quad (4.40)$$

$$\text{Fair Conversion Rate Phase b : } \Delta_{ib}^k = \tilde{\Delta}_b > 0 \text{ for } k \in \mathcal{K}, \quad (4.41)$$

where the fair conversion rate in phase- b is given by equation 4.42.

‘Fair’ conversion rates in the loss absorption phase

Splitting the phase-specific conversion rate in its fair (and unfair) part is relevant, because there is typically only one way to set the conversion rates Δ_{ia}^k in phase a without breaching principle I and II (as will be come clear in the next Section 4.5.1.4 on ‘unfair’ conversion rates and is argued in Appendix B.5.2). This is to set Δ_{ia}^k equal to the ‘fair’ conversion rate $\tilde{\Delta}_a$ in phase a (see equation 4.40), which corresponds to a pure write-down. The intuition behind this is as follows.

Principle I is not breached when conversion rates are set fairly, because creditors who face pure write-downs in a bail-in would have also faced pure write-downs (i.e. losses) in a liquidation, so are not ‘worse off’. Furthermore, applying positive conversion rates

$\Delta_{ia}^k > 0$ would risk making more senior creditors worse off in bail-in than in a liquidation (violating principle I), since allocating shares is a zero-sum game: if some creditors are allocated a larger number of shares, others will have a smaller fractional share ϵ_{ji} of a bank i 's equity (as is evident from equation 4.36). Hence, any other conversion rate but pure write-downs violate principle I.

Principle II holds if conversion rates in phase a are set fairly (see equation 4.41 and 4.40). Under fair conversion rates, the affected creditors in phase- a , who face pure write-downs, suffer larger net losses than affected creditors in phase- b , who receive an debt-to-equity conversion. Since creditors in phase a are always in lower (or equal) priority classes in the hierarchy of bail-inable claims than creditors in phase b , principle II stands (see details in Appendix B.5.2).⁵⁰

'Fair' conversion rates in the recapitalisation phase

There is also typically only one way to set conversion rate in priority class k in the recapitalisation phase Δ_{ib}^k without breaching principle I and II. This is to set it equal to the 'fair' conversion rate $\tilde{\Delta}_b$ given by

$$\tilde{\Delta}_b = \frac{\eta_i^{\tau_a}}{\tilde{E}_i^{\tau_a}} \mathbb{1}_{\{\tilde{E}_i^{\tau_a} > 0\}} + w \mathbb{1}_{\{\tilde{E}_i^{\tau_a} = 0\}}. \quad (4.42)$$

Equation 4.42 shows that the fair conversion rate in the recapitalisation phase $\tilde{\Delta}_b$ is equal to the number of shares per unit of phase- a equity ($\frac{\eta_i^{\tau_a}}{\tilde{E}_i^{\tau_a}}$) in case equity holders were not fully diluted in phase a (i.e. if phase a was not necessary since $\tilde{E}_i^{\tau} = \tilde{E}_i^{\tau_a} > 0$, see equation 4.14) and is equal to any desired number w in case equity holders were fully diluted after phase a . The idea is that by doing so, non-wiped-out equity holders are diluted by a 'fair' amount and not more. Any conversion rate may be chosen if existing equity holders are diluted.

When the conversion rate $\tilde{\Delta}_b$ is fair, then for any haircut $h_{ji}^{km\tau_b}$ a creditor receives in phase b it will receive an equal value $E_{ji}^{km\tau_b}$ of the bank's phase b CET1 equity $\tilde{E}_i^{\tau_b}$ (note that the value of the claim may change after the bail-in as is clear from equation 4.39). The conversion rate $\tilde{\Delta}_b$ is 'fair' in the sense that a haircut is compensated with an equal equity claim,⁵¹ so that no net losses are suffered as a direct consequence of the bail-in.

⁵⁰ In line with Hüser et al. (2017), we will by default set the phase- a conversion rates Δ_{ia}^k fairly (i.e. $\Delta_{ia}^k = \tilde{\Delta}_a$, $\forall k \in \mathcal{K}$, see equation 4.40), unless otherwise indicated. Hence, creditors will not receive any share as a compensation of their haircuts in phase a (i.e. $\epsilon_{ji}^{km\tau_a} = 0$ for $k \in \mathcal{K}$).

⁵¹ Chen et al. (2013) call the conversion rate of a CoCo 'fair' if the write-down is met an an equal value of the equity claim right after the conversion.

Roughly speaking, the reason why fair conversion rates typically do not breach principle I and II is as follows. Recapitalisation does not take place in liquidation, so haircuts must be compensated with an equal value of equity to ensure that creditors are ‘worse off’ in a bail-in than in a liquidation. Principle II also holds, as has already been explained in Appendix B.5.2).

Let us now formalise the above-mentioned. The ‘fair’ share of equity $\tilde{\epsilon}_{ji}^{km\tau_b}$, which a creditor who received a haircut $h_{ji}^{km\tau_b}$ in phase b receives of bank i ’s phase- b CET1 equity $\tilde{E}_i^{\tau_b}$ (corresponding to a fair phase- b conversion rate $\tilde{\Delta}_b$, see equation 4.41) is given by

$$\tilde{\epsilon}_{ji}^{km\tau_b} = \frac{h_{ji}^{km\tau_b}}{\tilde{E}_i^{\tau_b}}. \quad (4.43)$$

By multiplying equation 4.43 with the bank’s phase- b CET1 equity $\tilde{E}_i^{\tau_b}$ and using equation 4.39, we obtain the value $E_{ji}^{km\tau_b}$ of the creditor’s acquired equity claim. It is given by

$$E_{ji}^{km\tau_b} = h_{ji}^{km\tau_b}. \quad (4.44)$$

Equation 4.44 shows that the value of the new claim $E_{ji}^{km\tau_b}$ is equal to the value of the haircut $h_{ji}^{km\tau_b}$. Indeed, when the share $\epsilon_{ji}^{km\tau_b}$ is such that is is fair $\tilde{\epsilon}_{ji}^{km\tau_b}$ (see equation 4.43), then as a direct consequence of bail-in, a creditor who received a haircut in the recapitalisation phase $h_{ji}^{km\tau_b}$ does not suffer net losses, which can be understood from the equations below

$$E_{ji}^{km\tau_b} + B_{ji}^{km\tau_b} = B_{ji}^{km\tau_a} \iff \quad (4.45)$$

$$\tilde{\epsilon}_{ji}^{km\tau_b} \tilde{E}_i^{\tau_b} + (B_{ji}^{km\tau_a} - h_{ji}^{km\tau_b}) = B_{ji}^{km\tau_a}. \quad (4.46)$$

Equation 4.45 gives the no-net-loss condition for phase b , which says that the value of a creditor’s claim in phase a and b must be equal, so that no net losses are suffered due to phase b . In other words, the phase- b haircutted claim $B_{ji}^{km\tau_a}$ ($=B_{ji}^{km\tau_a} - h_{ji}^{km\tau_b}$) plus the equity compensation $E_{ji}^{km\tau_b}$ ($=\tilde{\epsilon}_{ji}^{km\tau_b} \tilde{E}_i^{\tau_b}$) must equal the phase- a claim $B_{ji}^{km\tau_a}$. The only share $\epsilon_{ji}^{km\tau_b}$ that satisfies the no-net-loss requirement (as given by equation 4.46) is the fair share $\tilde{\epsilon}_{ji}^{km\tau_b}$. The relation between the fair conversion rates $\tilde{\Delta}_b$ and $\tilde{\Delta}_a = 0$ and the fair share $\tilde{\epsilon}_{ji}^{km\tau_b}$ is given by equation

$$\tilde{\epsilon}_{ji}^{km\tau_b} = \frac{\tilde{\Delta}_b h_{ji}^{k\tau_b}}{\tilde{\Delta}_b b_i + \eta_i^{\tau_a}}, \quad (4.47)$$

which is a simplification of equation 4.36. By default we will set the conversion rates Δ_{ib}^k in phase b fairly (i.e. $\Delta_{ib}^k = \tilde{\Delta}_b$, for $k \in \mathcal{K}$, see equation 4.41), unless otherwise indicated.

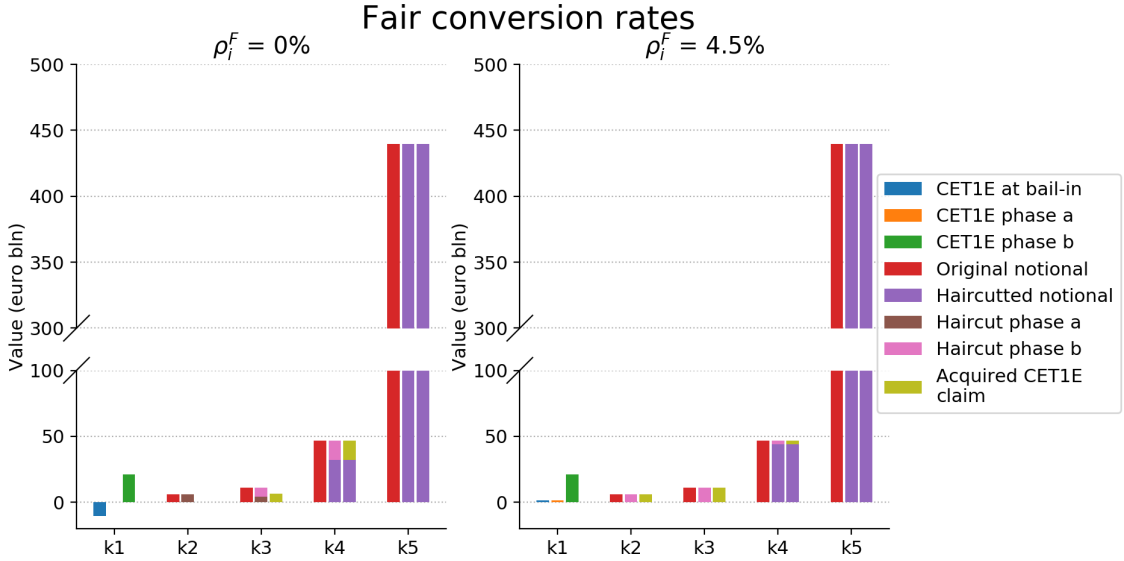


Figure 4.7: This figure provides an example of the application of a bail-in procedure to a bank's bail-inable debt. We take bank $i = FR12$ that partook in the EBA 2018 stress test. In the left plot the bank's assets have fallen below the FLTF ratio of $\rho_i^F = 0\%$ such that its risk-weighted capital ratio is $\rho_i^T = -4\%$ at the start of the bail-in (i.e. at time τ). Suppose that the regulator has determined that the bank should be recapitalised to its original capital ratio, that is, $\rho_i^T = \rho_i^{data}$. Further, we assume that the regulator applies a fair conversion rate in phase a and b (see equation 4.40 and 4.40). We observe the following. The bars associated to priority class k_1 capture the evolution of the bank's CET1 equity \tilde{E}_i over the course of the bail-in: at time τ_i (the start of the bail-in); at time τ_a (after phase a has been completed); and at time τ_b (after phase b has been completed). We see that the loss absorption phase (i.e. phase a) accomplishes wiping out any degree of insolvency by applying haircuts in priority classes that come lowest in hierarchy (k_2 and k_3). Creditors who faced haircuts in phase a suffer net losses that are not compensated with a claim of CET1 equity. Next, the recapitalisation phase (i.e. phase b) accomplishes to recapitalise the bank such that it regains a positive amount of CET1 equity (see green dark bar at k_1), by applying haircuts in phase b . It applies these haircuts to the priority class that is next in line in the hierarchy (i.e. k_4). For any haircut applied in phase b , creditors get an equal amount of CET1 equity in return (see light green bar at k_4 , which is equal in size to the pink bar at k_4), since the regulator's conversion rate is set fairly. In the right plot we show the same result, but now for the case where the regulator set the FLTF trigger at $\rho^F = 4.5\%$. Because of this choice, the bank is already bailed in when its risk-weighted capital ratio falls to $\rho_i^T = 0.5\%$. You can observe that no haircuts are applied in phase a in this case, because no degree of insolvency had to be wiped out. Only haircuts are applied in phase b to recapitalise the bank. We note that since no haircut was applied in phase a , creditors face a zero net loss as a direct consequence of the bail-in; any haircuts in phase b were replaced with equal claims in CET1 equity. In sum, we learn from the left and right plot that creditors only face net losses as a direct consequence of the bail-in when conversion rates are set fairly, whenever the bank is insolvent at the start of bail-in. On the other hand, if the bank is solvent at the start of bail-in each creditor that receives a haircut obtains an equal claim of CET1 equity of bank i , so no net losses are suffered.

Unfair Conversion Rates

Specification of Principle I: No-Creditor-Worse-Off

The no-creditor-worse-off condition, P1, requires that a creditor in a bail-in is at least as well off in a bail-in as in the case where the bank would have been liquidated upon failure (i.e. in a ‘hypothetical liquidation’). That is, for each creditor it must hold that its claim following a bail-in $B_{ji}^{km\tau_b,B}$ is at least as high as its claim in a hypothetical liquidation $B_{ji}^{km\tau_b,L}$, that is,

$$B_{ji}^{km\tau_b,B} \geq B_{ji}^{km\tau_b,L}. \quad (4.48)$$

The claim following a bail-in is given by

$$B_{ji}^{km\tau_b,B} = B_{ji}^{km\tau} - h_{ji}^{km\tau_a} - h_{ji}^{km\tau_b} + \epsilon_{ji}^{km\tau_b} \tilde{E}_i^{\tau_b} \quad (4.49)$$

$$= B_{ji}^{km\tau_b} + \epsilon_{ji}^{km\tau_b} \tilde{E}_i^{\tau_b}, \quad (4.50)$$

From equation 4.50 it can be seen that the claim $B_{ji}^{km\tau_b,B}$ in a bail-in is equal to the pay-off $P_{ji}^{km\tau_b}$ of a bail-inable claim valued at time τ_b (see equation 4.88).

The claim following a liquidation is given by one minus the loss given default (LGD) in liquidation $\zeta_i^{k\tau_b,L}$ times the original claim $B_{ji}^{km\tau}$, that is,

$$B_{ji}^{km\tau_b,L} = (1 - \zeta_i^{k\tau_b,L}) B_{ji}^{km\tau}. \quad (4.51)$$

The LGD $\zeta_i^{k\tau_b,L} \in [0, 1]$ in each priority class k in liquidation is given by

$$\zeta_i^{k\tau_b,L} = \begin{cases} 100\%, & \text{for } k = 2, \dots, k^l; \\ \frac{\theta_i^\tau - (\sum_{k=1}^{k^l} L_i^{k\tau} - L_i^{k^l\tau})}{L_i^{k^l+1,\tau}}, & \text{for } k = k^l + 1; \\ 0\%, & \text{for } k \geq k^l + 2, \end{cases} \quad (4.52)$$

where k^l denotes the first priority class that the loss that can be absorbed θ_i^τ in liquidation does not fully exceed, and is given by

$$k^l := \max\{x : \theta_i^\tau \geq \sum_{k=1}^x L_i^{k\tau} - L_i^{k^l\tau}\} \in k^1 \cup \mathcal{K}. \quad (4.53)$$

$L_i^{k\tau}$, shown in equation 4.53, signifies the value of bank i 's liabilities in priority class k at time τ . Note that $L_i^{k\tau}$ may be different from the bail-inable debt $B_i^{k\tau}$ in priority class k , due to fixed or ad-hoc exclusions of liabilities from the hierarchy of bail-inable claims

(see Section 4.5.1.1 and Section 4.5.1.2). In increasing order of seniority, the liquidation hierarchy (or in other words, the regulator insolvency hierarchy) is given by⁵²

$$\{L_i^{k_1}, L_i^{k_2}, L_i^{k_3}, L_i^{k_4}, L_i^{k_5}\} = \quad (4.54)$$

$$\{\tilde{E}_i, \tilde{E}_i^{AT1}, \tilde{E}_i^{T2}, \tilde{I}_i^{k_4} + \tilde{O}_i^{k_4}, \tilde{I}_i^{k_5} + \tilde{O}_i^{k_5}\} \quad (4.55)$$

The total loss that needs to be absorbed in liquidation, displayed in equation 4.53, is given by

$$\hat{\theta}_i^\tau = \begin{cases} \hat{l}_i^\tau, & \text{if } c_i^\tau = 0\%; \\ \hat{l}_i^\tau + c_i^\tau A_i^\tau \mathbb{1}_{\{\rho_i^\tau \leq 0\}} + \max\{L_i^\tau - A_i^\tau(1 - c_i^\tau), 0\} \mathbb{1}_{\{\rho_i^\tau > 0\}}, & \text{if } c_i^\tau > 0\%, \end{cases} \quad (4.56)$$

From equation 4.56 we observe that the loss that needs to be absorbed in liquidation $\hat{\theta}_i^\tau$ equals the loss that needs to be absorbed in a bail-in \hat{l}_i^τ (see equation 4.14) if the estimated liquidation cost c_i^τ is zero. The loss that needs to be absorbed $\hat{\theta}_i^\tau$ may not be equal to the loss that can be absorbed θ_i^τ by bank i 's liabilities, $L_i^\tau = \sum_{k=2}^7 L_i^k$. The loss that can be absorbed is given by

$$\theta_i^\tau = \min\{\hat{\theta}_i^\tau, \sum_{k \in \mathcal{K}} L_i^{k\tau}\}. \quad (4.57)$$

For Section 4.5.1.2 it will also be useful to define the LGD $\zeta_i^{k\tau a, B} \in [0, 1]$ incurred in the loss absorption phase of a bail-in⁵³; this is given by

$$\zeta_i^{k\tau a, B} = \begin{cases} 100\%, & \text{for } k = 2, \dots, k^b; \\ \frac{l_i^\tau - (\sum_{k=1}^{k^b} B_i^{k\tau} - B_i^{k^1\tau})}{B_i^{k^b+1, \tau}}, & \text{for } k = k^b + 1; \\ 0\%, & \text{for } k \geq k^b + 2, \end{cases} \quad (4.58)$$

where k^b signifies the first priority class that the feasible loss absorption in phase a of a bail-in does not fully exceed. It is given by

$$k^b := \max\{x : l_i^\tau \geq \sum_{k=1}^x B_i^{k\tau} - B_i^{k^1\tau}\} \in k^1 \cup \mathcal{K}, \quad (4.59)$$

Having defined the NCWO condition (see equation 4.48), let us now briefly comment on how the NCWO condition is measured in practise and what wiggle room a resolution authority has to try to make sure that the NCWO condition is satisfied.

⁵²For our purposes, we do not need to list the liabilities in priority classes k_6 and k_7 .

⁵³We note that in this equation (we set) $B_i^{k^1\tau} = 0$.

The claim following a liquidation $B_{ji}^{km\tau_b,L}$ is valued as follows. As eluded upon in Section 4.3.3, Article 74 of the BRRD tasks an independent valuer to estimate the liquidation value $B_{ji}^{km\tau_b,L}$ (or in other words the LGD $\zeta_i^{k\tau_b,L}$) in liquidation. It is left vague, however, what method the independent valuer should use to estimate the liquidation cost in a *hypothetical bail-in*. Given that the insolvency proceedings of Lehman Brothers are 11 years after its failure still ongoing, this does not seem an easy task altogether. To be able to evaluate the NCWO condition we propose a simple and reasonable way to estimate the liquidation value of a claim. We apply a uniform liquidation cost c_i^τ across assets A_i^τ to obtain the estimation (as we saw in equation 4.56). Further, given the considerable uncertainty around the estimates of the liquidation costs of a bank, creditors could easily argue that the liquidation cost c_i^τ in a liquidation would have been very low, so that they are in fact worse off in a bail-in than in a *hypothetical* bail-in. Therefore, P1 seems to act as a magnet for lawsuits.

The value of a claim following a bail-in $B_{ji}^{km\tau_b,B}$ is chosen and valued as follows by the resolution authority. The resolution authority picks its conversion rates (Δ_{ia}^k and Δ_{ib}^k) to aim to satisfy P1 (the NCWO condition) and P2. Concretely, given a chosen recapitalisation target ρ_i^T and given its chosen conversion rates Δ_{ia}^k and Δ_{ib}^k it can compute the value of the haircutted claim $B_{ji}^{km\tau_b}$ (using equation 4.50) and the value of the acquired equity share $\epsilon_{ji}^{km\tau_b}$ (using equation 4.36) of bank i 's estimated phase- b equity $\tilde{E}_i^{\tau_b}$ (using equation 4.27) to obtain the claim value $B_{ji}^{km\tau_b,B}$ in a bail-in (see equation 4.50). The resolution authority can keep changing the conversion rates until it has found a configuration that satisfies P1 and P2. In this Section we show that the regulator has not much wiggle room to set the conversion rates as they wish. Typically, only the 'fair' rates satisfy P1 and P2, sometimes 'unfair' rates are also permissible (but not desirable), and occasionally no configuration is possible, in which case the resolution financing funds need to come to aid. The scope to set unfair rates and the circumstances when the resolution financing fund must be used are discussed in what follows.

Specification of Principle II: Preservation of Hierarchy of Claims


To fulfil Principle II conversion rates must be set such that for any priority class k greater than priority class $k - 1$

$$\Delta_i^k \geq \Delta_i^{k-1} \quad (4.60)$$


holds.

Scope to set ‘Unfair’ Conversion Rates

In a limited set of circumstances the resolution authority is able to apply differential rates beyond that which is applied by setting conversion rates ‘fairly’ (i.e. $\tilde{\Delta}_b > \tilde{\Delta}_a = 0$) without breaching principle P1 and P2. Essentially, applying ‘unfair’ rates may be possible whenever the estimated liquidation cost in a hypothetical bail-in is large and whenever the size of the exclusions is relatively small (see Figure 4.8a).



(a)



(b)

Figure 4.8: Plot 4.8a shows that the larger the (assumed) liquidation costs c_i^τ (see equation 4.56) are and the smaller the exclusions X_i^τ (see equation 4.68) are, the more scope exists to apply extra write-downs $\psi_i^{\tau a}$ (see equation 4.67) to set unfair rates (see equation 4.38 and 4.61). The amount of extra write-downs is always capped to not exceed the amount necessary to reach the recapitalisation target. In other words, unfair conversion rates are only possible whenever the estimated liquidation costs c_i^τ in a hypothetical bail-in are sufficiently high and the exclusions X_i^τ of debt from bail-in are sufficiently low. The reason is that in such case creditors are worse off in liquidation than in bail-in, so that creditors in bail-in could be made as worse off as they would have been in liquidation (without breaching the no-creditor-worse-off condition) by applying $\psi_i^{\tau a^*}$ extra write-downs and allocating the extra shares to creditors in the highest priority classes (who are part of the recapitalisation phase). Plot 4.8b shows that the smaller the estimated liquidation costs c_i^τ are and the larger the exclusions X_i^τ are, the more likely it becomes that the resolution financing fund $F_i^{\tau b}$ (see equation 4.69) has to be employed to make a difference payment to the creditors who are worse off in bail-in than in liquidation (see Section 4.5.1.2). The resolution-financing-fund contribution is always capped by the amount necessary to reach the recapitalisation target. Interestingly, by comparing Plot 4.8a and 4.8b we observe that the size of the unfair write-down $\psi_i^{\tau a^*}$ and the size of the resolution financing fund contribution $F_i^{\tau b}$ resemble a mirror image: either the liquidation costs c_i^τ are high enough and the exclusions X_i^τ low enough leaving scope (but no need) to set unfair rates, or the liquidation costs c_i^τ are low enough and exclusions X_i^τ are high enough necessitating resolution financing fund contributions $F_i^{\tau b}$ to compensate creditors who are worse off in a bail-in. It never happens that unfair rates can be applied, while the resolution financing fund must pay, or visa versa. It is always one or the other. Finally, contrary to common believe, from Plot 4.8a and 4.8b it seems that for realistic liquidation costs (e.g. $c_i^\tau \geq 10\%$) resolution authorities do not have to greatly worry about breaching the no-creditor-worse-off condition even in the face of significant exclusions X_i^τ ; for realistic liquidation costs c_i^τ regulators typically have scope to set unfair rates and are not required to pay from the resolution financing fund, since liquidation costs typically overshadow exclusions in size.

The idea is that the ‘unfair’ rates can be applied by making those creditors who would have been worse off in a liquidation equally worse off in a bail-in (these are usually the more junior creditors). The freed capital can then be allocated among creditors who receive haircuts in phase b (these are usually the more senior creditors). We argue that such practise is disproportionate and risks causing contagion due to junior creditors who receive unnecessary large losses. Let us now formalise this intuition.

The ‘unfair’ extra share, which is allocated to a creditor j who holds a contract with time to maturity m of bank i in priority class k , is given by

$$\hat{\epsilon}_{ji}^{k\tau_b} = \frac{\psi_i^{\tau_a^*} \epsilon_{ji}^{k\tau_b}}{b_i}, \quad (4.61)$$

where the unfair recapitalisation amount b_i is explained in equation 4.64 below.

We recall from equation 4.38 that a creditor j receives the unfair share $\hat{\epsilon}_{ji}^{\tau_b, k}$ on top of its fair share $\epsilon_{ji}^{\tau_b, k}$, so that its total share received equals $\epsilon_{ji}^{km\tau_b}$. The term $\psi_i^{\tau_a^*}$ denotes the total size of the extra write-downs in (what we refer to as) phase a^* , to make creditors equally worse off in bail-in as in liquidation, which we will define shortly (in equation 4.67).

The procedure to apply the unfair rates is as follows:

1. Execute the loss absorption phase as usual (see Section 4.5.1.4).
2. Apply $\psi_i^{\tau_a^*}$ amount of extra write-downs in phase a^* . This means that the conversion rate in phase a^* is zero, $\Delta_{ia^*}^k = 0, \forall k \in \mathcal{K}$.
3. Proceed to the recapitalisation phase. The recapitalisation phase is executed in a regular way except for the following adjustments:

- (a) The necessary recapitalisation amount \hat{b}_i is reduced due to the extra write-downs $\psi_i^{\tau_a^*}$. Therefore, equation 4.24 updates to

$$\hat{b}_i = \rho_i^T \Omega_i^{\tau_a} - \tilde{E}_i^{\tau_{a^*}}, \quad (4.62)$$

where

$$\tilde{E}_i^{\tau_{a^*}} = \tilde{E}_i^{\tau_a} + \psi_i^{\tau_a^*}. \quad (4.63)$$

- (b) The feasible recapitalisation amount, given in equation 4.25, updates to

$$b_i = \min\{\hat{b}_i, B_i^{\tau_{a^*}}\}, \quad (4.64)$$

where $B_i^{\tau_{a^*}}$ is given by $B_i^{\tau_a} - \psi_i^{\tau_a^*}$.

- (c) The phase- b capital, given in equation 4.27, updates to

$$\tilde{E}_i^{\tau_b} = \tilde{E}_i^{\tau_{a^*}} + b_i \quad (4.65)$$

- (d) For any haircuts applied in phase b return a ‘fair’ share $\hat{\epsilon}_{ji}^{km\tau_b}$. The extra ‘unfair’ share is compensated based on the ‘fair’ share received (see equation 4.61). Hereby, it is critical note that some creditors who would have received a fair share in phase b if no ‘unfair’ rates were applied, now instead face a pure write-down in phase a^* . As such the pool of creditors who benefit from the extra unfair share $\hat{\epsilon}_{ji}^{km\tau_b}$ is confined to only those creditors who are subject to haircuts in phase b excluding those who received haircuts in phase a^* or phase a .

Let us now return to the extra write-downs $\psi_i^{\tau_a^*}$ that are permissible without breaching P1 and P2. The size of the extra write-down $\psi_i^{\tau_a^*}$ is given by

$$\psi_i^{\tau_a^*} = \min\{\hat{\psi}_i^{\tau_a^*}, y_i b_i\} \quad (4.66)$$

where the recapitalisation amount b_i is computed using equation 4.24 and 4.25 based on time τ_a values. Further, $y_i \in [0, 1]$ is some fraction that determines how unfair resolution authorities would like to make the shares of those who receive haircuts in phase b , and $\hat{\psi}_i^{\tau_a^*}$ is given by

$$\psi_i^{\tau_a^*} = \begin{cases} 0, & \text{if } k^b > k^l; \\ B_i^{k^b+1, \tau} \max\{\zeta_i^{k^l+1, \tau_a, L} - \zeta_i^{k^b+1, \tau_a, B}, 0\}, & \text{if } k^b = k^l; \\ (1 - \zeta_i^{k^b+1, \tau_a, B}) B_i^{k^b+1, \tau} + (\sum_{k=k^b+2}^{k^l} \zeta_i^{k\tau_b, L} B_i^{k, \tau}) \mathbb{1}_{k^l \geq k^b+2} + \zeta_i^{k^l+1, \tau_a, L} B_i^{k^l+1, \tau}, & \text{if } k^b < k^l, \end{cases} \quad (4.67)$$

From equation 4.67 we learn that ‘unfair’ rates are not possible whenever creditors in a bail-in are worse off than the LGD in liquidation (i.e. $k^b > k^l$), in which case the resolution financing fund has to be applied (discussed in Section 4.5.1.2). It is also not possible when the LGD in bail-in $\zeta_i^{k^b+1, \tau_a, B}$ is higher than in liquidation $\zeta_i^{k^l+1, \tau_a, L}$ in priority class $k^l + 1 = k^b + 1$, in the case where $k^l = k^b$. Unfair rates can only be applied whenever creditors in a bail-in are better off than they would have been in a liquidation (i.e. $k^b \leq k^l$). All the terms in equation 4.67 in essence tell you how much extra write-downs should be applied to make creditors who are better off in a bail-in than in liquidation equally worse off as in liquidation.

Equation 4.66 tells that the extra amount of write-down that is possible $\hat{\psi}_i^{\tau_a^*}$ without breaching the no-creditor-worse-off condition, should never be greater than the recapitalisation amount b_i , else the bank will be recapitalised excessively much (i.e. above the recapitalisation target ρ_i^T). Even more, the extra write-down $\hat{\psi}_i^{\tau_a^*}$ should not exceed a fraction y_i of the recapitalisation amount b_i to ensure that some creditors will remain subject to haircuts in phase b and thus will be reaping the benefits of the extra unfair

share (via equation 4.61). In other words, if it would be the case that $\hat{\psi}_i^{\tau^*} = b_i$, then the recapitalisation phase would be made redundant and no recipients of the extra unfair shares would exist, leaving the equity shares of a bank without owner.

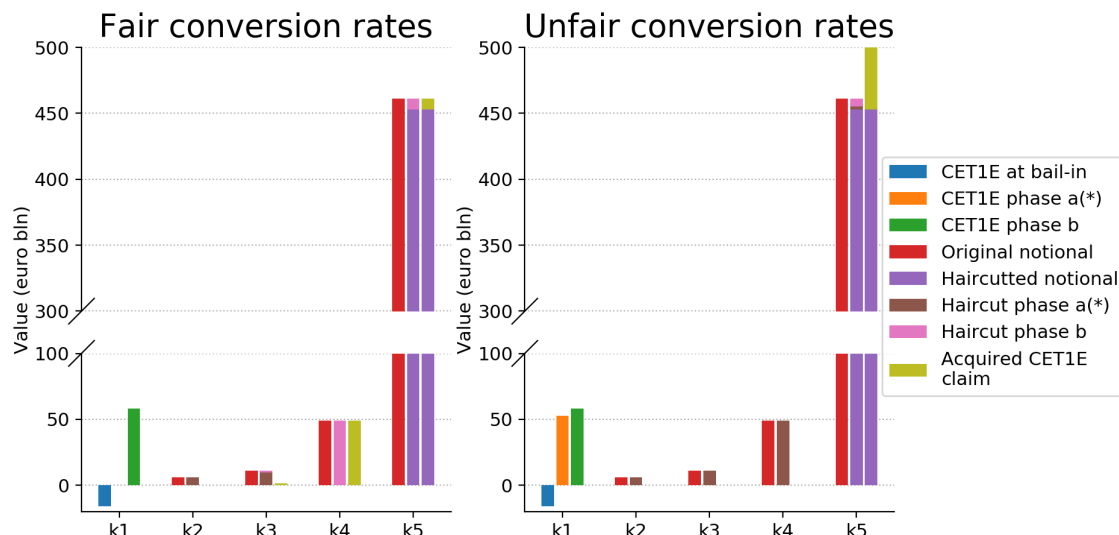


Figure 4.9: Compares the haircuts and acquired equity claims for ‘fair’- and ‘unfair’ conversion rates. For the case of fair conversion rates, we have explained in Figure 4.7 the application of haircuts in phase a and b and the debt-to-equity conversion. Here we will focus on the differences between the case of ‘fair’- and ‘unfair’ conversion rates. To evaluate the scope for setting unfair rates, resolution authorities must estimate the liquidation cost c_i^τ (see equation 4.56) in a hypothetical liquidation. Also, they must determine whether any debt will, on an ad hoc basis, be excluded from bail-in $U_i^{k\tau}$ (see equation 4.68). For the purposes of this illustration we suppose that the assumed liquidation cost is $c_i^\tau = 10\%$ and no ad hoc exclusions are made (i.e. $U_i^{k\tau} = 0$). Let us now proceed to inspect the bars associated to priority classes k_2 , k_3 and k_4 for both the fair and unfair case. We observe that for the unfair case an extra write-down $\hat{\psi}_i^{\tau^*}$ (see equation 4.66) has been applied to the bail-inable debt in priority class k_2 , k_3 and k_4 , resulting in extended haircuts in phase a (see the extended brown bars at k_2 , k_3 and k_4 in the right plot), which are not replaced by shares of CET1 equity (no equally sized pink bars). The extra write-down in phase a^* , $\hat{\psi}_i^{\tau^*}$, means that the remaining haircuts that need to be applied to fully recapitalise the bank have decreased. Therefore, you see that the haircuts in phase b (see the pink bar at k_4 in right plot) have become smaller in the ‘unfair’ case than in the ‘fair’ case. In the left plot, the total positive amount of CET1 equity of bank i , which is created by the haircuts in phase a^* and phase b , will be allocated as shares to creditors who were subject to phase b of the bail-in only (i.e. to creditors who received haircuts in phase b in k_4). This is unfair, because creditors in phase a^* receive no shares and creditors in phase b receive all the shares. In other words, creditors in phase b make a net profit, because the acquired equity claims is worth more than they haircuts they receive (observe that the light-green bar is larger than the pink bar in k_4 for the ‘unfair’ case). Compare this to the fair case. Here creditors in phase b obtain an equal amount of CET1 equity in return for any haircuts they receive in phase b (observe that the light-green bar is equal to the pink bar in k_2 , k_3 and k_4). In sum, unfair rates result in disproportionate losses to some (i.e. those who face pure write-downs in phase a^*) and disproportionate profits to others (i.e. those who remain part of phase b). Thus, unfair rates redistribute wealth unfairly. In contrast, fair rates ensure that all creditors in phase b make neither net profits nor net losses as a *direct* consequence of the bail-in. Obviously, both in the fair and the unfair case creditors who receive haircuts in the loss absorption phase a undergo pure write-downs. This is inevitable: a resolution authority cannot give shares in return as long as the bank is still insolvent.

Resolution Financing Agreements P1 stipulates that if the NCWO condition is violated, then creditors under Article 75 of the BRRD are entitled to the difference between the expected loss under resolution and liquidation if this is positive. In such case, resolution authorities have no choice but to resort to the resolution financing fund.

When the exclusions $X_i^{k\tau}$ of claims (from the hierarchy of bail-inable claims (see Section 4.5.1.1) relative to the liquidation hierarchy (see equation 4.55)) are positive and the liquidation costs c_i^τ (see equation 4.56) are small enough, only then it is likely that the resolution financing fund contribution is required. Without exclusions $X_i^{k\tau}$ the resolution financing fund never has to be tapped. The larger the exclusions $X_i^{k\tau}$ are and the smaller the estimated liquidation cost c_i^τ are in a hypothetical bail-in, the more likely it becomes that the resolution financing fund must be used (see Figure 4.8b). Given this observation, to avoid ad-hoc exclusions $U_i^{k\tau}$ (see Section 4.5.1.1), it makes complete sense that instruments must be subordinated to count towards TLAC T_i or MREL M_i (see Section 4.5.1.2).

The size of the exclusions $X_i^{k\tau}$ in priority class k at the start of bail-in τ is given by any positive difference of the claims that are by default included in the hierarchy of liquidation $L_i^{k\tau}$ and the hierarchy of bail-inable claims $B_i^{k\tau}$ plus the claims that are excluded on an ad-hoc basis $U_i^{k\tau}$ (e.g. because bail-in in these claims would give rise to contagion or be too difficult, see Section 4.5.1.1). That, the size of exclusions $X_i^{k\tau}$ is given by

$$X_i^{k\tau} := (L_i^{k\tau} - B_i^{k\tau}) + U_i^{k\tau} \quad (4.68)$$

For a given loss that must be absorbed in a bail-in \hat{l}_i (see equation 4.14), the larger the exclusions $X_i^{k\tau}$ are the more priority classes k^b will experience full write-downs and the more likely a breach of the NCWO condition becomes. The total size of the resolution fund contribution $F_i^{\tau b}$ is given by

$$F_i^{\tau b} = \begin{cases} (1 - \zeta_i^{k^l+1, \tau_a, L}) B_i^{k^l+1, \tau} + (\sum_{k=k^l+2}^{k^b} \zeta_i^{k\tau_a, B} B_i^{k\tau}) \mathbb{1}_{k^b \geq k^l+2} + \zeta_i^{k^b+1, \tau_a, B} B_i^{k^b+1, \tau}, & \text{if } k^b > k^l; \\ B_i^{k^b+1, \tau} \max\{\zeta_i^{k^b+1, \tau_a, B} - \zeta_i^{k^l+1, \tau_a, L}, 0\}, & \text{if } k^b = k^l; \\ 0, & \text{if } k^b < k^l, \end{cases} \quad (4.69)$$

The terms in equation 4.69 show how much worse off a creditor is in bail-in than in liquidation, and thus how much a creditor should be compensated. Comparing the size of the resolution fund contribution $F_i^{\tau b}$ (in equation 4.69) and the scope for extra write-downs $\psi_i^{\tau a^*}$ (see equation 4.67), we immediately observe the symmetries just discussed.

Importantly, whenever more creditors are affected by the loss absorption phase of a bail-in than by the loss absorption phase in liquidation (i.e. if $k^b > k^l$) or whenever

creditors are more severely affected (i.e. if $\zeta_i^{k^b+1,\tau_a,B} > \zeta_i^{k^l+1,\tau_a,L}$ if $k^b = k^l$), then the resolution authority *cannot* avoid a breach of the NCWO condition by setting rates ‘fairly’ or ‘unfairly’: in both cases the resolution financing fund has to be deployed. The reason that applying ‘unfair’ rates is not effective in avoiding having to tap into resolution financing fund, is that awarding some creditors more shares than their fair share makes others worse off.

Dilution of Existing Equity Holders If the bank is insolvent at the start of bail-in τ (i.e. when phase a applies, see equation 4.14), then the rule is that existing equity holders are completely wiped out (Klimek et al. (2015), Hüser et al. (2017)). We therefore set the unit holdings $\eta_i^{\tau_a}$ of existing equity holders to zero in such case. That is,

$$\eta_i^{\tau_a} = \begin{cases} \eta_i^{\tau}, & \text{if } \rho_i^{\tau} \geq 0; \\ 0, & \text{if } \rho_i^{\tau} < 0. \end{cases} \quad (4.70)$$

In that spirit, the shares of existing equity holders update to

$$\epsilon_{ji}^{k_1 m \tau_a} = \begin{cases} \epsilon_{ji}^{k_1 m \tau}, & \text{if } \rho_i^{\tau} \geq 0; \\ 0, & \text{if } \rho_i^{\tau} < 0. \end{cases} \quad (4.71)$$

If the bank is still solvent at the start of the bail-in (i.e. $\rho_i^{\tau} \geq 0$), then existing equity holders are diluted whenever the conversion rates Δ_i^k for $k \in \mathcal{K}$ are positive. In general, the dilution factor d_i , which tells by what fraction the non-wiped-out equity holders share $\epsilon_{ji}^{k_1 m \tau_a}$ of bank i ’s equity is reduced,⁵⁴ is given by

$$d_i = 1 - \frac{\eta_i^{\tau_a}}{\eta_i^{\tau_b}} \in [0, 1], \quad (4.72)$$

where $\eta_i^{\tau_a}$ is given by equation 4.70 and $\eta_i^{\tau_b}$ is given by equation 4.37. We note that if shareholders in phase a are completely wiped out ($\eta_i^{\tau_a} = 0$) then the dilution factor d_i is equal to 1 (i.e. a hundred percent). Given a dilution factor d_i , the share $\epsilon_{ji}^{k_1 m \tau_a}$ of existing equity holders updates to

$$\epsilon_{ji}^{k_1 m \tau_b} = (1 - d_i) \epsilon_{ji}^{k_1 m \tau_a}. \quad (4.73)$$

When the dilution factor \tilde{d}_i is ‘fair’, which is the case when fair conversion rates have been applied in phase a and b (see equation 4.40 and 4.41), then the claims $E_{ji}^{k_1}$ of existing equity holders at phase a and b are equal, so that non-wiped-out existing equity holders do not suffer net losses as a consequence of the bail-in. That is,

⁵⁴We recall that k_1 denotes priority class 1 associated to the existing CET1 equity holders (see Section 4.5.1.1).

$$E_{ji}^{k_1\tau_b} = E_{ji}^{k_1\tau_a}, \quad (4.74)$$

which we can rewrite as

$$\begin{aligned} \epsilon_{ji}^{k_1\tau_a} \tilde{E}_i^{\tau_a} &= \epsilon_{ji}^{k_1\tau_b} \tilde{E}_i^{\tau_b} \\ &= \epsilon_{ji}^{k_1\tau_a} (1 - d_i) \tilde{E}_i^{\tau_b}, \end{aligned} \quad (4.75)$$

using equation 4.39 and 4.73. Equation 4.75 implies that the ‘fair’ dilution factor $\tilde{d}_i \in [0, 1]$ is given by

$$\tilde{d}_i = 1 - \frac{\tilde{E}_i^{\tau_a}}{\tilde{E}_i^{\tau_b}} \quad (4.76)$$

$$= 1 - \frac{\tilde{E}_i^{\tau_a}}{\tilde{E}_i^{\tau_a} + b_i}, \quad (4.77)$$

where we used equation 4.27 to obtain equation 4.77. From equation 4.75 we observe that the ‘fair’ dilution factor \tilde{d}_i increases in the recapitalisation amount b_i .

We emphasise that setting a fair dilution factor \tilde{d}_i reduces (dilutes) the fair share $\tilde{\epsilon}_{ji}^{k_1\tau_b}$ of existing equity holders (i.e. $\tilde{\epsilon}_{ji}^{k_1\tau_b} \leq \epsilon_{ji}^{k_1\tau_a}$), but leaves their equity claim $E_{ji}^{k_1\tau_b}$ unaltered (i.e. $E_{ji}^{k_1\tau_b} = E_{ji}^{k_1\tau_a}$) producing no net losses as a *direct* consequence of the bail-in, since the bank’s CET1 equity $\tilde{E}_i^{\tau_b}$ increases (i.e. $\tilde{E}_i^{\tau_b} \geq \tilde{E}_i^{\tau_a}$) as a consequence of phase *b* of the bail-in. By default we will assume that the dilution factor d_i is fair.

Having modelled the bail-in design, we next elucidate how we model the contagious amplifications that this design might promote.

4.5.2 Modelling Multiple Contagion Mechanisms in the Financial System

The aim of this Section is to delineate how we model multiple interacting contagion mechanisms in a financial system consisting of banks and non-banks – where bail-in is the preferred method for dealing with bank failures. We will make intelligible how we model five prevailing contagion mechanisms: bail-in-induced exposure loss contagion, overlapping portfolio contagion, funding contagion, bail-inable debt revaluations and halts on roll-overs of bail-inable debt (referred to as bail-inable debt ‘runs’). We will treat these contagion mechanisms in turn.

4.5.2.1 Exposure Loss Contagion, Overlapping Portfolio Contagion & Funding Contagion

The previous Section 4.5.1 has discussed how the bail-in design may spark off bail-in-induced exposure loss contagion. Exposure loss contagion may arise if creditors of a bailed-in bank suffer exposure losses due to the bail-in and pass on the distress. We made it clear that creditors tend to suffer larger exposure losses if the failure threshold is low and the debt-to-equity conversion rate is ‘unfair’. If the failure threshold is low it is more likely that the bank is insolvent at the start of the bail-in, necessitating pure write-downs to raise it out of insolvency. If the bank is insolvent, it is never possible to apply positive debt-to-equity conversion rates to junior creditors that are subject the loss absorption phase – without making more senior creditors ‘worse off’, which would breach the NCWO condition. If the debt-to-equity conversion rate is ‘unfair’ then creditors in junior priority classes must endure large exposure loss with no compensation. Hence, ‘unfair’ conversion rates tend to fuel contagion.

We use the system-wide stress testing model developed by [Farmer et al. \(2020\)](#) to model funding contagion and overlapping portfolio contagion. We provide a synopsis of the approach here.

Banks seek to avoid default by aiming to fulfil contractual obligations and comply with regulatory constraints. If a bank receives a funding shock and has not sufficient cash to meet the withdrawal obligation, it will resort to liquidating assets to raise cash. Banks maintain a ‘pecking order’ that specifies which assets they liquidate first to meet contractual obligations. We assume that a bank liquidates the most liquid assets first to avoid unnecessary liquidation costs. In practise this means that banks exhausts withdrawing *maturing* loans (which can precipitate funding contagion), before liquidating tradable assets (which can provoke overlapping portfolio contagion). Other than in [Farmer et al. \(2020\)](#), we, in this system-wide stress test, introduce maturities to account for collapse-prone or collapse-proof nature of a bank’s bail-inable debt pile, as expressed by its maturity profile of bail-inable debt. To the best of our knowledge, we are the first to build a model of funding contagion that captures maturities, rather than assuming all short-term debt is overnight (and thus can be withdrawn every timestep).

To avoid breaching the minimum risk-weighted capital requirements $\rho^M = 4.5\%$ of Common Tier I Equity relative to risk-weighted assets, as well as to eschew punitive restrictions on the ability to make discretionary payments (such as dividend payments) when falling too many quantiles into the regulatory capital buffer, banks maintain an internal

‘buffer’ and target’. Whenever their risk-weighted capital ratio falls below this buffer they will seek to gradually return to a robust recapitalisation target. In line with [Farmer et al. \(2020\)](#), we assume banks are willing to use no more than $u = 50\%$ of their combined regulatory capital buffer ρ_i^{CB} – which consists of the capital conservation buffer, countercyclical buffer, G-SIB surcharge and systemic risk buffer. Whenever banks have absorbed losses with their regulatory buffer ρ_i^{CB} in excess of $u = 50\%$ (i.e. whenever $\rho_i < \rho_i^B = \rho^M + (1 - u)\rho_i^{CB}$), they will seek to return to a target risk-weighted capital ratio ρ_i^T that complies with the regulatory buffer standards (i.e. $\rho_i^T = \rho^M + \rho_i^{CB}$). Other than in [Farmer et al. \(2020\)](#), we, more in keeping with reality, assume they climb up to their target gradually: by no more than 0.5% per day – thereby limiting self-inflicted harm in terms of high liquidation cost and system-inflicted damage. With this approach the target strategy is given by $\rho_i^T = \min\{\rho_i + 0.5\%, \rho^M + \rho_i^{CB}\}$. The gradual (risk-weighted) deleveraging approach improves upon contagion models that unrealistically do not put a cap on the percentage of a bank’s balance sheet will be liquidated at once ([Bookstaber \(2012\)](#), [Duarte & Eisenbach \(2015\)](#), [Greenwood et al. \(2015\)](#), [Cont & Schaanning \(2017\)](#)). To increase the risk-weighted capital ratio to the target, we assume that banks maintain a ‘risk-weighted pecking order’ that determines that banks will liquidate the assets with the highest risk-weights first in order to quickly prop up the ratio. The liquidation of assets may instigate overlapping portfolio contagion or funding contagion, or both.

In line with [Farmer et al. \(2020\)](#), we model the behaviour of the leveraged non-bank by its actions to keep its leverage ratio relatively stable. Whenever its leverage falls to 90% of their initial leverage, it will seek to return to its initial leverage ratio $\lambda_i^{t_0}$. The non-leveraged non-bank absorbs losses without feeding it back onto the system.

We proceed to discuss our novel method for pricing bail-inable debt. This method is used to capture that revaluations of bail-inable debt lead to mark-to-market losses of debt holders. The incurred losses may pressurise these institutions delever, and thereby add to contagious pressures. It also can lead to further mark-to-market losses on bail-inable debt holdings.

4.5.2.2 Revaluation of Bail-Inable Debt

Risk-Neutral Valuation We price bail-inable debt B_{ji}^{km} as the discounted expected value of future payoffs, in line with the standard approach to pricing financial contracts (see e.g. [Black & Scholes \(1973\)](#), [Merton \(1974\)](#)). The time- t value V_{ji}^{kmt} of j ’s claim on bank i ’s bail-inable debt in priority class k which matures at time $T = t + m$ (where we recall that $m := T - t$ is the time to maturity) is given by

$$V_{ji}^{kmt} = \exp^{-r(T-t)} \mathbb{E}^{\mathcal{Q}}[P_{ji}^{kT}(A_i^{\tau}, A_i^T)]. \quad (4.78)$$

Equation 4.78 tells that the time- t value V_{ji}^{kmt} of bail-inable debt B_{ji}^{kmt} is given by the expected payoff P_{ji}^{kT} under the risk-neutral measure \mathcal{Q} discounted by the risk-free rate r (which we assume to be zero, in line with today's low interest rate environment). The payoff P_{ji}^{kT} at time T is a function of bank i 's asset value A_i at the time of bail-in τ and at the maturity T of the bail-inable contract B_{ji}^{km} .

Jump Process for the Evolution of the Asset Value To value a bail-inable contract B_{ji}^{km} we need a process that governs the evolution of the asset value A_i . A simple geometric Brownian motion (GBM) would not suffice, since it does not capture the fat tails observed in the banks' asset values A_i (Rachev et al. (2005)). Especially in financial crises, which we focus on, a bank may suffer sharp tail losses in its asset value A_i , for example due to counterparty defaults (giving exposure losses) or fire sales (giving marked-to-market losses) (Farmer et al. (2020)).

Therefore, in line with Pennacchi (2010), who introduced a way to price the contractual analogue of bail-inable debt, CoCos, we propose to model the asset value A_i evolution according to a jump process. Specifically, we apply Merton's jump-diffusion process (Merton (1976)) with log-normal jumps. That is, the risk-neutral jump process for the asset value A_i follows

$$\frac{dA_i^t}{A_i^t} = (r - \lambda_i \bar{j}_i) dt + \sigma_i dW_i^{\mathcal{Q}t} + j_i dq_i^t, \quad A_i^t = A_i^{t_s}. \quad (4.79)$$

We will explain the components of equation 4.79 now. The initial condition stipulates that asset value A_i at time t equals the asset value in stress test simulation at time t , which is denoted by t_s . The jump events of bank i are governed by a compound Poisson process q_i^t with jump intensity λ_i , which gives the mean number of arrivals per unit time t . The magnitude of bank i 's random jump is given by j_i (note, the magnitude should be interpreted as the fractional increase or decrease in the bank's total assets A_i), where $1 + j_i$ is log-normally distributed with mean μ_i^J and standard deviation σ_i^J , that is,

$$\ln(1 + j_i) \sim \mathcal{N}(\mu_i^J, \sigma_i^J) \quad (4.80)$$

The mean jump size \bar{j}_i is given by

$$\bar{j}_i = \exp(\mu_i^J + \frac{(\sigma_i^J)^2}{2}) - 1. \quad (4.81)$$

If the jumps are turned off (i.e. $\lambda_i = 0$), then equation 4.79 returns to a geometric Brownian motion identical to those posited in Black & Scholes (1973), Merton (1974). That is,

$$\frac{dA_i^t}{A_i^t} = rdt + \sigma_i dW_i^{\Omega t}, \quad A_i^t = A_i^{t_s}. \quad (4.82)$$

where σ_i^2 is the instantaneous variance of bank i 's asset returns conditional on the Poisson event not occurring, and $dW_i^{\Omega t}$ is a standard Gauss-Wiener process under the risk-neutral dynamics. Further, dq_i^t and $dW_i^{\Omega t}$ are assumed to be independent.

The solution to equation 4.79 is

$$A_i^{t+1} = A_i^t \exp\left((r - \lambda_i \bar{j}_i - \frac{\sigma_i^2}{2})\Delta t + \sigma_i \sqrt{\Delta t} Z_i\right) (1 + j_i)^{\phi_i}, \quad A_i^t = A_i^{t_s}, \quad (4.83)$$

where Z_i is distributed according to the standard normal distribution, and where ϕ_i equals to one with probability λ_i and equals to zero with probability $1 - \lambda_i$. Hence, a jump occurs if $\phi_i = 1$ and no jump occurs if $\phi_i = 0$. The magnitude of jump j_i is (similar to equation 4.80) log-normally distributed. That is,

$$\ln(1 + j_i) \sim \mathcal{N}(\mu_i^J, \sigma_i^J). \quad (4.84)$$

Payoff of Bail-Inable Debt Using equation 4.83, the path of bank i 's asset value A_i can be generated over the lifetime of the bail-inable contract $B_{j_i}^{k,m}$ ($s \in [t, T]$). Given the invariant-liability assumption, we can also run n Monte Carlo paths of the CET1 equity value \tilde{E}_i and risk-weighted capital ratio ρ_i (see Figure 4.10).⁵⁵ That is: $B_i^s = B_i^t$ and $L_i^s = L_i^t, \forall s \in [t, \min\{\tau_i^{m,n}, T\}]$; and $B_i^s = B_i^{\tau_b, n}$ and $L_i^s = L_i^{\tau_b, n}, \forall s \in [\tau_b^{m,n}, T]$. This assumption is in line with the invariant-liability assumption in Merton's structural credit risk model (Merton (1974)). To generate the equity \tilde{E}_i path we use the invariant-liability assumption and equation B.2 and B.3. To generate the risk-weighted capital ratio ρ_i path, we use the invariant-liability assumption, equation 4.10 and approximation $\Omega_i^{t+1} \approx \frac{A_i^{t+1}}{A_i^t} \Omega_i^t$. The generated sample paths are then given by

$$\{A_i^{t_s, n}, A_i^{t+1, n}, \dots, A_i^{T, n}\}; \quad (4.85)$$

$$\{\tilde{E}_i^{t_s, n}, \tilde{E}_i^{t+1, n}, \dots, \tilde{E}_i^{T, n}\}; \quad (4.86)$$

$$\{\rho_i^{t_s, n}, \rho_i^{t+1, n}, \dots, \rho_i^{T, n}\}. \quad (4.87)$$

⁵⁵Our invariant-liability assumption, which we use for the purposes of valuation only, says that the bail-inable debt and liabilities remain equal to their time- t value up to bail-in, and stay equal to their post-bail-in value τ_b up to maturity T .

Using these generated sample paths (equation 4.85, 4.86 and 4.87), we can compute the payoff $P_{ji}^{kT,n}$ for each Monte Carlo run n . The payoff $P_{ji}^{kT,n}$ is given by

$$P_{ji}^{kT,n} = \begin{cases} B_{ji}^{kmT,n}, & \text{if } \rho_i^{s,n} \geq \rho_i^F, \forall s \in [t, T] \\ B_{ji}^{kmT,n} + E_{ji}^{kmT,n}, & \text{if } \rho_i^{s,n} < \rho_i^F, \text{ for a } s \in [t, T] \end{cases} \quad (4.88)$$

where the payoff if no bail-in occurs (i.e. if $\rho_i^{s,n} \geq \rho_i^F, \forall s \in [t, T]$) can, using the invariant-liability assumption, be rewritten as

$$B_{ji}^{kmT,n} = B_{ji}^{kmt}, \quad (4.89)$$

and the payoff if a bail-in occurs (i.e. if $\rho_i^{s,n} < \rho_i^F$, for a $s \in [t, T]$) can be rewritten as

$$\begin{aligned} B_{ji}^{kmT,n} + E_{ji}^{kmT,n} &= (B_{ji}^{km\tau,n} - h_{ji}^{km\tau_a,n} - h_{ji}^{km\tau_b,n}) + \epsilon_{ji}^{km\tau_b,n} \hat{E}_i^{T,n} \text{56} \\ &= B_{ji}^{km\tau_b,n} + \epsilon_{ji}^{km\tau_b,n} (\tilde{E}_i^{\tau_b,n} + \sum_{s=\tau_i^{m,n}}^T (\tilde{E}_i^{s+1,n} - \tilde{E}_i^{s,n})) \\ &= B_{ji}^{km\tau_b} + \epsilon_{ji}^{km\tau_b,n} (\tilde{E}_i^{\tau_b,n} + \tilde{E}_i^{T,n} - \tilde{E}_i^{\tau,n}). \text{57} \end{aligned} \quad (4.90)$$

Let us now explain the composition of the payoff $P_{ji}^{kT,n}$ in equation 4.88, which we have just decomposed using equation 4.89 and equation 4.90. Equation 4.89 tells us that the payoff $P_{ji}^{kT,n}$ in simulation run n is equal to the notional B_{ji}^{kmt} , if no bail-in occurs over the lifetime of the contract, which is so if a simulated sample path n of the risk-weighted capital ratio $\rho_i^{s,n}$ (see equation 4.87) remains at all times during $s \in [t, T]$ above the FLTF trigger ρ_i^F .

Equation 4.90 tells us that if a bail-in occurs within the lifetime of the contract (i.e. if $\rho_i^{s,n} < \rho_i^F$, for a $s \in [t, T]$), then the payoff $P_{ji}^{kT,n}$ of the bail-inable claim is equal to the time- T value of the bail-inable claim $B_{ji}^{kmT,n}$ plus the acquired equity claim $E_{ji}^{kmT,B,n}$.

Equation 4.90 can be further decomposed as follows. The bail-inable claim $B_{ji}^{kmT,n}$ at maturity T equals the haircutted bail-inable claim $B_{ji}^{km\tau_b}$ right after bail-in (using the invariant-liability assumption). The time- τ_b bail-inable claim $B_{ji}^{km\tau_b}$ equals the time- t bail-inable claim (using the invariant-liability assumption) B_{ji}^{kmt} minus the haircuts applied in phase a and b . The value of the time- T equity claim $E_{ji}^{kmT,n}$ in run n is given by the equity share $\epsilon_{ji}^{km\tau_b,n}$ acquired in phase b times the time- T value of bank i 's equity $\hat{E}_i^{T,n}$, where $\hat{E}_i^{T,n}$ denotes the time- T value of a bank's CET1 equity that is potentially modified by bail-in. In turn, $\hat{E}_i^{T,n}$ is derived from its CET1 equity right after bail-in $\tilde{E}_i^{\tau_b,n}$ plus the sample-path difference of its CET1 equity between time T , $\tilde{E}_i^{T,n}$ (not modified by bail-in) and time τ , $\tilde{E}_i^{\tau,n}$.

We highlight the two implicit assumptions we make when valuing the CET1 equity

claim $E_{ji}^{km,B,T}$ at the maturity T of the original bail-inable contract B_{ji}^{km} . First, the CET1 equity claim $E_{ji}^{km,B,T}$ can be converted to cash at time T at no liquidation cost (i.e. no price impact). Second, the new equity holder only wants to convert its CET1 equity claim $E_{ji}^{km,T}$ to cash at time T rather than later or earlier. Alternatively, you can also value the payoff of a CET1 equity claim E_{ji}^{km} right after bail-in (τ_b), if you believe creditors would want to convert their acquired equity claim $E_{ji}^{km\tau_b}$ to cash right away. In such case, $E_{ji}^{mT,B,n}$ in equation 4.88 is replaced with $E_{ji}^{m\tau_b,n}$ (assuming a zero risk-free rate r) and equation 4.90 simplifies to

$$B_{ji}^{kmT,n} + E_{ji}^{km\tau_b,n} = B_{ji}^{km\tau_b} + \epsilon_{ji}^{km\tau_b,n} \tilde{E}_i^{\tau_b,n} \quad (4.91)$$

If the acquired claim is immediately converted to cash, then the creditor j is no longer exposed to bank i 's equity fluctuations $\tilde{E}_i^{s,n}$ for $s \in [\tau_b, T]$, which could alter j 's payoff P_{ji}^{kmT} . To avoid any exposure to bank i 's equity fluctuations, creditors may prefer to convert their CET1 equity claim straight away, especially when they believe the bank's CET1 equity value \tilde{E}_i could plummet after bail-in ($t > \tau_b$).

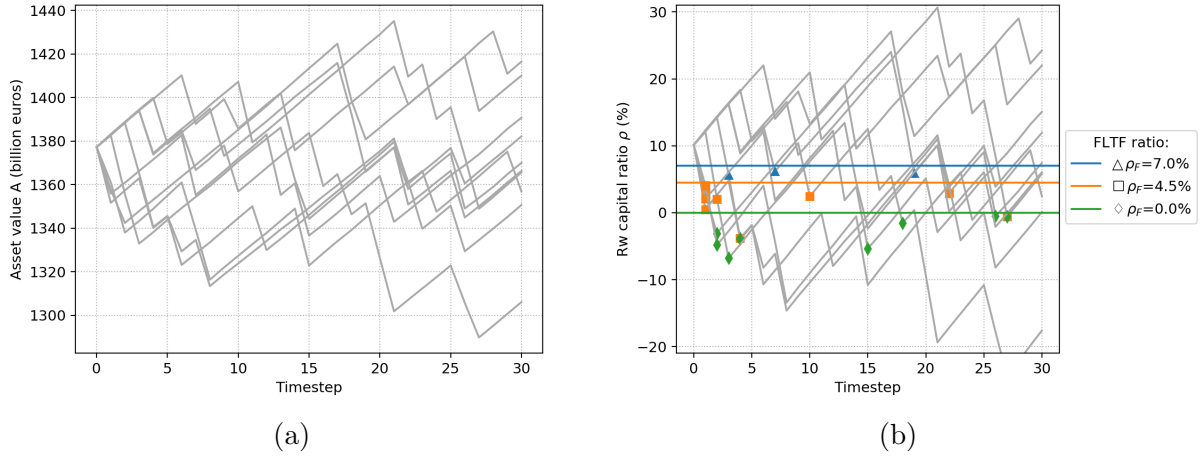


Figure 4.10: Figure 4.10a illustrates the evolution of bank $i = FR12$'s asset value A_i according to the jump process described in equation 4.83 over the course of 30 time steps for $N=10$ Monte Carlo runs. The jump-process parameters for this illustration are: $\sigma_i = 8\%$, $\mu_i^J = -2\%$, $\sigma_i^J = 2$ and $\lambda_i = 50$. Plot 4.10b shows the corresponding evolution of the risk-weighted capital ratio ρ_i . The value of bail-able debt is approximated as the average discounted payoff of N scenarios of the evolution of the risk-weighted (rw) assets ρ_i . The payoff in run n is equal to the original payoff (the notional) if the rw capital ratio did not fall below the failing-likely-to-fail trigger (FLTF) ρ^F before the maturity T (where $T = 30$ here) of the contract. The triangles, squares and spades represent the first time the risk-weighted capital ratio ρ fell in a run n below the FLTF trigger ρ^F , for the case where ρ^F is set to 7%, 4.5% and 0%, respectively. That is, these shapes represent the risk-weighted capital ratio at the start τ of the 'hypothetical bail-in' in run n , $\rho_i^{\tau,n}$. We note that when the FLTF trigger is set to $\rho^F = 0\%$, it can happen that the risk-weighted capital ratio at the start of bail-in $\rho^{\tau,n}$ is negative. This can happen as a consequence of sudden asset drops (due to exposure losses, for instance). Whereas if the FLTF trigger is set significantly above 0%, then it is less likely that a drop in the asset value causes the risk-weighted capital ratio to be negative at the start of bail-in (e.g. the spades are less likely to be below zero). Hence, the level of the FLTF trigger matters: only if the risk-weighted capital ratio is negative at the start of bail-in do some creditors face losses in phase a of the 'fair' (hypothetical) bail-in, which are typically not compensated with equity shares. Thus when the FTFT is set lower it is more likely that the payoff is reduced. This in turn affects the pricing of bail-inable debt.

Hypothetical Bail-In in a Stress Test to Price Bail-In Debt The value V_{ji}^{kmt} (see equation 4.78) of a bail-inable debt claim B_{ji}^{km} is approximated by the average payoff $P_{ji}^{kT,n}$ (see equation 4.88) over N Monte Carlo runs. That is,

$$V_{ji}^{kmt} \approx \exp^{-r(T-t)} \frac{1}{N} \sum_{n=1}^N P_{ji}^{kT,n}(A_i^{\tau,n}, A_i^{T,n}). \quad (4.92)$$

Figure 4.11 and 4.5.2.2 show the sensitivity of the valuation V_{ji}^{kmt} to the jump process parameters and to the bail-in parameters, respectively. Figure 4.5.2.2 shows the difference in valuation V_{ji}^{kmt} for different priority classes k , given a set of bail-in parameters and jump process parameters.

In the stress test, we will value V_{ji}^{kmt} each bail-inable debt contract B_{ji}^{km} of every bank $i \in \mathbb{B}$ at every time step t according to equation 4.92. We obtain the relevant

inputs to price each bail-inable debt contract B_{ji}^{km} at each time t as follows. The initial conditions for the pricing of bail-inable debt B_{ji}^{km} update at each time step t using the time- t balance sheet information in the stress test. Specifically, at each time step the inputs of the jump process and payoff function (i.e. A_i^t , \tilde{E}_i^t , ρ_i^t , L_i^t and B_i^t) update. When a bail-in event occurs in run n a ‘hypothetical bail-in’ is executed at time $\tau_i^{m,n}$ using the following information: the asset value at bail-in $A_i^{\tau,n}$, the CET1 equity value at bail-in $\tilde{E}_i^{\tau,n}$, the risk-weighted capital ratio at bail-in $\rho_i^{\tau,n}$, the time- t composition of bail-inable debt B_i^t (see equation 4.2), the time- t composition of liabilities L_i^t , the conversion rates in the loss absorption phase Δ_{ia}^k , the conversion rate in the recapitalisation phase Δ_{ib}^k and the recapitalisation target ρ_i^T (see Section 4.5.1.4). The hypothetical bail-in gives us the value of the haircuts in the loss absorption phase $h_{ji}^{km\tau_a,n}$ and the recapitalisation phase $h_{ji}^{km\tau_b,n}$ in run n . Together, this allows us to compute the value of bail-inable debt V_{ji}^{kmt} ($\forall t, k \in \mathcal{K}, i \in \mathbb{B}$).




Figure 4.11: Shows the sensitivity of the relative price fall of the bail-inable debt value V_{ji}^{kmt} (see equation 4.92) (relative to time t_0) to the four jump process (see equation 4.83) parameters: (4.11a) the jump intensity λ_i ; (4.11b) the jump mean j_i^J ; (4.11c) the jump volatility σ_i^J ; and (4.11d) the volatility σ . Several observations are noteworthy. First, the price of bail-inable debt V_{ji}^{kmt} falls relatively more as the FLTF trigger is set lower. The lower the FLTF trigger ρ_i^F is set the higher the chance that bank i is insolvent at the start of bail-in (i.e. ρ_i^F). When this is the case, creditors face pure write-downs in the loss absorption phase when conversion rates are set fairly. The consequence is that the payoff is lower and hence that the bail-inable debt value falls relatively more. Furthermore, we observe that when the jump volatility σ_i^J is positive a higher jump intensity λ_i means that the bail-inable debt value falls more, since the asset value A_i^n diverges more across N Monte Carlo runs and hence more asset paths experience a bail-in. A higher jump volatility also makes a bail-in more likely as paths diverge more, leading to more paths with a lower payoff, giving a lower bail-inable debt value. The following default set-up is used. The contour plots are shown for bail-inable debt of bank $i = FR12$ in priority class k_2 . The default jump process parameters are: $\lambda_i = 50$, $\mu_i^J = -2\%$, $\sigma_i^J = 2\%$, $\sigma_i = 5\%$. The default bail-in parameters are: initial shock $x = 1$ (equal to 2018 EBA stress test scenario), time to maturity $m = 7$ days, $\rho_i^T = \rho_i^{data}$, and the conversion rates in phase a and b are fair (i.e. $\Delta_{ia}^{k_2} = \tilde{\Delta}_a$ and $\Delta_{ib}^{k_2} = \tilde{\Delta}_b$, see equation 4.40 and 4.41). $N = 100$ Monte Carlo runs are used to compute the relative price fall at each point.

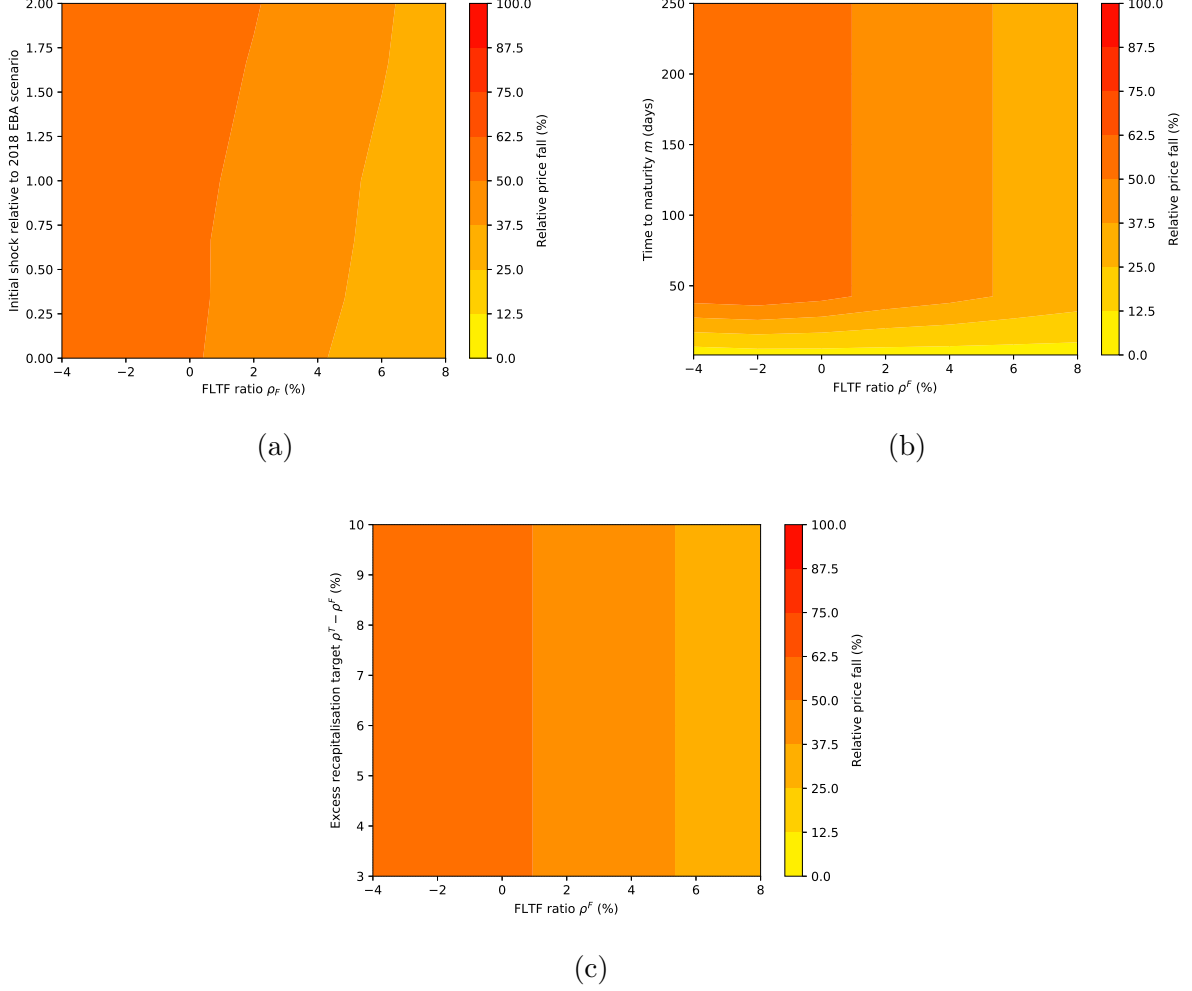


Figure 4.12: Shows the sensitivity of the relative price fall of the bail-inable debt value V_{ji}^{kmt} (see equation 4.92) (relative to time t_0) to the shock and bail-in parameters: (4.12a) the initial shock size x ; (4.12b) the time to maturity m ; and (4.12c) the excess recapitalisation target above the FLTF trigger $\rho_i^T - \rho_i^F$. Several observations are noteworthy. First, the price of bail-inable debt V_{ji}^{kmt} falls relatively more as the FLTF trigger is set lower. The reason is that the lower the FLTF trigger ρ_i^F is set the higher the chance is that bank i is insolvent at the start of bail-in (i.e. ρ_i^F). When this is the case, creditors face pure write-downs in the loss absorption phase when conversion rates are set fairly. The consequence is that the payoff is lower and hence that the bail-inable debt value falls relatively more. Furthermore, we observe that more severe initial shocks make the bail-inable debt fall more, because large initial shocks push the risk-weighted capital ratio $\rho_i^{t_1}$ down making a bail-in more likely. We also observe that bail-inable contracts with a longer time to maturity m are worth less, because it is more likely that a bail-in takes place in the elongated time period before maturity T . The default settings are as follows. The contour plots are shown for bail-inable debt of bank $i = FR12$ in priority class k_2 . The default jump process parameters (see equation 4.83) are: $\lambda_i = 50$, $\mu_i^J = -2\%$, $\sigma_i^J = 2\%$, $\sigma_i = 5\%$. The default bail-in parameters are: initial shock $x = 1$ (equal to 2018 EBA stress test scenario), time to maturity maturity $m = 7$ days, $\rho_i^T = \rho_i^{data}$, and the conversion rates in phase a and b are fair (i.e. $\Delta_{ia}^{k_2} = \tilde{\Delta}_a$ and $\Delta_{ib}^{k_2} = \tilde{\Delta}_b$, see equation 4.40 and 4.41). $N = 100$ Monte Carlo runs are used to compute the relative price fall at each point.




Figure 4.13: Shows the sensitivity of the relative price fall of the bail-inable debt V_{ji}^{kmt} (see equation 4.92) (relative to time t_0) to the jump volatility σ_i^Y for several bail-inable debt priority classes k . Plot 4.13a, 4.13b, 4.13c, and 4.13d show the the relative price fall of bail-inable debt in priority classes k_2 , k_3 , k_4 and k_5 respectively. We observe that debt in higher priority classes revalues less easily. It is logical that the lowest priority classes are the most subject to revaluations. When conversion rates are fair, these lower priority classes are the first to face pure write-downs to absorb losses, leading to payoff reductions and hence bail-inable debt reductions. Higher priority classes are less easily subject to loss absorptions and hence revalue less. Although higher priority classes may be subject to haircuts in the recapitalisation phase, this does not lead to payoff declines when the conversion rates are fair, since any haircut in phase b is replaced with an equal claim of bank i 's CET1 equity. The finding that the relative price fall becomes more pronounced as the priority class becomes lower does not only hold as a function of the jump volatility σ_i^Y and the FLTF ratio ρ_i^F , but it holds generally (i.e. across jump- and bail-in parameters). The default settings are as follows. The contour plots are shown for bail-inable debt of bank $i = FR12$ in priority class k_2 . The default jump process parameters (see equation 4.83) are: $\lambda_i = 50$, $\mu_i^J = -2\%$, $\sigma_i^J = 2\%$, $\sigma_i = 5\%$. The default bail-in parameters are: initial shock $x = 1$ (equal to 2018 EBA stress test scenario), time to maturity maturity $m = 7$ days, $\rho_i^T = \rho_i^{data}$, and the conversion rates in phase a and b are fair (i.e. $\Delta_{ia}^{k_2} = \tilde{\Delta}_a$ and $\Delta_{ib}^{k_2} = \tilde{\Delta}_b$, see equation 4.40 and 4.41). $N = 100$ Monte Carlo runs are used to compute the relative price fall at each point.

4.5.2.3 ‘Runs’ on Bail-Inable Debt

We consider the possibility that bail-inable creditors j run in anticipation of a bail-in. This is problematic for two reasons. First, runs may cause the bank’s bail-inable debt B_i to collapse, which makes it less likely that a resolution authority can successfully complete the loss absorption and liquidation phase of a bail-in (see Section 4.5.1.4 and 4.5.1.4). Second, runs may trigger a bank to liquidate assets (e.g. by fire selling securities) to repay the withdrawn bail-inable debt, which could give contagion. Given that the conditions under which bail-inable creditors run are unknown to resolution authorities, we posit three plausible run scenarios.

Run Scenarios

Uncertainty A creditor j holding bail-inable debt B_{ji}^{km} attempts to run if bank i ’s risk-weighted capital ratio ρ_i gets too close to its FLTF ratio ρ_i^F , that is if

$$\rho_i < \rho_{ji}^R := \rho_i^F + r_{ji}, \quad (4.93)$$

where r_{ji} presents some margin (e.g. 2%) on top of the FLTF trigger ρ_i^F below which a creditor attempts to run.

Creditors who attempt to run based on condition 4.93 do so irrespective of their bail-inable debt payoff P_{ji}^{mk} and hence regardless of potential losses. A reason why a creditor may run based on such a crude run condition is that large uncertainty may exist over the losses it could suffer in a bail-in, as is the case in the current regulatory regime. Resolution authorities have not provided a precise specification of the bail-in parameters, which determine the payoff losses. In particular, merely indicative guidelines exist for the FLTF ratio ρ_i^F (see Section 4.5.1.3), the recapitalisation target ρ_i^T (see Section 4.5.1.4) and the conversion rates Δ_i^k (see Section 4.5.1.4), leaving a wide range of possible values these bail-in parameters could take. Such uncertainty makes it impossible for creditors to accurately price bail-inable debt and estimate their losses if a bail-in occurs. Another factor adding to the uncertainty is the time regulators expect it takes to complete bail-in (see Section 4.3.3), which would leave creditors paralysed in the process.

Expected Losses A creditor may also attempt to run whenever its expected loss Λ_{ji}^{kmt} exceeds a certain threshold ψ_{ji} (e.g. 5%), that is if

$$\Lambda_{ji}^{kmt} := 1 - \frac{V_{ji}^{kmt}}{B_{ji}^{kmt}} > \psi_{ji}, \quad (4.94)$$

where the expected loss Λ_{ji}^{kmt} is given by the relative difference in the valuation V_{ji}^{kmt} at time t (see equation 4.92) compared to the notional of bail-inable debt B_{ji}^{kmt} time t . Other than in run condition one (see Section 4.5.2.3), under this run condition creditors would only attempt to run when the bail-in parameters (i.e. ρ_i^F , ρ_i^T , Δ_{ia}^k and Δ_{ib}^k) are set such that creditors in priority class k are expected to suffer larger losses and when such a bail-in becomes likely (i.e. if ρ_i gets close to ρ_i^F). Further, creditors in different priority classes $k \in \mathcal{K}$ would attempt to run at different times, since the expected loss Λ_{ji}^{kmt} in each priority class k is different. When fair conversion rates ($\tilde{\Delta}_a$, $\tilde{\Delta}_b$, see Section 4.5.1.4) are applied, more junior creditors would typically try to run sooner than more senior creditors, as the lowest priority classes k are the first in line to absorb losses via pure write-downs (see Section 4.5.1.4).

Value-at-Risk Losses A creditor may attempt to run whenever its value-at-risk (VaR) losses $\tilde{\Lambda}_{ji}^{kmt}$ exceed a certain threshold $\tilde{\psi}_{ji}$ (e.g. 5%), that is if

$$\tilde{\Lambda}_{ji}^{kmt} := 1 - \frac{\tilde{V}_{ji}^{kmt}}{B_{ji}^{kmt}} > \tilde{\psi}_{ji}, \quad (4.95)$$

where the worst-case value \tilde{V}_{ji}^{kmt} at time t is approximated by

$$\tilde{V}_{ji}^{kmt} \approx \frac{1}{\alpha N} \exp^{-r(T-t)} \sum_{n=1}^N P_{ji}^{kT,n}(A_i^{\tau,n}, A_i^{T,n}) \mathbb{1}_{\{P_{ji}^{kT,n} \in \mathcal{P}^{\alpha,N}\}}, \quad (4.96)$$

where $\mathcal{P}^{\alpha,N}$ is the set of the $\alpha\%$ lowest payoffs in N Monte Carlo runs. The VaR run condition is similar to run condition two (see Section 4.5.2.3), except that creditors who apply this condition are typically more risk-averse or more prudent: they run when their VaR loss $\tilde{\Lambda}_{ji}^{kmt}$ rather than their expected loss Λ_{ji}^{kmt} exceeds a certain threshold.

Payment Suspensions and Public Funding Backstop Mechanism While runs in anticipation of a bail-in are both a credible threat to financial stability and to a successful recapitalisation in a bail-in, runs during a bail-in are of less concern, in part due to proper regulatory frameworks that are in place to address liquidity issues. First, the resolution authority has the power to suspend certain payment obligations for one business day.⁵⁸⁵⁹ However, if the resolution authority has not affected bail-in within the suspension day, the bail-inable debt of the bank may still deteriorate the next day, since obligations that would have been due during the suspension period will fall due immediately upon

⁵⁸See: Article 69 of the BRRD.

⁵⁹The suspension power does not apply to certain instruments, such as eligible deposits and payments to operating systems, presumably to ensure depositors retain access to cash and financial markets continue to operate.

the expiry of the suspension period.⁶⁰ Second, even if runs take place during a bail-in, a ‘public funding backstop mechanism’ is available to G-SIBs (FSB (2016a)), which does not prevent a bail-inable debt \hat{B}_i collapse but does ensure that the bank can pay back withdrawn bail-inable debt without having to resort to liquidating assets, which potentially gives contagion. The public funding backstop mechanism can be used if a (recapitalised) firm cannot maintain private sector access to refinance its liabilities as they fall due. The term of the funding is typically no longer than needed to maintain continuity of critical functions to achieve resolution, but sufficiently long to allow the G-SIB in resolution to regain access to private sources of funding (FSB (2016a)). It is not clear whether D-SIBs and non-SIBs would also enjoy access to the public funding backstop mechanism. Hence, if they are also bailed in, a run on them during a bail-in could prove harmful.

Another part of the reason why runs during a bail-in are less of a concern is that in our model a bail-in is concluded in one time step (see section 4.3.3, as we believe it ideally should be) making runs less likely. We also consider runs following a bail-in. However, these only occur in so far as the run conditions (specified in Section 4.5.2.3) are breached again and do not occur due to fears about asset quality or concerns about long-term profitability. These concerns should be alleviated by a successful restructuring of the bank to restore long-term viability (see Section 4.3.3).

When is a Bail-Inable Creditor Able to Run? Even if a bail-inable creditor wants to run, the bail-inable creditor may not be able to run. A creditor can only run on the maturity date T of the contract. Therefore, long-term bail-inable debt $B_{ji}^{m_2k}$ is less prone to runs than short-term bail-inable debt $B_{ji}^{m_1k}$ (where the time to maturity m_2 is such that $m_2 > m_1$). In light of this, bail-inable debt B_{ji}^{mk} only counts towards a bank’s loss absorbing requirements (see Section 4.5.1.2 on TLAC & MREL, and see specifically Section B.5.1) if the time to maturity m is greater than or equal to one year. Hence, creditors that hold bail-inable debt that counts towards the loss absorbing requirements are only able to run (i.e. not roll over their maturing debt) in anticipation of a bail-in, if they anticipate bail-in at least a year in advance, making such runs less likely and less likely harmful.

However, many liabilities that do not count towards the loss absorbing requirements do count towards bail-inable debt \hat{B}_i (see Section 4.5.1.1). Only liabilities that have a time to maturity smaller than seven days are excluded from bail-inable debt \hat{B}_i (see Section 4.5.1.1). While this inconsistency may serve a purpose (e.g. banks have extra debt that is in principle bail-inable, which is useful in case the debt that counts towards the

⁶⁰See: Article 69(2) of the BRRD.

loss absorbing requirements turns out to be insufficient), this inconsistency also provides an opportunity for creditors of bail-inable debt that does not count towards the loss absorbing requirements to run in anticipation of a bail-in, since they are subject to potential bail-in losses too. These creditors are also more likely to be successful in running since their contracts mature between seven days and one year.

To conclude, let us formalise when a bail-inable creditor B_{ji}^{km} runs in our model. A bail-inable creditor runs if the following two conditions hold. First, a bail-inable creditor attempts to run if the applicable run condition (see Section 4.5.2.3) is satisfied. Second, running becomes mandatory if the maturity T of the bail-inable contract B_{ji}^{mk} coincides with the current time t in the stress test, which is the case if the time to maturity m is equal to zero.

4.6 Link to the Literature

Our contribution adds to the nascent network literature on the systemic effects of bail-in. Klimek et al. (2015) employ an agent-based network model to evaluate the economic and financial ramifications of bail-in. They compare its performance against other resolution mechanisms. Hüser et al. (2017) evaluate the systemic implications of bail-in in the EU, drawing on a calibrated multi-layered network model to bank debt and equity cross-holdings. Bernard et al. (2017) investigate the incentives for banks to contribute to a voluntary bail-in arise from their exposure to credit and price-mediated contagion. These papers neither investigate the systemic impact of the bail-in design, nor include the prevailing contagion mechanisms and non-banks in this analysis. Instead, they take the bail-in design as is and merely explore the repercussions of exposure loss contagion (Klimek et al. (2015), Hüser et al. (2017) and Bernard et al. (2017)) and overlapping portfolio contagion (Bernard et al. (2017)).⁶¹ By ignoring a set of prevailing interacting contagion mechanisms, they risk underestimating the systemic footprint of the bail-in design. Though bail-in has been designed with systemic considerations in mind,⁶² it is not enough to assert its suitability on a system-wide scale. As Aymanns et al. (2016) have shown for the case of the Basel II leverage requirements, well-intended microprudential regulation may undermine financial resilience when systemic feedbacks are taken into account. This makes the investigation of the stability implications of the bail-in design in a networked financial system a necessary gap to fill.

⁶¹To the best of our knowledge, other prior papers on bail-in discuss its risks and some of its design issues (see e.g. Eichengreen & Ruehl (2001), Rutledge et al. (2012), Persaud (2014), Conlon & Cotter (2014), Sommer (2014), De Spiegeleer et al. (2014), Avgouleas & Goodhart (2016), Schäfer et al. (2016), Zenios (2016)), but typically do not evaluate the implications of the bail-in design in a *networked* system.

⁶²See: Directive 2014/59/EU of the European Parliament and of the Council.

Our system-wide stress testing methodology extends that of [Farmer et al. \(2020\)](#). This work in turn unites the modelling of heterogeneous institutions, contracts, markets, constraints and behaviours – which can produce amplifications among multiple contagion mechanisms. Many of these have, individually or in part, though not fully jointly, been modelled in previous papers (see e.g. [Amini et al. \(2013\)](#), [Caccioli et al. \(2013, 2014, 2015, 2014\)](#), [Kok & Montagna \(2013\)](#) for models overlapping portfolio contagion, exposure loss contagion and funding contagion). Our novel contributions – in particular to [Farmer et al. \(2020\)](#), but also to the wider financial contagion literature – are to model the bail-in design, revaluations of bail-inable debt (we take a similar approach to valuing bail-inable debt as [Pennacchi \(2010\)](#), [Chen et al. \(2013\)](#) take to price contingent convertibles), multiple maturities (which are necessary to understand the stability of bail-inable debt), and gradual risk-weighted deleveraging (rather than unbounded deleveraging as [Bookstaber \(2012\)](#), [Duarte & Eisenbach \(2015\)](#), [Greenwood et al. \(2015\)](#), [Cont & Schaanning \(2017\)](#) allow for).

In terms of enhancing understanding of the bail-in design, we are the first to have worked out the implications of the two principles for setting debt-to-equity ratios: the NCWO principle and the preservation-of-hierarchy-of-claims principle. We find that regulators always have the option to set fair conversion rates whenever they do not have to resort to the resolution financing fund; and they never have the obligation to set unfair conversion rates –which are unnecessarily detrimental to financial stability – in such cases. We work out what scope regulators have to set unfair conversion rates, while remaining compliant with these two principles. We show that debt-to-equity conversion rates may be set unfairly if the assumed liquidation costs in the hypothetical disorderly liquidation of the bank are large enough and the debt exclusions from the application of the bail-in tool are small enough. Conversely, we compute the size of the resolution financing fund contribution, which is positive whenever these assumed liquidation costs are sufficiently small and the debt exclusions are sufficiently large. We are also the first, to the best of our knowledge, to precisely stipulate the TLAC and MREL requirements for banks in formula form. We also have more precisely specified in formulas the bail-in design than we have seen in any previous paper – which also contributes to the understanding of the bail-in design.

4.7 Results

Our aim is to provide evidence addressing the main question of this paper: what are the systemic implications of the bail-in design? Our three-pronged approach to answer this

enquiry takes the following logic.

First, we need to examine the joint stability impact of the parameters set by regulators when designing each individual bank bail-in: the failure threshold, the recapitalisation target and the debt-to-equity conversion rate. Since these three settings form the principle bail-in design for a specific bank, we refer to this trio as the ‘primary’ bail-in parameters. We ask: are there ‘good’ and ‘bad’ sets of primary parameters from the systemic-risk perspective? And, what is the fragility difference between the ‘good’ and ‘bad’ parameters choice, if any?

This enquiry imparts the overall impact of the primary parameters, but enmeshes the stability impact of individual primary bail-in parameters. Therefore, we proceed to disentangle the sway that each primary parameter holds over stability.

Second, to more comprehensively answer our basic research question, we ought to study supplementary stability effect that the structural (rather than bank-specific) bail-in design disseminates. In particular we focus on the systemic impact of debt exclusions, loss absorbing requirements and bail-in ‘uncertainty’. Since these parameters are structural in their nature (rather than upfront in the individual bail-in design), we refer to these as the ‘secondary’ bail-in parameters. We then pose the following questions: do secondary parameters matter on a system-wide scale? If so, to what extent could the choice of secondary parameters further widen the stability wedge between ‘good’ and ‘bad’ primary parameters?

This analysis conveys the umbrella impact of the secondary bail-in design. But falls short of unravelling the system-wide impact of the individual secondary parameters. Hence, the next plots focus on untangling the power that each of the three secondary parameters holds to modify stability.

The systemic footprint of the primary and secondary bail-in design is inadequately examined when the germane mechanisms that endogenously amplify shocks are not taken into account. Therefore, our third, and final, analysis focuses on appraising the role that each contagion mechanism plays in enlarging the systemic footprint of the bail-in design. Such system-wide analysis would not be complete without the inclusion of non-bank holdings of bail-inable debt.

To assert the robustness of our three main results and their decomposition into parts, we conduct extensive sensitivity analyses throughout the result section. To be clear, while our extensive experiments and sensitivity analyses substantiate the qualitative validity

of our results. They do not underscore the quantitative validity of our findings. For the latter to gain credence, we require a calibration of our model to more detailed data on the bail-inable debt holdings and to asset liquidities. All network models in the financial stability space, including ours, also likely benefit from better informed behaviour of financial institutions.

The first way, in which we facilitate ease of composition and cross-comparison across results, is to employ the same set of default parameters across all plots, unless otherwise stated. The default parameters are summarised in Table 4.1.

It is worthwhile to highlight our purpose of defining a set of ‘good’ and ‘bad’ primary and secondary bail-in parameters. First, it allows us to do sensitivity analyses by exploring the impact of deviations from the ‘good’ or ‘bad’ baseline. For the sake of consistency, *in all our plots we look at deviations from the ‘good’ primary parameters and status-quo secondary parameters going forward – unless otherwise indicated.* But in Appendix B.4, which exhibits our extensive sensitivity analyses, we also show that our results stand when deviations from the ‘bad’ baseline case are taken. Second, it allows us to show the stability impact of either of the extreme, but still plausible, parameters that regulators could pick. We discover that ‘good’ and ‘bad’ parameters lie at the extreme ends of the stability spectrum, while staying on the edges of the regulatory range that they could plausibly attain. Outside these ranges, regulators are not likely to set the bail-in parameters.

To be sure, we define ‘good’ and ‘bad’ parameters to sharply accentuate how stability hinges on the bail-in design. We neither intend to prescribe the regulatory application of ‘good’ bail-in parameters, nor intend to invariably dissuade wielding ‘bad’ parameters. We solely provide evidence that suggests that stability tends to ameliorate when parameters advance in the ‘good’ direction and deteriorate when parameters move in the ‘bad’ direction.

Here we summarise the settings of the ‘good’ and ‘bad’ ‘primary’ and ‘secondary’ **bail-in design**:

The triple of **good primary parameters** is given by:

1. Failure threshold at the minimum capital requirements (i.e. $\rho^F = 4.5\%$).
2. Recapitalisation target equal to minimum capital requirements plus two times combined regulatory buffers (i.e. $\rho^{RT} = \rho^M + 2\rho^{CB}$).

Table 4.1: Default settings for the figures in the result section.

Parameter Category	Default Settings	Detailed Default Settings	Brief Description & Motivation
Institutions	Banks turned on	Our research question focusses on stability in the banking sector.	
	Non-banks turned on	<ul style="list-style-type: none"> Leveraged non-banks vs. non-leveraged non-banks: $\chi = 50\%$. Initial leverage of leveraged non-banks ($\frac{E}{A}$): $\lambda_i^0 = 30\%$. 	<p>Turned on to take contagious feedback loops between the banks and non-banks into account.</p> <p>We do not know what percentage of bail-inable debt held by the non-banks is respectively held by leveraged vs. non-leveraged non-banks. We roughly know leverage of non-banks.</p> <p>(Sensitivity analysis in result section.)</p>
Contracts & Contagion Mechanisms	Bail-in induced exposure loss contagion, Overlapping portfolio contagion & Funding Contagion turned on.	We include ('turn on') all relevant contagion channels, because modelling a subset of contagion channels may lead to an underestimation of systemic risk (see e.g. Kok & Montagna (2013) , Caccioli et al. (2013) , Farmer et al. (2020)).	
	Revaluation of bail-inable debt turned on.	<p>Jump process parameters:</p> <ul style="list-style-type: none"> Volatility $\sigma = 8\%$; Jump volatility $\sigma_J = 2\%$; Jump mean $\sigma_\mu = -2\%$; Jump intensity $\lambda = 50$. 	<p>The jump process parameters are chosen in line with Chen et al. (2013), Pennacchi (2010), who model contingent convertibles (CoCos), the contractual analogue of a bail-in.</p> <p>(Sensitivity analysis in main body.)</p>
	Loss-concern-induced halts of rolling over bail-inable debt ('bail-in runs')	<p>Run scenario:</p> <ul style="list-style-type: none"> Runs based on expected losses (threshold=10%) turned on; Runs based on VaR losses (threshold=20%) turned off; Runs based on 'uncertainty' turned off. 	<p>In the stress test exercise we pose three 'what-if' scenarios for why creditors may decide to stop rolling-over bail-inable debt.</p> <p>(Sensitivity analysis in result section & Appendix.)</p>
Constraints	Risk-weighted (rw) capital ratio	Regulatory stress tests focus on assessing whether the banks' rw capital ratios remain strong enough to survive an severely adverse scenario. We could also easily include the leverage ratio and liquidity coverage ratio, building forth upon Farmer et al. (2020) .	
Market	Asset price fall is $x = 5\%$ if 5% of the market capitalisation has been sold.	This is in line with a standard assumption in the literature, see e.g. Schnabel & Shin (2004) , Cifuentes et al. (2005) Gai & Kapadia (2010) , Caccioli et al. (2014) and Farmer et al. (2020) .	
Behaviour	Seek to avoid default:	<ul style="list-style-type: none"> Fulfilling CO takes priority over complying with RC. Pecking orders: liquidate most liquid assets first (for CO), liquidate assets with the highest risk-weight first (for the rw capital ratio). 	<p>In line with Farmer et al. (2020) we assume banks are willing to use $u = 50\%$ of their combined regulatory buffer ρ_i^{CB}. Once they fall below this they will gradually seek to return to a stable rw capital ratio target ρ_i^T.</p> <p>(Sensitivity analysis in Farmer et al. (2020) on which this paper builds forth.)</p>
	Meet contractual obligations (CO) and regulatory constraints (RC).	<p>Internal rw capital buffer & target:</p> <ul style="list-style-type: none"> $\rho_i^B = \rho^M + (1 - u)\rho_i^{CB}$, where $u = 50\%$; $\rho_i^T = \min\{\rho_i + 0.5\%, \rho_i^B\}$, where $\rho_i^B = \rho_i + \frac{1}{2}\rho_i^{CB}$. 	
Failure Method & Design	Bail-in (disorderly liquidation).	<p>Good 'primary' parameters:</p> <ul style="list-style-type: none"> Failure threshold $\rho^F = 4.5\%$; Recapitalisation target $\rho_i^{RT} = \rho^M + 2\rho_i^{CB}$; Fair conversion rates. 	<p>Elected to show the financial stability impact of a well-designed 'primary' bail-in.</p> <p>(Sensitivity analysis in results section & Appendix.)</p>
	Bail-in design: Good 'primary' parameters & <i>Status quo</i> secondary parameters.	<p><i>Status Quo</i> 'secondary' parameters:</p> <ul style="list-style-type: none"> Debt excluded from bail-in with a TTM less than 7 days; Current loss absorbing requirements; 'Uncertain' bail-in design. 	<p>We take the 'secondary' bail-in design in accordance with the regulatory <i>status quo</i>.</p> <p>(Sensitivity analysis in results section & Appendix.)</p>

3. Fair conversion rates.

The triple of **bad primary parameters** is comprised of:

1. Failure threshold at insolvency (i.e. $\rho^F = 0\%$).
2. Recapitalisation target equal to minimum capital requirements plus a half of the combined regulatory buffers (i.e. $\rho^{RT} = \rho^M + \frac{1}{2}\rho^{CB}$).
3. Unfair conversion rates.

The **bad (*status quo*) secondary parameters** are:

1. Debt excluded from bail-in with a time to maturity less than 7 days.
2. Current loss absorbing requirements.
3. Certain bail-in design.

The **good secondary parameters** are composed of:

1. Debt excluded from bail-in with a time to maturity less than 1 year.
2. Double the current loss absorbing requirements (but short-term debt is excluded from bail-in, see above item).
3. Uncertain bail-in design.

The second way in which we promote compositional ease and cross-comparison of findings is to consistently displaying results.

y-axis:

In order to appraise the stability impact of the bail-in design we show a measure of systemic risk on the y-axis. A suitable systemic risk measure is the average size of the contagious asset losses (in trillion euros) in the European banking system. To place losses in perspective the total asset value in the banking system prior to shocks is 22 trillion

euros.⁶³ ⁶⁴ The average contagious asset loss is computed across $N=50$ simulation runs, where – consistent with Gai & Kapadia (2010), Gai et al. (2011), Caccioli et al. (2014), Paulin et al. (2018) – in each run the asset holding network and interbank network are randomly reconstructed. In line with Cont et al. (2010), we intentionally exclude the impact of the initial stress scenario from this measure – to be able to evaluate the endogenous systemic-risk amplification of the shock rather than the severity of the shock itself.

x-axes:

Since our paper seeks to study the impact of the bail-in design both in case of idiosyncratic failures and in the case of system-wide crises, we always display our findings as a function of each. On the left x-axis, we display the number of idiosyncratic bank failures, by picking the $x = 1, \dots, 5$ largest banks by asset value to exogenously default.⁶⁵ On the right x-axis, we show a system-wide shock in severity ranging from zero to two times the severely adverse initial scenario applied in the 2018 European Banking Authority stress test.⁶⁶ Critically, the left and the right subplots of each figure answer two supplementary research questions. The left plot asks if the bail-in design could induce contagion. Whereas the right subplot inquires how the bail-in design might either exacerbate or contain earlier contagious losses incurred in the system-wide tribulation. Aside from this purpose, the x-axes also enable the sensitivity analysis of results to the shock size and type. It appears that such sensitivity analysis is key: although we find that our results are robust to different shocks, the systemic damage is heavily dependent on the adverse shock scenario.

⁶³We note that the ‘average extent of contagion’, a commonly used systemic risk measure that counts the average fraction of bank failures in a systemic event (Gai & Kapadia (2010), Gai et al. (2011), Caccioli et al. (2014), Paulin et al. (2018)), is a suitable instability yardstick in financial systems where liquidation is the preferred mode of bank failure. But, it does not suit well for the systems where bail-in is the norm. The reason is that we only concern about bail-ins occurring in so far as they impose significant losses on the system. With some bail-in designs losses might be contained while bail-ins may occur frequently. In contrast, we invariably worry about disorderly liquidations of SIBs because they almost always impose large creditor and systemic losses. Hence, counting failures in the case of disorderly liquidations makes sense, but doing so for bail-ins may give a distorted view of financial distress.

⁶⁴We also note that measuring systemic risk by the losses as a percentage of the initial CET1 equity in the banking sector, as for instance Cont & Schaanning (2017, 2019) have done, does not suit a bail-in study. Bailed-in banks may regain capital even though their bail-in could impose large system-wide losses.

⁶⁵An exogenous bank failure is brought about by imposing an asset loss to a bank’s external assets such that its risk-weighted capital ratio falls $f\%$ below the failing-likely-to-fail ratio at which a bank is bailed in. Given that a bank usually suffers an asset loss that cause it to fall well-below the bail-in threshold and given that asset losses typically are only fully recognised once the bail-in process is started (Chennells & Wingfield (2015)), we believe 4% is reasonable.

⁶⁶See: <https://eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing/2018>.

4.7.1 Stability Impact of the Primary Bail-In Design

Our main research question is: what are the systemic implications of the bail-in design? Figure 4.14 is the first central plot of our paper in answer to this question.




Figure 4.14: Shows the stability impact of the ‘primary’ bail-in design.

It suggests that financial stability hinges on the primary bail-in design. Figure 4.14 reveals how the value of a set of ‘primary’ bail-in parameters – which include the failure threshold, the recapitalisation target and the debt-to-equity conversion rate – sweepingly alters systemic risk. A notable exception is the case of smaller SIB failures. Figure B.1 in Appendix B.4 shows that bail-in usually works for idiosyncratic failures of *smaller* European SIBs, implying that the elected ‘primary’ parameters do not matter much here – which is consistent with empirics. For all other cases, however, the primary bail-in parameters *do* crucially matter. The left plot in Figure 4.14 exhibits that ill-designed bail-ins may induce widespread contagion if *larger* European SIBs idiosyncratically fail. The right plot in Figure 4.14 finds that bank bail-ins may heftily exacerbate financial fragility in financial crisis episodes if ‘primary’ parameters are ‘badly’ chosen.

Strikingly, we witness a phase shift from an unstable (orange lines) to stable system (blue lines), if resolution authorities do choose ‘good’ primary bail-in parameters. For large idiosyncratic failures, ‘good’ parameters cut down contagion (by up to 3.5 trillion euros relative to the ‘bad’ case) so as to altogether extinguish it. For system-wide shocks, ‘good’ parameters curb instability rather than exacerbate it as ‘bad’ parameters do (reducing contagious asset losses by down to 1.5 trillion euros). This result qualitatively holds even if epidemic contagion causes multiple banks to be bailed-in amid pervasive

distress.

To place the first main result in perspective, it is essential to benchmark the systemic implications of a ‘good’ and ‘bad’ primary bail-in design against the main alternative modes of dealing with bank failure: disorderly liquidation and bail-out. Figure B.2 & B.3 in Appendix B.4 confirm regulators’ hopes that bail-in substantially brings down contagion relative to the case where failed banks are disorderly liquidated: a *sine qua non* for finding bail-in a preferential failure mode over disorderly liquidation. The reaped stability benefits of changing the failure method from disorderly liquidation to bail-in are much less substantial when the primary parameters are ‘badly’ chosen: it achieves a one-trillion-euro reduction in losses in the case of idiosyncratic failures, rather than the 4.5 trillion euro decline that ‘good’ parameters result in.

Figure B.2 & B.3 in Appendix B.4 also confirm our expectation that bail-out outperforms bail-in on the financial stability front, which was anticipated. Bail-outs – though politically and economically undesirable – recapitalise a bank by inserting cash without imposing losses on creditors. These are financial benefits that bail-in does not enjoy. It is worth noting that there are in fact two forms of bail-in. ‘Bail-in disorderly liquidation’ is the form where a bailed-in bank whose loss-absorbing debt proves insufficient to recapitalise it to meet its authorisation conditions is disorderly liquidated. ‘Bail-in bail-out’ is the form where a bank in such circumstance is further bailed-out instead. As noted in the default settings, we work under the premise that a bank is disorderly liquidated; because we are interested in evaluating the systemic success of bail-in in absence of government aid. Unsurprisingly, from a financial stability perspective ‘bail-in bail-out’ gains the palm. Overall, this result implies that a well-designed bail-in, in particular,⁶⁷ is far more efficacious in taming systemic risk than failing a bank disorderly. It is also much more desirable than a bank bail-out.

We have seen the sway that the primary bail-in parameters hold over stability. And weighed up the merits of a ‘good’ and ‘bad’ bail-in against other failure modes. But, we have yet to disentangle how each primary policy parameter bends the system towards stability or peril. We proceed to focus on the influence that the first primary bail-in parameter, the failure threshold (also referred to as the FLTF ratio), exerts on systemic risk.

⁶⁷More so than an ill-designed bail-in.

4.7.1.1 Stability Impact of the FLTF Ratio

Figure 4.15 unfolds the stability impact of the failing-likely-to-fail ratio (ρ): *the closer a bank's capital ratio is to insolvency before it is bailed-in, the higher the system's contagious losses are.*

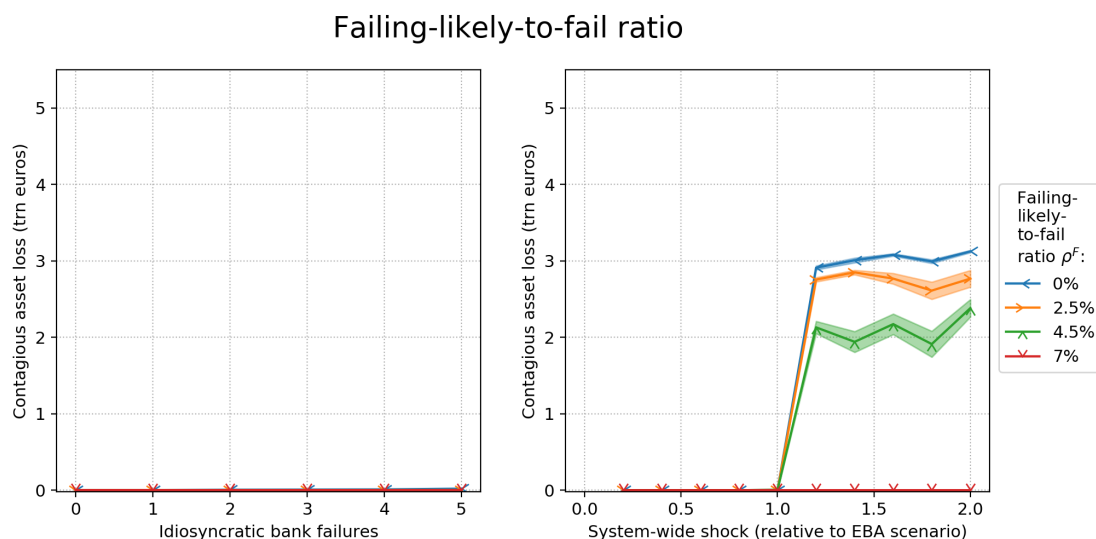


Figure 4.15: Shows the stability impact of the first ‘primary’ bail-in parameter: the failing-likely-to-fail ratio (or in other words, the failure threshold).

This finding appear counterintuitive. On the one hand, postponing the bail-in of a bank until it is nearly insolvent tends to reduce the number of bank failures, as less banks see their capital ratios fall that far. Even though this may seem beneficial, from the perspective of systemic risk it is usually not. Even though bail-ins may happen less frequently, the damage they do is more severe. Banks that are bailed-in only when their capital ratio falls below a FLTF threshold that is set close to insolvency (i.e. not far off $\rho = 0\%$) are very likely to be somewhat, or quite possibly severely, insolvent at the start of a bail-in.⁶⁸ This impairs stability. Recapitalising an insolvent bank leaves the regulators with no other option (without violating the no-creditors-worse-off principle or the preservation of hierarchy of claims principle) but to apply pure write-downs to reduce liabilities to eliminate the insolvency. Only after the insolvency is wiped out, is the regulator in a position to compensate loss-bearing creditors with equity claims via a positive debt-to-equity conversion rate. Hence, forced pure write-downs upon insolvent banks impose net exposure losses on junior creditors. This may galvanise exposure loss contagion.

⁶⁸That is for two reasons. First, an asset loss is always comes in lumps, entailing that a bank that falls below the FLTF cut-off due to an asset loss typically plummets significantly rather than slightly below the threshold. Second, asset losses are frequently fully recognised only when the bail-in has commenced (Chennells & Wingfield (2015)) – though we do not model this here.

On the other hand, a bail-in that is launched once a bank is still solvent – which is more likely the case if the FLTF ratio is set high enough – maintains stability. In such case, a bank can be recapitalised by applying solely positive debt-to-equity conversion rates. This implies that creditors do not have to suffer net exposure losses: a haircut is replaced with an equal equity claim, resulting in no net loss, if conversion rates are selected to be fair.

Surprisingly, even if conversion rates are ‘unfair’ as is the case under a ‘bad’ bail-in design – so creditors of solvent banks suffer some losses too (see Figure B.4 in Appendix B.4) – we observe that an ‘early’ bail-in triumphs a ‘late’ bail-in. Presumably, the reason is that an early recapitalisation restores a bank to good health, making it less likely to act in destabilising ways.

Overall, this result implies that regulators who bail-in banks ‘too late’ risk jeopardising their stability. We will now proceed to unknot the systemic ripple effects of the second ‘primary’ bail-in parameter, the recapitalisation target.

4.7.1.2 Stability Impact of the Recapitalisation Target

Figure 4.16 unveils the stability impact of the second ‘primary’ bail-in parameter, the recapitalisation target. *It shows that contagion tends to be more pronounced, if bailed-in banks are less strongly recapitalised.*

Recapitalisation target

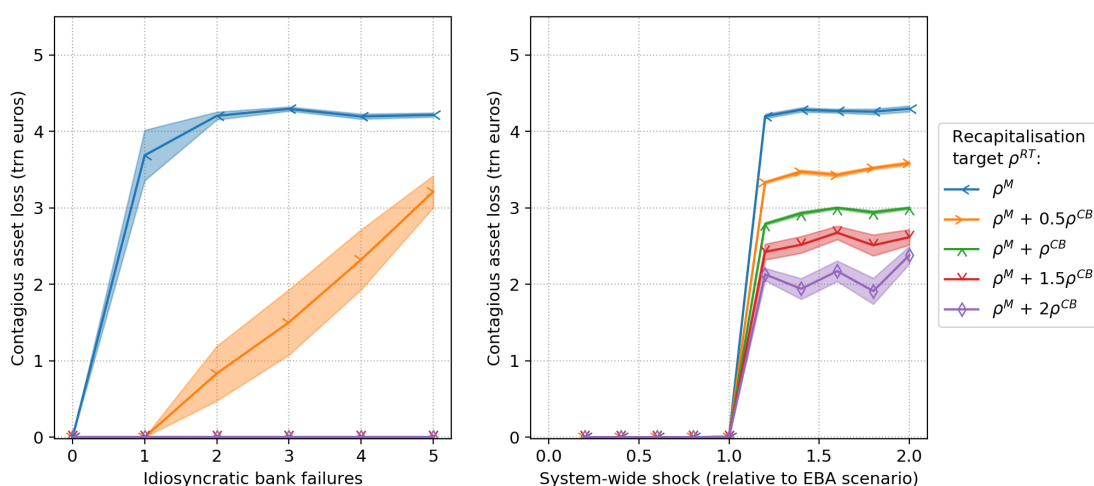


Figure 4.16: Shows the stability impact of the second ‘primary’ bail-in parameter: the recapitalisation target ρ^{RT} . The label ρ^M means that a bank is recapitalised to just meet its authorisation condition: its minimum capital requirement of $\rho^M = 4.5\%$ of CET1 equity relative to risk-weighted assets. Whereas the label $\rho^M + \rho^{CB}$ means that a bank is recapitalised to not only meet its minimum capital requirement ρ^M , but also (just about) comply with its combined regulatory capital buffer ρ^{CB} .

This result is perhaps surprising. A higher recapitalisation target implies a larger recapitalisation amount (recall equation 4.24). It thus entails more colossal haircuts resulting in larger losses that afflict creditors. So, one might think that increasing the recapitalisation target increases the exposure losses to creditors. Yet, this does not have to be true if the conversion rates are fair – as they are in a well-designed bail-in. In such cases, any haircut in the recapitalisation phase is replaced with an equal claim of the bank’s equity, so that no net losses are suffered. Obviously, the creditor is now lower in the insolvency hierarchy, so that raises its exposure to future – though not current – losses, were it to hold on to the equity stake rather than convert it to cash. Also, if the creditor were to sell the equity stake, he may incur liquidation costs. In sum, on account of no net losses being (directly) suffered when conversion rates are fair, a higher recapitalisation target does not inflict higher losses on creditors, and therefore does not compromise stability.

One might speculate that this result may not uphold in the face of unfair, rather than fair, conversion rates. Unfair conversion rates exact pure write-downs, and thus losses, on most creditors except the most senior ones that were included in the bank recapitalisation. Hence, a higher recapitalisation target means more severe losses for most, and higher profits for a very few, presumably wreaking more instability.

Unexpectedly, even when unfair conversion rates are employed as part of a ‘bad’ bail-in design, we observe that a lower recapitalisation target tends to aggravate instability

(see Figure B.5 in Appendix B.4). Why is this so? The reason is twofold. First, banks that are inadequately recapitalised are prone to funding flights. It is widely known that a market-imposed minimum capital constraint exists (Burrows et al. (2012)): banks whose capital ratio falls below this threshold typically encounter a halt to the roll-over of their short-term funding. Funding shocks typically elicit disorderly liquidations to meet the repayment obligations. So, a stronger recapitalisation minimises the risk that new asset losses following a bail-in cause the bank to tumble below the market constraint, fortifying stability.

In our model, specifically, the market constraint comes into force via creditors who pull back their maturing bail-inable debt whenever its expected loss or VaR loss exceeds some threshold. This expected-loss threshold is more easily exceeded if the bank is ill-capitalised, as the expected loss increases in the likelihood of a costly bail-in; or, if creditors do not have enough information to price bail-inable debt and thus compute expected or VaR losses. In such cases, the creditors in our model use a rough rule of thumb to pull back funding if the bank's capital ratio is very close to its failure threshold.

The second reason why a lower recapitalisation target tends to aggravate instability is that ill-capitalised banks are susceptible to act in destabilising ways to further bolster up their capital ratio. Banks typically maintain a stable 'capital ratio target' (Adrian & Shin (2010)). And hold a 'buffer' below which they will take actions to, over time, revert to their target capital ratio (Cont & Schaanning (2017), FSB (2017), Farmer et al. (2020)). There are various reasons for this; one of which is the survival instinct to not fall below the market-imposed and regulatory-imposed minimum capital constraints, which can lead to failure. Another is not to eat into the regulatory capital buffers, which come with restrictions on the ability to make discretionary dividend and bonus payments (Goodhart (2013), Farmer et al. (2020)). Though such motives may be individually rational, they can be destabilising from the perspective of the system as a whole. Delevering to shore up the risk-weighted capital ratio entails liquidating assets with a non-zero risk weight.⁶⁹ Liquidating assets depresses the value of common asset holdings, if tradable assets are sold at discounts. Or, it induces funding shocks, if maturing loans are withdrawn. Indeed, the risk-weighted delevering in our model provokes common asset holding contagion and funding contagion. Hence, in this second way too, the by-product of a poor bank recapitalisation is often system-wide contagion.

⁶⁹Or issuing new equity. But this is typically not feasible during financial distress (Greenwood et al. (2015)).

To sum up, this finding implies that resolution authorities who are bent on safeguarding stability should richly recapitalise a bank. We next proceed to decipher the system-wide effects of the third primary bail-in parameter, the debt-to-equity conversion rate.

4.7.1.3 Stability Impact of the Debt-to-Equity Conversion Rate

Figure 4.17 untangles the stability impact of the third primary bail-in parameter, the debt-to-equity conversion rate. *It shows that fair conversion rates preclude contagion in the case of idiosyncratic cases of SIB failures and curtail contagion in system-wide crises, while unfair conversion rates severely deepen financial distress in both circumstances.*




Figure 4.17: Shows the stability impact of the third ‘primary’ bail-in parameter: the debt-to-equity conversion rate.

Unsurprisingly, the benchmark case of ‘zero conversion rates’, merely plotted to facilitate comparison, undermines stability most. This is logical, since zero conversion rates imply pure write-downs – meaning that each creditor in the bail-in suffers maximal exposure losses.

Fair conversion rates fare best. This is understandable, given that with a fair conversion rate a loss-bearing creditor in the recapitalisation phase is fully compensated by an equally sized equity claim. As a result creditors suffer no net losses due to the bail-in mechanism, thereby ruling out exposure loss contagion and the further contagious ampli-

fications that may ensue.⁷⁰ This does not mean, however, that a fair conversion rate can fully pre-empt exposure loss contagion. If a bank is insolvent at the onset of bail-in, then the regulator usually has no choice – without violating the conversion rate principles — but to impose pure write-downs on the most junior creditors until the bank is taken out of insolvency.

Unfair conversion rates fare worse than fair conversion rates, but – perhaps needless to say – better than zero conversion rates. Why are unfair rates a detriment to financial stability, while fair rates are not? On account of the unfair distribution of losses that unfair rates impinge on creditors, some debtors face excessive losses, while others reap excessive profits. Junior creditors are lined-up to face pure write-downs – whereas under fair rates they would have been compensated enduring no net loss. While senior creditors are orchestrated to benefit; they acquire a larger equity claim than under the fair scheme, and so make a net profit. Imposing excessive losses on junior creditors threatens to them to transgress their constraints – be it internal, regulatory, or contractual. This is especially probable in a financial crisis, in which institutions typically drift close to their constraints (Aymanns et al. (2018)). The accrual of net profits to a small group of the most senior creditors is not only unfair, but also proves inconsequential with respect to ameliorating stability. Most it can do is to move senior creditors further away from their constraints, which is a good consequence. However, that evidently does not outweigh the cost of bringing the rest of the creditors closer to their constraints – as we observe. Institutions that encroach their constraints or even breach them are prone to enfeeble stability by acting in contagious manners. For all these reasons, fair conversion rates – which impose no more than the minimum necessary net losses to wipe out insolvency – outperform unfair conversion rates on the financial stability front.

The first central result in this paper (see Figure 4.14) showed that stability hinges on the triple of primary bail-in parameters, but did not reveal the role herein of the conversion rates. In this plot, we disentangled the stability impact of the third and final primary bail-in parameter, the debt-to-equity conversion rate. To more exhaustively answer our basic research question concerning the systemic implications of the bail-in design, it is necessary to understand how resilience may be further pivoted by the secondary, more structural, bail-in parameters. This brings us to our second central result in this paper.

⁷⁰Obviously, equity claims can revalue after bail-in, risking losses at a later stage if the equity claim is not converted to cash.

4.7.2 Stability Impact of the Secondary Bail-In Design

Figure 4.18 displays the financial stability impact of also tuning the secondary bail-in design on top of calibrating the primary design. Comparing the case where both primary and secondary parameters are ‘well’ or ‘ill’ calibrated (this figure) the case where solely the primary parameters are suitably or ill-suitably calibrated (the first main result, Figure 4.14) we observe that *the selection of secondary parameters may significantly stretch the wedge between ‘good’ and ‘bad’ primary parameters.*

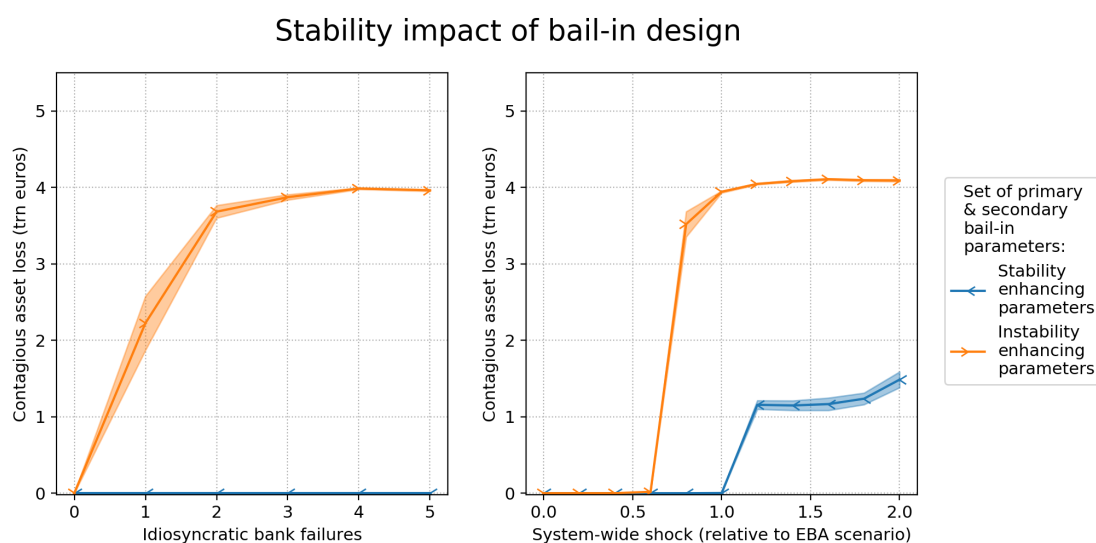


Figure 4.18: Shows the stability impact of also calibrating the ‘secondary’ bail-in design on top of tuning the ‘primary’ bail-in design.

Interestingly, while even ‘good’ primary bail-in parameters were not able to fully contain contagion during system-wide distress (see Figure 4.14), adding ‘good’ secondary parameters achieves to suppress contagion to acceptable levels. On the other hand, tallying ‘bad’ secondary parameters to already ‘bad’ primary parameters, aggravates financial instability across shock types (see Figure 4.14).

This result implies that regulators who suitably calibrate secondary parameters stand strong to tame systemic risk – even under severely adverse systemic scenarios. On the other hand, regulators who do not, risk structurally eroding resilience. Though this result elegantly captures the joint impact of the secondary parameters, it shrouds the part each of the three ‘secondary’ parameters plays in amending stability. The aim of the next results is to unveil this.

4.7.2.1 Stability Impact the Loss Absorbing Requirements & Debt Exclusions from Bail-In

Figure 4.19 uncovers the stability impact of debt exclusions from bail-in. *It shows that systemic risk drops sharply if debt with a time to maturity (TTM) less than a year is excluded from bail-in.*

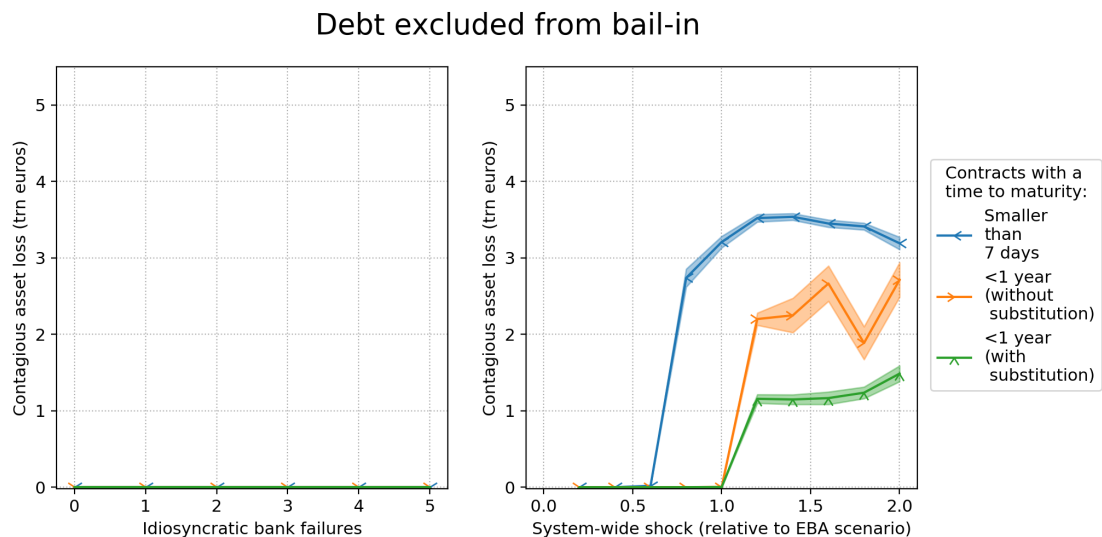


Figure 4.19: Shows the stability impact of two of the ‘secondary’ bail-in parameters: debt exclusions and loss absorbing requirements.

Remarkably, excluding short-term debt (orange line) succeeds in subduing contagion to (more) acceptable levels in periods of system-wide turmoil. Whereas even good ‘primary’ parameters were not able to proficiently subjugate contagion in such cases (see the first central plot, Figure 4.14).

This result runs contrary to expectations. Excluding debt from the application of the bail-in tool reduces the stack of bail-inable debt that regulators have at their disposal to absorb losses and recapitalise a failing bank. This makes it more likely that a bailed-in bank cannot be recapitalised to meet its authorisation conditions (i.e. its minimum capital requirements) and must harmfully be disorderly liquidated (our default assumption) or bailed-out instead.

On the other hand, excluding short-term debt should not dent stability, if regulators adequately calibrated the ‘loss absorbing requirements’ – which are meant to ensure that banks have sufficient ‘loss absorbing capacity’. According to current regulatory judgement, banks that comply with their loss absorption requirements should already hold enough loss-absorbing debt. So, including short-term debt in bail-in would be superfluous. (We recall that only debt with a TTM greater than one year counts towards the loss absorbing requirements, so excluding short-term debt does not alter the loss absorbing

requirements.) So, provided that regulators have calibrated the requisite loss absorbing requirements well, excluding short-term debt should at least do no harm.

Why do we then observe that excluding short-term debt actually does good? The explanation is that roll-over risk of bail-inable debt falls sharply, if only long-term bail-inable debt remains – this is understandable. Short-term creditors will feel secure to keep lending to a bank, even if it is on the brink of a bail-in, if they know that their contracts are by definition excluded from bail-in losses. Indeed, the creditors in our model only halt funding to bank's whose bail-inable debt is expected to suffer grave losses, and these creditors never withdraw funding of non-bail-inable debt purely out of loss concerns. Excluding short-term debt not only makes bail-inable debt collapses in short anticipation of an impending bail-in impossible, but also eliminates the consequent risk of disorderly liquidations of assets (at fire sale prices) to meet the obligation to pay back the withdrawn bail-inable debt. Unfortunately, under the BRRD that governs EU bail-ins, only debt with a TTM less than 7 day is excluded – corresponding to the worse-off blue line in Figure 4.19 .

Figure 4.19 further demonstrates that if the excluded short-term bail-inable debt (with a TTM between 7 days and 1 year) is accompanied by an equal increase in long-term bail-inable debt with a TTM greater than a year, then resilience further upgrades. Such a shift can be achieved by increasing banks' loss absorbing requirements, since only long-term bail-inable debt with a TTM greater than a year counts towards the requirements, as we previously noted. Hence, our result suggests that *increasing the loss absorbing requirements significantly shrinks systemic risk*.

This finding belies regulatory judgement that the current loss absorbing requirements are sufficiently high. Our finding highlights the merits of increased loss absorbing requirements in terms of notching up banks' bail-inable debt pile; which reduces both the risk of necessitated alternative means of dealing with failure and the risk that an ill-recapitalised bank will be prone to funding flights and engage in destabilising actions to prop up its capital ratio (e.g. delevering). Our finding also bespeaks a more structural point – which bears no reference to bail-in in particular. Increasing the loss absorbing requirements structurally shifts the maturity profile of banks' liabilities towards more stable long-term funding. Banks who issue longer-term debt and invest in longer-term debt issued by other banks are both less vulnerable to funding shocks and less infectious in imposing funding shocks on others. The reason is that long-term debt cannot easily be withdrawn.

The two findings based in Figure 4.19 imply that regulators can significantly subdue contagion by excluding short-term debt from bail-in and by increasing the loss absorb-

ing requirements. And that tuning these secondary parameters well is especially critical in system-wide crisis episodes: in such cases, as our first main result (see Figure 4.14) showed, even good primary parameters are not able to curb contagion to tolerable levels.

Having unmasked the stability contribution of two of the ‘secondary’ parameters, we proceed to examine how the third and last ‘secondary’ parameter turns stability. It proved most efficient to do this as part of our third central result – which focuses on studying the role of each contagion mechanism and bail-inable debt holder in modifying the systemic footprint of the ‘primary’ and ‘secondary’ bail-in design.

4.7.3 Contagious Amplifications of the Bail-in Design

Our third and final main figure (Figure 4.20) attests that it is imperative to take multiple contagion mechanisms and non-bank holdings of bail-inable debt into account to avoid underestimating the system-wide implications of a bail-in design. To reduce our results to essentials, we show this result here for the default settings of ‘good’ primary parameters and a system-wide shock – for now not tuning secondary parameters. (In Appendix B.4 we also show this result holds for a ‘bad’ bail-in design and idiosyncratic shocks.)




Figure 4.20: Shows the contagious amplifications of the primary bail-in design – as a function of system-wide shock.

Our result shows that merely considering the exposure loss contagion that could ensue from bail-ins – as Hüser et al. (2017) have done – would falsely suggest the EU financial system becomes resilient to severe system-wide shocks with the introduction of bail-ins to deal with failures of SIBs. Instead, if, rightly, four more prevailing contagion mechanisms

are taken into account, then stability may be compromised for severe enough shocks – especially with a poor primary bail-in design (see Figure B.6 in Appendix B.4). The reason is that contagion mechanisms may mutually amplify each other, as we next explain.

Exposure Loss Contagion & Overlapping Portfolio Contagion

Exposure loss contagion amplifies overlapping portfolio contagion, and visa versa. Exposure losses incurred in a bail-in may prompt institutions to delever to strengthen their capital ratios. To delever, institutions may liquidate tradable assets at fire-sale prices. This will lead other institutions that hold these assets in common to suffer mark-to-market losses. In turn, this can trigger these institutions to take destabilising actions, such as embarking on further fire sales. In our model, institutions will delever when their capital ratio falls below their safety buffer (Farmer et al. (2020)). When this happens, they will seek to gradually return to a more prudent capital ratio target.

On the flipside, overlapping portfolio contagion can also foster exposure loss contagion. Mark-to-market losses may shift a bank’s capital ratio below its failure threshold. The consequent bail-in could impose exposure losses.

Figure B.6 in Appendix B.4 shows that the two-sided amplification between exposure loss contagion and overlapping portfolio contagion is substantially magnified if primary bail-in parameters are ill-chosen. It is not difficult to understand why: exposure losses tend to be more extreme when conversion rates are ‘unfair’. Because in such a case, most junior creditors who bear haircuts will not be compensated by equity claims, making their *net* exposure loss more grievous.

Fragility also easily escalates when the bail-in recapitalisation target is low, as it is under the ‘bad’ primary settings. Following a weak recapitalisation, banks easily suffer further losses that shift their capital ratio below their safety buffer, triggering delevering via for instance overlapping portfolio contagion.

Low failure thresholds, under an ill-designed primary bail-in, further reinforce the amplification between exposure loss contagion and overlapping portfolio contagion. This is because more banks are insolvent at the start of bail-in when thresholds are low. Raising a bank out of insolvency requires applying pure write-downs. This troubles counterparties with larger exposure losses, which could in turn more sharply stir up overlapping portfolio contagion.

Exposure Loss Contagion, Funding Contagion & Overlapping Portfolio Contagion

Exposure loss contagion, funding contagion and overlapping portfolio contagion are mutually amplifying. Exposure losses afflicted in a bail-in may motivate a bank to delever to a stronger capital ratio, as just discussed. An alternative way to delevering by liquidating tradable assets is to withdraw maturing loans. The counterparties of the bank are then obliged to repay the notional of the loan. If they do not have a thick enough cash buffer to do so, they are forced to disorderly liquidate assets to raise cash. If they choose to halt rolling over loans to their counterparties, funding contagion arises. If instead, they opt to fire sell assets, overlapping portfolio contagion may emerge. So, bail-ins can induce funding contagion and overlapping portfolio contagion in sequel.

In the other direction, a system-wide shock may spur funding and overlapping portfolio contagion, which may in turn precipitate bail-ins that could induce exposure losses.

In our model, which builds forth upon [Farmer et al. \(2020\)](#), banks maintain a ‘pecking order’ that determines which type of assets they will liquidate first. In line with the literature (e.g. [Halaj \(2018\)](#)), we assume that they start with liquidating assets with the lowest liquidation cost. Hence, the banks in our model will opt to liquidate maturing loans before they undertake to liquidate tradable assets.

Revaluations of Bail-inable Debt & Interaction with Other Contagion Mechanisms

The bail-inable debt value plummets if two scenarios are jointly true. One, a bail-in becomes more likely. Two, the debt pay-off in a bail-in is appreciably lower in a bail-in than in the absence of one. This is logical, since the value of a contract is given by the discounted value of future pay-offs. If many future pay-offs are low, the contract’s value will reflect it.

The set of circumstances under which these two scenarios are simultaneously true determines when one should expect sharp revaluations. For the first scenario, a bail-in becomes more likely if a bank’s capital ratio falls closer to the failure threshold. Hence, the mixture of overlapping portfolio contagion, funding contagion and exposure loss contagion, which may pull down a bank’s capital ratio, accelerates declines in the bail-inable debt value. For the second scenario, the bail-in pay-off dwindles if conversion rates are unfair and the bank is bailed-in late – since these ‘primary’ bail-in settings inflict exposure losses destitute of compensation. Hence, ‘bad’ primary bail-in parameters could intensify destabilising spirals – wherein debt revaluations prompt more contagion, foster more debt revaluations and so on. Therefore, ‘bad’ primary parameters tend to ignite mutual amplifications among contagion mechanisms.

We have explained how the three forms of contagion can stimulate bail-inable debt prices to fall. In reverse, declines in the bail-inable debt value lead to mark-to-market

losses. This could precipitate bail-ins and delevering, which in turn bring about the aforementioned forms of contagion.

Revaluations of Bail-inable Debt & Runs

Bail-inable debt revaluations may spur creditors to stop rolling over bail-inable debt, imposing destabilising funding shocks.

A revaluation of a bail-inable debt contract not only imposes mark-to-market losses on banks and non-banks who hold that contract, but also increases the expected loss – equal to the notional minus the value of the contract – and VaR ‘worst-case’ loss on that contract. As the expected loss and VaR loss on a bail-inable debt contract increases, a creditor is more inclined to withdraw this debt if it matures to reduce exposure to the loss-heavy investment.

In our model funding is cut when the expected loss or VaR loss exceeds a certain threshold. This inflicts a funding shock on the bank that issued the bail-inable debt. In turn, the shock may prompt the bank to undertake costly asset liquidations to fulfil the obligation to repay the contract; thereby kindling overlapping portfolio contagion or funding contagion, for instance.

Hence, sharp declines in the bail-inable debt values (brought about by widespread contagion) may precipitate destructive bail-inable debt collapses – caused by creditors who do not roll-over this debt for fear of losing money. A bail-inable debt collapse not only spurs disorderly asset liquidations, but also makes it harder to successfully recapitalise the bank would it need to be bailed-in. According to [Avgouleas & Goodhart \(2015\)](#) the risk of bail-inable debt collapses is one of the most salient reasons why bailing-in banks in system-wide crises might destroy rather than ameliorate stability.

Figure 4.21 shows that he is right. We observe that contagious asset losses in system-wide crises fall dramatically if the ‘secondary’ bail-in design is also well-tuned, in addition to the calibration of the ‘primary’ design (shown above in Figure 4.14). This drop can be explained by collapse-prone short-term debt that is excluded from bail-in and replaced with collapse-proof long-term bail-inable debt, which reduces the risk of evaporating bail-inable debt piles. More generally, suitably tuning the ‘secondary’ bail-in design, on top of the ‘primary’ design, strikingly shrinks contagious amplifications among contagion channels.

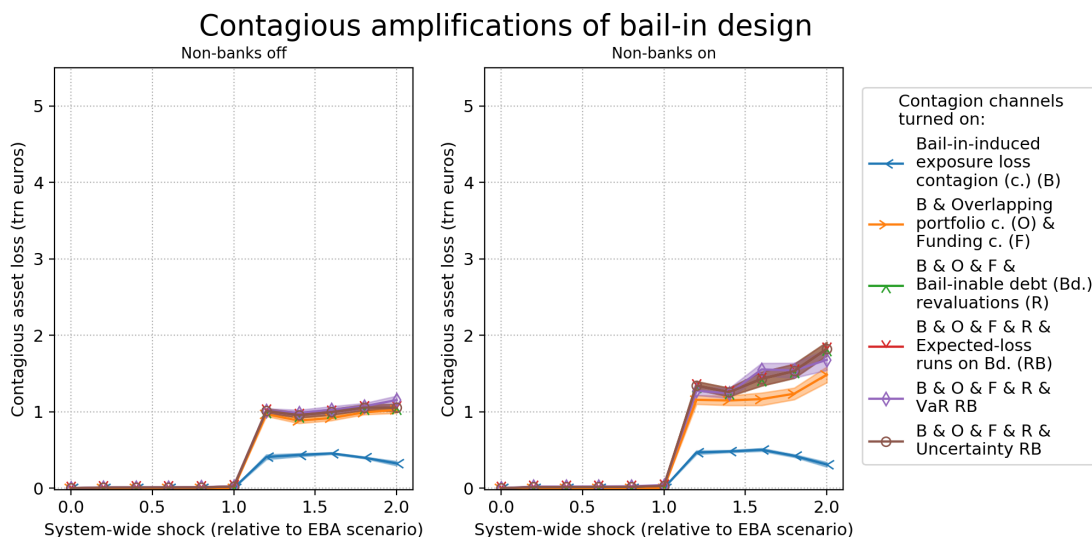


Figure 4.21: Shows the contagious amplifications of fine-tuning the ‘secondary’ bail-in design, in addition to the ‘primary’ design – as a function of the system-wide shock.

Uncertain Bail-In Design

We now revert to our endeavour of examining how the third and last ‘secondary’ parameter – bail-in uncertainty – influences stability. A funding withdrawal of a bail-inable debt contracts inspired by increases in the contract’s expected loss and VaR loss require the creditor to be able to *price* bail-inable debt. This is only possible if the regulator has *a priori* precisely proclaimed the primary bail-in parameters that will apply in a prospective bail-in of the bank who has issued the bail-inable debt contract. In the absence of such information – and instead in the presence of loose avowals about the ranges of these parameters – pricing bail-inable debt with any accuracy, and thus estimating loss exposures becomes impossible. Therefore, creditors instead might resort to crude measures to decide when to stop rolling over short-term bail-inable debt – such as a cut-off capital ratio, below which bail-inable funding is retracted, as is the case in our model. If substantial ‘uncertainty’ exists about the primary bail-in design, adverse shocks may not only induce an indiscriminate collapse of bail-inable debt across priority classes, irrespective of the prospective losses inflicted in each priority class. It can also provoke bail-ins by inducing a downward spiral of capital ratios brought about by calamitous asset liquidations to meet the ‘uncertainty’ funding shocks. If uncertainty in the bail-in design prevents creditors from seeing that regulators will in fact apply ‘good’ primary parameters (making prospective losses low), such indiscriminate runs would be not only unwarranted, but also unnecessarily cataclysmic.

Non-banks

This figure (Figure 4.20) illustrates that ignoring non-bank holdings of bail-inable debt

may lead one to materially underestimate the systemic impact of bail-ins on the banking system.⁷¹ The reason is that non-banks hold the overwhelming majority of bail-inable debt. Thereby non-banks are susceptible to contagiously feedback onto the banking system any bail-inable debt losses thrust upon them. Non-banks hold much bail-inable debt, because the loss absorbing requirements only count debt that is not cross-held by banks. Hence, all the eligible debt that is short of a bank’s loss absorbing requirements is necessarily held by non-banks. Only eligible debt in excess of the requirements is potentially cross-held by banks.

Given that a substantial amount of bail-inable debt is held by non-banks, any exposure loss resulting from a bail-in tends to be disproportionately felt by non-banks. Also, since debt must be long-term to count towards the loss absorbing requirements, non-banks are particularly exposed to vicious revaluations of bail-inable debt – especially when primary bail-in parameters are ‘bad’ (see Figure B.6 in Appendix B.4). Debt contracts with longer maturities are more prone to revaluations, since, other things equal, a bail-in is more likely to occur over a longer than shorter horizon.

Exposure losses and mark-to-market revaluation losses in the non-banking sector can either befall leveraged non-banks or non-leveraged non-banks.

Leveraged non-banks, impaired by exposure losses and mark-to-market revaluation losses, are likely to delever to, for instance, meet margin requirements. In line with [Cont & Schaanning \(2017\)](#), these leveraged institutions in our model delever to return to target whenever their leverage falls to 90% of their leverage target. The consequent disorderly liquidation of tradable assets by leveraged non-banks regresses to hurt banks that hold these assets in common, potentially setting off pervasive contagion in the banking system. As non-banks become more highly leveraged they grow more susceptible to forced delevering. Hence, non-bank holdings of bail-inable debt are a more pernicious threat to banking sector stability if non-banks are highly leveraged. We show this in Figure B.7 in Appendix B.4.

On the other hand, most non-leveraged non-banks impaired by exposure and mark-to-market losses are unlikely to bounce back losses onto the banking sector, which improves resilience. This gain in stability, however, comes with a cost: real-economy individuals will tend to shoulder the losses of bank bail-ins. The reason is that non-leveraged non-banks are institutions such as pension funds and insurance companies, in which real

⁷¹Note that our measure of systemic risk estimates the contagion losses that befall the banking system. Purposefully, it does not measure contagious losses afflicting the non-banking sector. It solely aims to measure to what extent contagious feedback loops between banks and non-banks could exacerbate banking-sector losses.

economy individuals have invested their savings and protected themselves against casualties.

Hence, as a larger percentage of the bail-inable debt held by the non-banking sector is possessed by non-leveraged non-banks, systemic risk tends to fall (see Figure B.8 in Appendix B.4). However, the repellent downside is that real-economy individuals will suffer more losses (see Figure B.9 in Appendix B.4). Aside from the consideration that imposing losses on taxpayers may be undesirable, it may also be politically difficult to pull off – as previous experience has shown (WBG (2017)).

The result on non-banks implies the following. While regulators evidently believe that discouraging cross-holdings of bail-inable debt by banks improves stability – as is clear from the design of the loss absorbing requirements – they may have deluded themselves: it merely pushes bail-inable debt holdings out into the non-banking system, but does not shield the banking system from bail-in-induced contagion. It begs the question: who can best bear the bail-in losses?

4.8 Link to the Literature

Our contribution adds to the nascent network literature on the systemic effects of bail-in. Klimek et al. (2015) employ an agent-based network model to evaluate the economic and financial ramifications of bail-in. They compare its performance against other resolution mechanisms. Hüser et al. (2017) evaluate the systemic implications of bail-in in the EU, drawing on a calibrated multi-layered network model to bank debt and equity cross-holdings. These papers neither investigate the systemic impact of the bail-in design, nor include multiple contagion mechanisms and non-banks in this analysis. Instead, they take the bail-in design as is and merely explore the repercussions of exposure loss contagion. By ignoring multiple interaction contagion mechanisms, they risk underestimating the systemic footprint of the bail-in design. Though bail-in has been designed with systemic considerations in mind,⁷² it is not enough to assert its suitability on a system-wide scale. As Aymanns et al. (2016) have shown for the case of the Basel II leverage requirements, well-intended microprudential regulation may undermine financial resilience when systemic feedbacks are taken into account. This makes the investigation of the stability implications of the bail-in design in a networked financial system a necessary gap to fill.

Our system-wide stress testing methodology extends that of Farmer et al. (2020). This work in turn unites the modelling of heterogeneous institutions, contracts, markets, constraints and behaviours – which can produce amplifications among multiple contagion mechanisms. Many of these have, individually or in part, though not fully jointly,

⁷²See: Directive 2014/59/EU of the European Parliament and of the Council.

been modelled in previous papers (see e.g. [Amini et al. \(2013\)](#), [Caccioli et al. \(2013, 2014, 2015, 2014\)](#), [Kok & Montagna \(2013\)](#) for models overlapping portfolio contagion, exposure loss contagion and funding contagion). Our novel contributions – in particular to [Farmer et al. \(2020\)](#), but also to the wider financial contagion literature – are to model the bail-in design, revaluations of bail-inable debt (we take a similar approach to valuing bail-inable debt as [Pennacchi \(2010\)](#), [Chen et al. \(2013\)](#) take to price contingent convertibles), multiple maturities (which are necessary to understand the stability of bail-inable debt), and gradual risk-weighted deleveraging (rather than unbounded deleveraging as [Bookstaber \(2012\)](#), [Duarte & Eisenbach \(2015\)](#), [Greenwood et al. \(2015\)](#), [Cont & Schaanning \(2017\)](#) allow for).

In terms of enhancing understanding of the bail-in design, we are the first to have worked out the implications of the two principles for setting debt-to-equity ratios: the NCWO principle and the preservation-of-hierarchy-of-claims principle. We find that regulators always have the option to set fair conversion rates whenever they do not have to resort to the resolution financing fund; and they never have the obligation to set unfair conversion rates –which are unnecessarily detrimental to financial stability – in such cases. We work out what scope regulators have to set unfair conversion rates, while remaining compliant with these two principles. We show that debt-to-equity conversion rates may be set unfairly if the assumed liquidation costs in the hypothetical disorderly liquidation of the bank are large enough and the debt exclusions from the application of the bail-in tool are small enough. Conversely, we compute the size of the resolution financing fund contribution, which is positive whenever these assumed liquidation costs are sufficiently small and the debt exclusions are sufficiently large. We are also the first, to the best of our knowledge, to precisely stipulate the TLAC and MREL requirements for banks in formula form. We also have more precisely specified in formulas the bail-in design than we have seen in any previous paper – which also contributes to the understanding of the bail-in design.

4.9 Discussion and Policy Implications

The major contribution of this paper is the examination of the systemic implications of the bail-in design in our networked financial system. We have developed a multi-layered network model of the European financial system, which extends the one in [Farmer et al. \(2020\)](#). The model also captures the systemic footprint of the ‘primary’ and ‘secondary’ bail-in design parameters by jointly including the chief endogenous amplification mechanisms: exposure loss contagion, overlapping portfolio contagion, funding contagion, bail-

inable debt revaluations, and bail-inable debt runs.

Our results shed light onto the controversy whether bail-in is feasible in epochs of system-wide financial tumult and incidents of major SIB failures without further exacerbation of financial distress. Importantly, it enquired how the viability of bail-ins in such cases might depend, if at all, on the elected bail-in design.

We discovered that financial stability hinges on the selected bail-in policy parameters. Specifically, we showed that ‘good’ primary bail-in parameters altogether avert contagion for any idiosyncratic SIB failure, while ‘bad’ parameters induce significant contagious asset losses (on the order of multiple trillions) if larger European SIBs fail. Smaller European SIB failures are not powerful enough to bring about sweeping contagion, regardless of the bail-in design – consistent with experience (WBG (2017)). In cases of dire system-wide trouble, ‘good’ primary parameters accomplish to restrain contagious losses (extinction is usually not possible, since even in the absence of bail-ins do we observe institutions in turmoil that act in systemically unsettling ways), whereas ‘bad’ parameters substantially exacerbate distress. Supplementing the calibration of the bail-in design with ‘secondary’ bail-in parameters widens the stability wedge between a fit and an unfit primary bail-in design. ‘Good’ secondary bail-in parameters further improve resilience and ‘bad’ ones further deteriorate it.

Our results suggest that a crisis-proof bail-in design involves the following ‘good’ primary parameters: a timely bail-in, strong recapitalisations, and fair conversion rates. It is also reinforced by a trio of ‘good’ secondary policy settings: short-term (< 1 year) debt exclusions from the application of the bail-in tool, elevated loss absorption requirements, and certainty about the primary bail-in design. In contrast, a crisis-unfit bail-in design encompasses the following ‘bad’ bail-in settings: late bail-ins, weaker recapitalisations, unfair conversion rates, short-term ($7 \text{ days} < \text{maturity} < 1 \text{ year}$) debt inclusions, moderate loss absorption requirements and uncertainty in the bail-in design.

These qualitative findings are robust to extensive sensitivity analysis. Though, quantitative estimates would benefit from a more detailed network calibration, including to the cross-holdings of bail-inable debt. (Our current calibration simply deploys the balance sheet information on debt holdings in each priority class for European banks. It then reconstructs, instead of intrinsically ‘knowing’, bail-inable debt cross-holdings among banks and non-banks). A future study that amalgamates such calibration would quantitatively

strengthen our qualitatively robust results.

Our evidence fortunately suggests that the pivot for stability is in the hands of policy-makers. It also suggests, however, that the current policy parameters might be in the regime of instability. Towards this end, we briefly review the regulatory policy guidelines for each of the ‘primary’ and ‘secondary’ bail-in parameters. We then infer their systemic implications based on our findings.

Failure threshold

The key condition that must be fulfilled for the bank to be bailed-in is that it is deemed failing or likely to fail (FLTF).⁷³ Under Basel III rules, these guidelines imply that relevant authorities should intervene to resolve a bank whenever its risk-weighted capital ratio falls below the minimum requirement of 4.5% (Avgouleas & Goodhart (2015)). It is not clear that the European resolution authority will uphold this ‘good’ FLTF threshold. According to Avgouleas & Goodhart (2015), it is in the (short-term) interest of the stakeholders (bank, resolution authority, creditors, etc.) to act too late. The US authorities, according to New York Federal Reserve Staff (NY FED), will even explicitly disregard the Basel III requirement of an ‘early’ bail-in. The FED staff states that “[t]he resolution authority in the US model is ‘slow’ in the sense that it will shut down and resolve a firm only once its (book) equity capital is exhausted (McAndrews et al. (2014)).” Their preference corresponds to a ‘bad’ risk-weighted failure threshold of 0%. Given that our results unveil that regulators who bail-in banks ‘too late’ unnecessarily risk compromising resilience, the stability implications of such a ‘bad’ policy choice could prove dismal.

Recapitalisation target

A bank should be recapitalised by a sufficient amount such that it complies with the conditions for authorisation and sustains or regains market confidence.⁷⁴ This means that a recapitalisation must at least lift the bank above the minimum capital ratio of 4.5% of core tier I equity over risk-weighted assets.⁷⁵ Technical guidelines leave it debatable, however, whether recapitalisation just above the minimum requirements is enough; or, whether a bank should be further recapitalised to either a ratio in line with the average of its peers (SRB (2017)), or to a ratio that also (partially) meets its regulatory buffer standards – which sit on top of its capital requirements (Farmer et al. (2020)). Even

⁷³See: Article 33(1) of the BRRD.

⁷⁴See: Article 43(2)a of the BRRD.

⁷⁵And 8% of Tier I capital relative to risk-weighted assets.

though our results show that a strong recapitalisation strengthens stability, it is far from obvious that resolution authorities will elect this option given that a higher (‘good’) recapitalisation is affected by imposing heavier haircuts on creditors – which, notwithstanding that creditors are typically compensated by equity claims, may be politically inexpedient.

Debt-to-equity conversion rates

Article 50 of the BRRD sets out the principles for the debt-to-equity conversion ratio that regulators should use when determining how to set the conversion rate. The debt-to-equity conversion rate specifies how many shares a creditor of bank receives per unit haircut to the principle amount of claim. Article 50 tasks the European Banking Authority (EBA) with the duty to provide guidelines regarding how creditors may be appropriately compensated by means of the conversion rate. The EBA’s guideline provides two principles for setting the conversion rate in each priority class: the no-creditor-worse-off (NCWO) principle and the preservation-of-hierarchy-of-claims (PHC) principle (EBA (2017c)). Based on these two principles it is not evident that the resolution authority will elect ‘fair’ conversion rates. Instead, it is plausible that regulators deviate from the ‘fair’ rates to apply ‘unfair’ rates in order to attempt to meet NCWO principle, which we showed poses a threat to stability.⁷⁶

Debt exclusions

Certain liabilities that would be part of normal insolvency proceedings are excluded from the application of the bail-in tool either *a priori* or on an *ad-hoc* basis.⁷⁷⁷⁸ Our results have suggested that including in bail-in debt with a time to maturity between seven days and one year – as is the case in the current regulatory approach – is a ‘bad’ idea. It undercuts stability compared with the case where debt with such maturities is *a priori* excluded from bail-in. Inclusion makes the funding pile of bail-inable debt unstable. Short-term bail-inable debt is prone collapse, because creditors might avoid debt exposures to bank impending to be bailed-in. Such a debt collapse renders a successful bail-in recapitalisation infeasible and provokes disorderly asset liquidations. We have shown that

⁷⁶We recall that ‘fair’ conversion rates are such that, at the point of conversion, a creditor of bank receives an equal amount of principle amount of claim per unit haircut in the bail-in recapitalisation phase. Contingent convertibles (CoCos) with such a conversion rate are also called ‘fair’ (Chen et al. (2013)). Conversion rates are called ‘unfair’ when they deviate from this scheme. To place a bound on the impact of unfair rates, we in this paper take the maximally unfair rate that still satisfies the two principles: junior creditors are made as worse off as in liquidation and senior creditors are excessively compensated.

⁷⁷See: Article 44(1) of the BRRD.

⁷⁸See: Article 44(2) of the BRRD.

regulators would be better off excluding from bail-in debt with a time to maturity between seven days and one year, in line with the eligibility prescriptions for debt to count towards the loss absorption requirements (FSB (2015b)). The loss absorption requirements rightly acknowledge that only long-term debt with a time to maturity greater than one year is truly a stable source of loss-absorbing debt. Given the demonstrated downsides of including short-term debt, this concern is better addressed by notching up the loss absorbing requirements, as we discuss next. In absence of ‘good’ choice by regulators for debt exclusions, our result suggest that the system may be headed towards instability.

Loss absorbing requirements

The bail-in tool can only be efficacious if banks have a sufficient loss absorbing capacity to absorb losses and be recapitalised. In recognition of this fact, the Financial Stability Board (FSB) and the Bank Recovery and Resolution Directive (BRRD) have established minimum requirements for banks’ loss absorbing capacity: TLAC (for globally systemically important banks) and MREL (for European banks). Our results have provided evidence that increasing the loss absorbing requirements, relative to the *status quo*, can supremely bring down financial vulnerability as well as keeping contagion subdued following severely adverse system-wide shocks. We have also shown that an increase in the loss absorbing requirements is especially warranted if debt with a time to maturity between seven days and one year is excluded from bail-in. It is not obvious that regulators are willing to enlarge the loss absorbing requirements to serve stability. The reason is that hiking up the requirements amounts to exhorting structural changes in the maturity composition of banks’ debt: given that solely long-term debt counts towards the requirements, raising this would impel banks to swap short-term for longer-term debt. Aside from financial stability considerations, we do not consider broader economic factors that can inform whether such a change is desirable. It suffices to note that the new bundle of post-crisis regulations has already brought structural network changes. Think of the liquidity coverage ratio (LCR) that stimulates banks to hold sufficient highly-liquid assets. And thereby shifts the liquidity profile of banks’ assets (BIS (2013)). Or, think of the introduction of the loss absorbing requirements, which is already shifting the maturity profile of banks’ liabilities (FSB (2015b)). Since regulators have wrought structural changes before, they have the power to cultivate structural changes again. And hence bolster stability: banks who are obliged to hold longer-term debt are less susceptible to funding shocks and bail-inable debt collapses.

Bail-in design uncertainty

Uncertainty induces instability. The experience of the last 2007-2009 crisis has taught us once more. Disregarding this widely accepted phenomenon ([Geithner \(2014\)](#), [Bernanke \(2015\)](#)), resolution authorities have left the primary bail-in uncertain. As we have explained: the failure threshold could minimally range from a capital ratio of 0% to 4.5%; the recapitalisation target could at least range from a capital ratio 4.5% (meeting authorisation conditions) to 12% (also meeting buffer standards); and the conversion rate could attain any value that satisfies the two principles, be it fair, fully unfair or somewhere in between.

We finalise this paper by returning to Ben Bernanke's words: "Have we ended bail-outs? [...] We cannot guarantee that a future administration, fearful of the economic consequences of a building financial crisis, will not authorise a financial bail-out. But the best way to reduce the odds of that happening is to have in place a set of procedures to deal with failing financial firms that those responsible for preserving financial stability expect to be effective" ([Bernanke \(2017\)](#)). Moral hazard is reduced if bail-ins are a credible alternative to bail-out. Our paper showed that the credibility of bail-in critically depends on the bail-in design.

Chapter 5

Concluding Remarks

The overarching research question in this thesis is whether new building blocks – expressing the heterogeneity of institutions, contracts, markets, constraints and behaviour in the interconnected financial system – can be supplied for system-wide stress tests to better capture the endogenous amplification of shocks in order to improve the assessment of systemic risk and the evaluation of prudential policies to address financial fragility.

In this thesis I have proposed a novel methodology for system-wide stress testing taking into account crucial determinants of systemic risk – interacting contagion mechanisms, constraints, institution types and behaviours (in Chapter 3) – which may give rise to endogenous amplifications of valuation and liquidity shocks. This method builds forth upon the previous literature discussed in Chapter 2. This approach to system-wide stress testing proves to be powerful in answering some of the most salient questions of financial stability today – for instance, revolving around the regulatory standard design (Basel III; in Chapter 3) and novel resolution frameworks (bail-in; in Chapter 4). The application power of these methods to inform prudential policies appeals to central banks, as is evident from the collaborations I have formed with the Bank of England, European Central Bank, International Monetary Fund and South-African Reserve Bank.

While my thesis has made some important steps to advance the assessment of systemic risk and evaluate policies to address it, some shortcomings as well as exciting areas of future research remain. The limitations of my research have been elaborately discussed in the discussion section of each chapter. So I will primarily focus on how some of the deficiencies in my work may be addressed in the future research. I list some prime topics that hold my interest here:

1. Calibrating system-wide models.

- Although outside the scope of this current research, my models could benefit from a calibration to more granular data and to market liquidity.
2. System-wide stress testing the derivative markets jointly with other asset classes.
 - To holistically account for systemic risk, derivative markets should be taken into account in system-wide assessments – in my models that is currently not the case.
 3. Learning agents in financial stability models.
 - Our system-wide stress tests assume that institutions act according to heuristics, while in reality institutions are boundedly-rational learning agents.
 4. Market ecology models of systemic risk.
 - Institutions attempt to avoid default during severe distress in our system-wide model, which causes selling pressures. While in reality, when some are forced to sell others might be willing and able to buy. We do not capture the reinforcing effect this may have on stability.
 5. Feedbacks between the financial system and the real economy.
 - Ultimately the reason why we care about systemic risk in the financial system is that it may impair the real economy. Our stress tests falls short of assessing the negative effects of financial distress on the real economy.
 6. Integrating regulatory and market-based stress tests.
 - Regulatory stress tests based on accounting data may be slow to recognise risks, if any, that have evidently been building up in the financial system as evinced by market indicators. Our stress test can be improve by incorporating in marked-based stress indicators.
 7. Hedging & systemic risk.
 - Systemic risk might be misjudged if hedges that cancel risks, or that may break down in distress, are not taken into account – as is the case in our current models.

8. Reverse stress test of the South-African financial system.

- Our models investigate the impact of an adverse scenario on the financial system. Another approach to assess systemic risk is to evaluate what shock type and severity is needed to cause a system-wide collapse.

9. Climate-change stress test.

- System-wide stress test models can be naturally adopted to see how climate-risk may propagate through the system. These stress tests can in turn help to better price exposures to climate risk.

These topics will partially be studied as part of my postdoctoral research, which I will conduct at the Massachusetts Institute of Technology (MIT) and Oxford University starting in late October 2019.

Appendix A

Foundation of System-Wide Stress Testing

A.1 Further Details on Model Implementation

A.1.1 Notation

Table A.1 and A.2 give the definition of the variables used in this paper.

Table A.1: Shows the definition of notation.

Category	Subcategory (if any)	Variable	Definition
Institutions & Contracts	Cash	C_i	Cash of institution i .
		C_i^u	Unencumbered cash of institution i .
		C_{ij}^e	Encumbered cash of institution i provided to institution j .
		$C_{ij}^{e,t,E}$	Extra encumbered cash of institution i provided to institution j .
$C_{ij}^{e,t,R}$		Encumbered cash of institution i returned by institution j .	
Tradable Assets		T_i	Tradable assets of institution i .
		T_{ia}	Tradable assets of institution i of type a .
		T_{iam}	Tradable asset m of institution i of type a .
		s_{iam}	Encumbered tradable asset m of institution i provided to institution j of type a .
		s_{iam}^u	Unencumbered tradable asset m of institution i of type a .
		s_{ijam}^e	Encumbered tradable asset m of institution i provided to institution j of type a .
		$s_{ijam}^{e,t,E}$	Extra encumbered tradable asset m of type a of institution i provided to institution j .
		$s_{ijam}^{e,t,R}$	Encumbered tradable asset m of type a of institution i returned by institution j .
Repurchase Agreements		\tilde{R}_i	Reverse repo contract of institution i .
		\hat{R}_i	Repo contract of institution i .
		R_{ij}	Reverse repo contract of institution i to institution j .
		h_{am}	Haircut applicable to tradable asset m of type a .
Other Items		M_{ij}	Margin call from institution i to institution j .
		Y_i	External assets of institution i .
		D_i	Deposits of institution i .
Markets		O_i	Other assets of institution i .
		\tilde{O}_i	Other liabilities of institution i .
		p_{am}	Price of asset m of type a .
		β_{am}	Price impact parameter associated to asset m of type a .
		f_{am}^t	Cumulative fraction sold of asset m of type a up to time t .

Table A.2: Shows the definition of notation (this table is a continuation of Table A.1).

Category	Subcategory (if any)	Variable	Definition
Con- straints	Risk- weighted capital ratio	ρ_i	Risk-weighted (rw) capital ratio of bank i .
		ρ_i^M	Regulatory minimum for the risk-weighted capital ratio.
		ρ_i^B	Buffer value of the risk-weighted capital ratio where bank i acts to return to target.
		ρ_i^T	Target value of the risk-weighted capital ratio of bank i .
ρ_i^{CB}		Combined regulatory risk-weighted capital buffer of bank i .	
ρ_i^{CCB}		Capital conservation buffer of bank i .	
ρ_i^{G-SIB}		Globally systemically important bank (G-SIB) surcharge of bank i .	
ρ_i^{D-SIB}		Domestically systemically important bank (D-SIB) surcharge of bank i .	
ρ_i^{SR}		Systemic risk buffer applicable to bank i .	
ρ_i^{CCyB}		Countercyclical capital buffer applicable to bank i .	
ρ_i^{data}		Rw capital ratio of bank i by 2017Q4 <i>S&P Global Market Intelligence</i> data.	
ρ_i^{EBA}		EBA 2018 microprudential stress test outcome of bank i for its rw capital ratio.	
\tilde{E}_i		Common Tier I (CET1) equity of bank i .	
Ω_i	Risk-weighted assets of bank i .		
A_{ip}	Asset value of type p of bank i .		
ω_p	Risk weight associated to assets of type p .		
Leverage ratio	Leverage ratio	λ_i	Leverage ratio of bank i .
		λ_i^M	Regulatory minimum for the leverage ratio.
		λ_i^B	Buffer value of the leverage ratio where bank i acts to return to target.
		λ_i^T	Target value of the leverage ratio of bank i .
		λ_i^{CB}	The (combined) regulatory leverage buffer of bank i .
		λ_i^{data}	The leverage ratio of bank i according to 2017Q4 <i>S&P Global Market Intelligence</i> data.
		λ_i^{EBA}	EBA 2018 microprudential stress test outcome of bank i for its leverage ratio.
		\hat{A}_i	Asset exposure of bank i .
LCR	LCR	Λ_i	Liquidity coverage ratio (LCR) of bank i .
		Λ^S	Regulatory standard for the LCR.
		Λ_i^{data}	LCR of bank i according to the 2017Q4 <i>S&P Global Market Intelligence</i> data.
		Q_i	High-quality-liquid-assets (HQLA) of bank i .
		Θ_i	Net outflows of bank i under a 30-day period of financial distress.
		Θ_i^I	Inflows of bank i under a 30-day period of financial distress.
		Θ_i^O	Outflows of bank i under a 30-day period of financial distress.
$\tilde{\omega}_p$	Inflow rate associated to assets of type p .		
$\tilde{\omega}_l$	Outflow rate associated to assets of type l .		
NAV	NAV	η_i	Net asset value (NAV) of asset manager i .
		χ_i^t	Relative loss in NAV at time t of asset manager i in comparison with time t_0 .
Behaviour	Behaviour	$u_i^{\%}$	Usability of buffers (percentage of regulatory buffers that banks are willing to use).
		y^{ρ}	Size of combined risk-weighted buffer ρ_i^{CB} relative to Basel III standard.
		y^{λ}	Size of combined leverage buffer λ_i^{CB} relative to Basel III standard.
		Δ_i^{ρ, t_0}	Distance of pre-stress (t_0) rw capital ratio of bank i to its regulatory rw buffer.
		Δ_i^{λ, t_0}	Distance of pre-stress (t_0) leverage ratio of bank i to its regulatory leverage buffer.
		d_i	Amount bank i aims to delever.
		\hat{r}_{ip}	Amount bank i liquidates of assets of type p to raise its risk-weighted capital ratio.
		q_i	Amount bank i liquidates of non-HQLA assets to raises its LCR.
f_i^t	Fraction of the initial number of outstanding shares withdrawn up to time t .		
Systemic risk measure	Systemic risk measure	\mathbb{E}	Average extent of a systemic event (average fraction of bank defaults in a systemic event).
		\mathbb{P}	Probability of a systemic event.
		\mathbb{S}	Set of simulation runs in which a systemic event (if at least 5% of banks default) occurs.
		$f_{\mathbb{D}}(n)$	Fraction of bank defaults in case of a systemic event in simulation run n .
		N	Number of simulation runs.
Sets	Sets	\mathcal{F}	Set of financial institutions.
		\mathcal{B}	Set of banks.
		\mathcal{M}	Set of asset managers.
		\mathcal{A}	Set of different asset types (gov. bonds, corp. bonds, equities, other tradable assets).
		\mathcal{N}	Set of non-banks that do not partake in our stress test.
		\mathcal{P}	Set of different types of assets.
		\mathcal{L}	Set of different types of liabilities.
		\mathcal{D}	Set of defaulted banks.
\mathcal{J}	Set of banks that defaulted due to the adverse scenario (set of initially defaulted banks).		

A.1.2 Initialisation

Financial Institutions

Hedge Funds Due to limited available information regarding hedge funds, we make some assumptions to initialise the balance sheet of each hedge fund $i \in \mathcal{H}$ (see Section 3.5.1). Specifically, we base our initialisation on the [IOSCO \(2017\)](#) and [FCA \(2015\)](#) surveys. It is useful to summarise their main findings to support our approach to modelling hedge funds. Hedge funds can attain two types of leverage: financial leverage (i.e. that acquired through borrowing) and synthetic leverage (i.e. that obtained through derivative exposures). While the [IOSCO \(2017\)](#) and [FCA \(2015\)](#) surveys indicate that the synthetic leverage can be substantial, the financial leverage of hedge funds is typically limited. The mean financial leverage of hedge funds based on the [FCA \(2015\)](#) survey is found to be 2.3. According to [IOSCO \(2017\)](#) and [FCA \(2015\)](#), hedge funds acquire almost all their financial leverage through collateralised lending, and hardly any through unsecured funds. Collateralised lending comes either in the form of repo contracts or in the shape of margin lending. The survey finds that the split between these is about 60 to 40 percent. This funding is typically provided by the prime broker of the hedge fund, which is usually a bank. Both forms of secured lending can lead to margin calls (defined in equation 3.8), which may trigger the hedge fund to engage in fire sales.

Given the above survey information and using *ECB Statistical Warehouse Data* on the aggregate hedge fund size and its aggregate asset allocation,¹ we decided to initialise the balance sheet of each hedge fund $i \in \mathcal{H}$ as follows. We impose the (heroic) assumptions that each bank $i \in \mathcal{B}$ acts as a prime broker to one hedge fund $i \in \mathcal{H}$,² so that $|\mathcal{H}| = |\mathcal{B}|$. We set the leverage of each hedge fund $i \in \mathcal{H}$ equal to the hedge funds' mean financial leverage (i.e. $\lambda_i = 2.3, \forall i \in \mathcal{H}$). We assume all funding from a bank $i \in \mathcal{B}$ to a hedge fund $i \in \mathcal{H}$ happens via reverse repos R_i .³ As we have data on the reverse repo R_i position of each bank $i \in \mathcal{B}$ (see Section 3.5.1), the total estimated size of the hedge fund sector in Europe (from the *ECB Statistical Warehouse Data*), and the leverage λ_i of each hedge fund $i \in \mathcal{H}$ ([FCA \(2015\)](#)), we can derive the asset value A_i and repo size

¹See *ECB Statistical Warehouse*: <https://sdw.ecb.europa.eu/browse.do?node=9691340>.

²In reality, the largest hedge funds may have multiple brokers. Given the significant data limitations and the subsequent necessity to take a stylised approach, for simplicity we choose to allocate one hedge fund counterparty to each bank. In practice, where hedge funds have multiple brokers, in the case that in our simulation we had one bank withdrawing funding from a hedge fund, we should then model the appetite and capacity of that hedge fund's other prime brokers to extend their exposure to that hedge fund - or even allow the hedge fund to seek a new prime broking relationship. The exclusion of this type of behaviour means that we are likely to overstate the impact of hedge funds' defensive actions in our simulations, all else being equal.

³As explained, in reality funding to hedge funds also goes via margin lending. We do not model margin lending for two reasons. First, we do not have data on the size of margin lending banks engage in. Two, margin lending does not affect systemic risk materially differently than repo lending does: in both cases, margin calls may trigger hedge funds to engage in fire sales. For a detailed discussion of margin lending, hedge funds and stability, see [Paulin et al. \(2018\)](#).

\tilde{R}_i ⁴ of each hedge fund $i \in \mathcal{H}$. The value of each asset type A_{ip} of a hedge fund $i \in \mathcal{H}$ is approximated, by multiplying the asset size A_i of a hedge fund $i \in \mathcal{H}$, with the aggregate asset value held by hedge funds of that asset type relative to the aggregate asset value of hedge funds.

Further Details on Financial Contracts

Tradable Assets As discussed in Section 3.5.2, we consider different types of tradable assets $a \in \mathcal{A}$. Specifically, we consider four types: $\mathcal{A} := \{\text{government bonds, corporate bonds, equities, other tradable assets}\}$. For each financial institution $i \in \mathcal{B} \cup \mathcal{M} \cup \mathcal{H}$, our balance sheet data (see Appendix A.1.2) allows us to initialise the value of each tradable asset type T_{ia} , for $a \in \mathcal{A}$. We set the number of individual securities M^a per type $a \in \mathcal{A}$ in line with the number of securities that Cont & Schaanning (2017) construct per type $a \in \mathcal{A}$ in the EU network. Specifically, this means setting $M^a = 37, \forall a \in \mathcal{A}$, which corresponds to the 37 geographical regions that Cont & Schaanning (2017) consider for their four types of marketable securities.

The *ECB Statistical Warehouse* also gives an estimate of the aggregate EU tradable asset positions for each non-bank sector T_{ia} ($a \in \mathcal{A}$ and $i \in \mathcal{N}$, where \mathcal{N} denotes the set of different types of non-banks not considered in our stress test) not included in our system-wide stress test (e.g. pension funds, insurance companies, financial vehicle corporations). Together this allows us to reconstruct the common asset holdings network (i.e. $T_{iam}, \forall i \in \mathcal{B} \cup \mathcal{M} \cup \mathcal{H} \cup \mathcal{N}, \forall a \in \mathcal{A}$, for $m = 1, \dots, M^a$) using the reconstruction method employed in Kok & Montagna (2013). In essence, this method allows us to reconstruct $|\mathcal{A}|$ number of random bipartite networks between the $|\mathcal{B} \cup \mathcal{M} \cup \mathcal{H} \cup \mathcal{N}|$ nodes and M^a securities, for each $a \in \mathcal{A}$. In each network, a link from an institution i to a security a means that the institution has that particular security in its portfolio. The amount of the shares is represented through the weight of the edge. Each link in a bipartite network has the same probability p to exist.⁵ In line with Kok & Montagna (2013), we assume that, for each institution, all its out-going links have the same weight.

Markets To estimate the price impact (see Section 3.5.2), we set the denominator of the cumulative fraction of net asset sales f_{am}^t , which appears in the price impact function (see equation 3.7) to the total market capitalisation in asset m of type $a \in \mathcal{A}$, which includes the holdings of non-banks that are not included in our stress test (see Appendix A.1.2).

⁴Namely, the repo size \tilde{R}_i of each hedge fund $i \in \mathcal{H}$ equals the reverse repo size R_i of its corresponding prime-broker bank $i \in \mathcal{B}$.

⁵We set $p = 0.3$ as a baseline. We find that our qualitative results are robust to varying p . We note that varying p influences the sparseness of the network and concentration of asset holdings.

That is, the denominator of f_{am}^t is given by $\sum_{i \in \mathcal{B} \cup \mathcal{M} \cup \mathcal{H} \cup \mathcal{N}} \frac{T_{iam}^{t_0}}{p_{am}^{t_0}}$, where in line with the contagion literature (e.g. Caccioli et al. (2014)) the initial price of each tradable asset is normalised to $p_{am}^{t_0} = 1$.

Constraints Here we discuss the regulatory parameters that are associated to the Basel III regulatory capital requirements and buffer standards discussed in Section 3.5.1. Let us start with explaining how the risk-weights ω_p in the risk-weighted capital requirement ρ_i (see equation 3.1) are set. In line with the Basel III standardised approach, we set the risk weights ω_p for $p = 1, \dots, 8$ (i.e. except $p = 9$) equal to $\{0, 0.35, 0, 1, 0.75, 1, 0.4, 0.1\}$. We are able to compute the risk-weight ω_{p9} for other assets O_i by solving one equation is one unknown as $\omega_{p9} = (\frac{\tilde{E}_i}{\rho_i} - \sum_{p=1}^8 \omega_p A_{ip}) \frac{1}{A_{i9}}$. Once we have computed ω_{p9} , we keep it constant throughout the stress test. Setting the risk-weight ω_{p9} as such makes sure that the CET1 ratio ρ_i at time t_0 of the stress test aligns with the 2017Q4 data. It makes sense to not set a fixed risk-weight for ω_{p9} , as other assets O_i is a collection of a variety of assets that would bear different risk-weights under the Basel III standardised approach.

We will now discuss the parameters associated to the LCR Λ_i (see equation 3.5). The net outflows Θ_i in the LCR denominator were defined as a function of the inflows $\Theta_i^I := \sum_{p \in \mathcal{P}} \tilde{\omega}_p A_{ip}$ and outflows $\Theta_i^O := \sum_{l \in \mathcal{L}} \tilde{\omega}_l L_{il}$ under distress. Here $\tilde{\omega}_p$ is the inflow rate for asset type $p \in \mathcal{P}$ and $\tilde{\omega}_l$ is the outflow rate for liability type $l \in \mathcal{L}$. The inflow rates $\tilde{\omega}_p$ and outflow rates $\tilde{\omega}_l$ associated to the LCR Λ_i are set in line with BIS (2013). The outflow rates $\tilde{\omega}_l$ associated with $\{D_i, \tilde{I}_i, \tilde{R}_i, \tilde{O}_i\}$ are respectively set to $\{0.05, 1, 1, 0.5\}$.⁶ Other liabilities \tilde{O}_i is a mix of different liabilities, so we cannot precisely determine the outflow rate. Hence, we set it equal to the (approximate) average outflow rate: 0.5. The inflow rates $\tilde{\omega}_p$ associated with $\{C_i, Y_i, T_i, I_i, R_i, E_i\}$ are respectively set to $\{0, 0.5, 0, 1, 1, 0\}$.⁷ The inflow rate $\tilde{\omega}_p$ associated with other assets O_i cannot be precisely determined as other assets consists of a mix of different asset types. Hence, we set it such that the LCR at time t_0 , $\Lambda_i^{t_0}$, matches the 2017Q4 data for each bank $i \in \mathcal{B}$. We keep the outflow rate $\tilde{\omega}_p$ associated with other assets O_i constant throughout the stress test. Whenever the LCR Λ_i of a bank is not reported we set it equal to the average LCR of the other banks $i \in \mathcal{B}$ for which the LCR Λ_i was reported.

The bank-specific standards for the components of the risk-weighted capital buffer ρ_i^{CB} (i.e. the G-SIB surcharge ρ_i^{G-SIB} , the D-SIB surcharge ρ_i^{D-SIB} , the systemic risk buffer ρ_i^{SR} , and the CCyB ρ_i^{CCyB} , see equation 3.3) are publicly listed.⁸

⁶If repo contracts \tilde{R}_i are secured with HQLA assets the outflow rate is zero instead of one.

⁷Again, if reverse repo contracts R_i are secured with HQLA assets the inflow rate is zero instead of one.

⁸See the list of G-SIB surcharges here: <http://www.fsb.org/wp-content/uploads/P211117-1.pdf>. See the list of D-SIB surcharges here: <https://www.eba.europa.eu/risk-analysis-and-data/>

Behaviour No data available, as discussed in Section 3.4.1.

A.1.3 Default Configuration

y-axis: Systemic Risk Measure In line with, but a generalisation upon Schnabel & Shin (2004), Cifuentes et al. (2005), Gai & Kapadia (2010), Caccioli et al. (2014), Paulin et al. (2018), we use the ‘average extent of a systemic event \mathbb{E} ’ to measure systemic risk. The systemic risk measure \mathbb{E} gives the average fraction of (contagion) defaults *given that a systemic event occurs*, which is said to be so if at least $\gamma = 5\%$ (contagion) defaults occur. That is, \mathbb{E} is given by

$$\mathbb{E} := \frac{1}{|\mathcal{S}|} \sum_{n \in \mathcal{S}} f_{\mathcal{D}}(n), \quad (\text{A.1})$$

where the terms of equation A.1 are defined as follows:

- \mathcal{S} denotes the set of simulations runs (out of N simulation runs in total) in which a systemic event occurs. That is, \mathcal{S} is defined by

$$\mathcal{S} := \{n \in [1, N] : \mathbb{1}_{SE}(n) = 1\}, \quad (\text{A.2})$$

where $\mathbb{1}_{SE}(n) = 1$ is an indicator variable denoting the occurrence of a systemic event in simulations run n , and is given by

$$\mathbb{1}_{SE}(n) = \begin{cases} 1, & \text{if } f_{\mathcal{D}}(n) > \gamma, \\ 0, & \text{otherwise,} \end{cases}$$

where γ denotes the threshold above which a systemic event is said to occur.

- $f_{\mathcal{D}}(n)$ denotes the fraction of (contagion) defaults in run n , defined as

$$f_{\mathcal{D}}(n) = \frac{1}{|\mathcal{D}|} \sum_{i \in \mathcal{D}} \mathbb{1}_{\mathcal{D}}(i, n), \quad (\text{A.3})$$

where we set $\mathcal{D} = \mathcal{B}$ (\mathcal{B} is the set of banks) or $\mathcal{D} = \mathcal{B} \setminus \mathcal{J}$ (\mathcal{J} is the set of initial defaults, so $\mathcal{B} \setminus \mathcal{J}$ is the set of banks that could default due to contagion). Hence, $f_{\mathcal{B}}(n)$ gives the fraction of total (i.e. initial defaults plus contagion defaults) in run n and $f_{\mathcal{B} \setminus \mathcal{J}}(n)$ gives the fraction of contagion defaults in run n . $\mathbb{1}_{\mathcal{D}}(i, n)$ is an indicator variable indicating whether a bank defaults (due to contagion) in run n . That is,

$$\mathbb{1}_{\mathcal{D}}(i, n) = \begin{cases} 1, & \text{if institution } i \in \mathcal{D} (= \mathcal{B}, \mathcal{B} \setminus \mathcal{J}) \text{ defaults in run } n, \\ 0, & \text{otherwise.} \end{cases}$$

other-systemically-important-institutions-o-siis-/2017. See the list of of applicable systemic risk buffers here: https://www.esrb.europa.eu/national_policy/systemic/html/index.en.html. See the list of CCyB here: https://www.esrb.europa.eu/national_policy/ccb/html/index.en.html.

For completeness, the probability of a systemic event \mathbb{P} is given by

$$\mathbb{P} := \frac{1}{N} \sum_{n=1}^N \mathbb{1}_{SE}(n), \quad (\text{A.4})$$

although we do not use this measure in our results. The randomness arises from the random redraw in every simulation run $n = 1, \dots, N$ of the interbank- and common asset holdings network (see Appendix A.1.2).

To interpret our results (see Section 4.7) correctly it is important to note the following. The coloured lines, which correspond to the system-wide stress test outcomes, give the average extent of a systemic event \mathbb{E} (see equation A.1 applying $\mathcal{D} = \mathcal{B}$). That is, it shows the average fraction of total (i.e. initial + contagion) defaults in a systemic event. The grey lines (associated to the coloured lines), which correspond to the microprudential stress test outcome, also display systemic risk measure \mathbb{E} for the case where $\mathcal{D} = \mathcal{B}$. However, since by design a microprudential stress test only captures initial defaults and no contagion defaults, the systemic risk measure \mathbb{E} in fact displays the average fraction of initial defaults in a systemic event (which is not random as it does not depend on the redrawing of the network). The difference between the coloured and the grey lines (i.e. between the system-wide and microprudential stress test outcome) typically corresponds to the average extent of a systemic event \mathbb{E} when $\mathcal{D} = \mathcal{B} \setminus \mathcal{J}$. When $\mathcal{D} = \mathcal{B} \setminus \mathcal{J}$, the average extent of a systemic event could also be called ‘the average extent of contagion (in a cascade)’, as Schnabel & Shin (2004), Cifuentes et al. (2005), Gai & Kapadia (2010), Caccioli et al. (2014), Paulin et al. (2018) refer to the systemic risk measure \mathbb{E} .

A.2 System-Wide Stress Testing Software

We developed state-of-the-art system-wide stress testing software, which lives up to today’s standards in (scientific) computing. This software can be used by regulators (and researchers) to build their own system-wide stress test models and flexibly adjust these depending on the stress test exercise or policy question at hand. Good software is critical to run robust stress tests on big data. In this Appendix will provide the links to the software packages and motivate their design principles. Detailed documentation is found on the Github links provided. Furthermore, we will discuss how we ensure that the institutions act in synchronous ways when this would be the case in financial markets.

A.2.1 Design Principles

We will now discuss the five design principles for robust system-wide stress testing code listed in Section 3.4.3.

Transparency The design principle transparency says that the model’s specification has to be fully documented. This is done by publishing a complete description of the model and by making the library (if any) underpinning the code and the code of the model (built within the library) publicly available. Additionally, we are putting emphasis on modularity and readability to further improve on transparency. Our model is fully described in this paper and our code (with a detailed code documentation) is published under the Apache License.⁹ The link to the system-wide stress test library and model are found here:

- **System-Wide Stress Test Library:**
<https://github.com/ox-inet-resilience/resilience>
- **System-Wide Stress Test Model:**
github.com/ox-inet-resilience/sw_stresstest

The library repository consists of reusable and extensible building blocks. The model repository is built in this library and consists of the system-wide stress testing model on randomised data (as not all data used for the paper is publicly available). The system-wide stress test library itself is built on top of the Economic Simulation Library (ESL), which contains a system to make the simulation order independent (see Appendix ??).¹⁰ Further, to give a broad overview of the structure of the code Figure A.1 displays the class diagram of the code.




Figure A.1: A system-wide stress test consists of five building blocks and so does its code. We have three main classes: First, *institutions*. The different institution types (e.g. bank, asset managers, and hedge funds) inherit from the parent class institutions. Since (regulatory) *constraints* are typically institution-type specific these inherit from the institution-specific classes. Second, *contracts* and its associated class *contract* 'obligations'. Each type of contract has its own class (e.g. tradable asset and loan (repo loan and interbank loan)) and inherits from the generic contract class. Third, *markets* and its associated class 'order book'. There can be many types of markets among which an asset market, which inherits from the generic market class. In addition the code also has a separate section (i.e. file) dedicated to the building block *behaviour*. Since behaviour only consists of behavioural functions it does not have its own class.

⁹Alternatively, transparency of the code can be achieved by publishing a virtual machine containing the code and environment. Such a virtual machine could also contain a detailed description of the model (see e.g. ?).

¹⁰The link to the Economic Simulation Library is given by: <https://github.com/ox-inet-resilience/py-distilledESL>.

Reproducibility The design principle reproducibility says that the reader should be able to run the simulations and obtain an identical outcome. Although it is often not possible to reproduce plots in a paper, since they use confidential data, it is possible to reproduce the same plots as in the paper, but run on random data. This can be done as follows. First, the same version of the code has to be used (while our code continues to develop over time, Github has version control that allows you to obtain the version that was used to produce the plots in the paper). The version of the code used to produce the the plots in this paper is tag ‘v0.2’ (find in Github of System-Wide Stress Test Library). The tag is the shorthand for the hash. Second, the same fixed seed has to be used as for the plots in the paper. Third, the reproduced plots have to be compared against the random data plots made available on Github and must be found to be identical.

Modularity The design principle modularity says that the code must be composed of self-contained modules. Modularity implies *flexibility*, which is the ability to easily adjust the model, replace components of the model with others, and extend (or reduce) the model. The five building blocks for system-wide stress testing (see Section 3.4.1) are chosen in such a way to maximise the modularity of the code. This is beneficial for various reasons, including the following:

1. It allows one to turn institutions, constraints, contagion mechanisms and behavioural strategies on and off, so that (among others):
 - (a) The financial system’s dynamics can be studied both holistically and in part.
 - (b) The contributions of each components to stability can be detected.
 - (c) (Simpler) contagion models can be replicated.
 - (d) The validity of the model can be checked and enforced one component at a time.
2. It allows one to model (some parts of) the system in a more abstract and (some parts of the system in) a more detailed way, depending on the granularity of data available, the assumptions being made, or the research or policy question being asked. For instance, a price impact function could be replaced by an order book where it would be advantageous for a particular research question.
3. It facilitates the adjustment of the stress test to a changing financial system. This is indispensable for the tool to have longevity in the macroprudential policy toolkit, since the structure of the financial system and the amplification mechanisms that it comprises constantly change over time ([Anderson et al. \(2018\)](#)).

The behaviours of institutions are deliberately separated into their own building block (see Section 3.4.1) because these vary most across models and because behavioural strategies are typically assumptions (by lack of data), whose sensitivity to the stress test outcome should be studied.

To illustrate how the five building blocks contribute to modularity and make it easy to implement many other contagion and stress test models in the literature, we implemented an overlapping portfolio contagion model (also referred to as fire sale contagion model) similar to Cont & Schaanning (2017), using the organising principles of the system-wide stress test framework, see:

- **Fire Sale System-Wide Stress Test (Learning Module):**
https://github.com/ox-inet-resilience/firesale_stresstest

We highly recommend the reader to go through this simplified model in order to grasp the structure of the framework. The full model essentially uses the same class structures and common application programming interface (API) but only with more extensive implementations for each building blocks.

Readability The design principle readability says that the reader should be able to read and understand the implementation in a short amount of time. We prioritise the code to be readable over inherent performance. We do so by using Python over other compiled languages in order to avoid verbosity in expression of the framework. (Another reason python is our preferred language since it has extensive scientific libraries ecosystem and is most widely used (Economist (2018)).) Further, we make our code readable by choosing intuitive variable names, commenting the code where necessary, and structuring our code logically (see e.g. Appendix A.2.1 on modularity).

Performant The design principle performance says that the code should execute fast as possible as long as it does not sacrifice readability (see Appendix A.2.1). The prime way to address performance is to use parallelisation across multiple central processing units (CPUs), which in colloquial language means that that computations are distributed across multiple brains (computational units). Parallelisation across N number of CPUs has the benefit of reducing the computation time by about N times. Cloud computing services enable you to run stress testing code at multiple CPUs.¹¹

¹¹Amazon EC2 c5.2xlarge and Google Cloud Platform n1-highcpu-16, are cloud computing services. Amazon EC2 c5.2xlarge consists of 16 vCPUs of 3.0 GHz Intel Xeon Platinum 8000 series, boostable to 3.4 GHz. See: <https://aws.amazon.com/ec2/instance-types/c5/>. Google Cloud Platform n1-highcpu-16 consists of 16 vCPUs of available Intel Xeon platforms. See: <https://cloud.google.com/compute/docs/machine-types#highcpu>, and <https://cloud.google.com/compute/docs/cpu-platforms>

In our model a figure consists of multiple lines, which extend the x-axis based on x computed points (where currently $x = 11$), where each line is computed as the average over N independent simulation runs (where currently $N = 100$). In such case, there are two ways to parallelise the computations to produce the figure:

1. Parallelise across independent simulation runs $n = 1, \dots, N$, where each simulation run has different random seed.
2. Parallelise across institutions within a simulation run n .

We chose the first for two reasons. First, a simulation run typically completes in approximately 3s. Hence, it is costlier to spawn processes dedicated to each institutions within 3s. Second, in order to parallelise across institutions, the code has to be designed in such a way that enables live objects to be serialised into a file, which often complicates the implementation.¹² While the second parallelisation technique would have provided the same amount of speed up if the code were to be written in C++, the crux of the point is that the first technique speeds up our current runs to the point where speed is no longer an issue. It takes typically 5-10 min to produce a figure which is run on Amazon EC2 c5.4xlarge (the figure would have taken about 8 times longer on a single core), for instance.

Other ways in which the performance of the code can be enhanced include caching commonly repeated computations and commonly called variables. For example, once an institution's total assets A_i have been computed ones (which is a relatively expensive operation according to the profiling results) and are known to be invariant over the next steps of the computations, its value is passed over directly to the next function

As part of our future development of the system-wide stress test library, we plan to maintain two versions of the library, a front-end library and back-end library, which will display identical behaviour. The Python implementation will focus on readability. The C/C++/Cython/Julia implementation will focus on performance. Two-language software is commonly observed in scientific computing. For instance, the linear algebra subset of Numpy library has various performant back-end choices, such as LAPACK, ATLAS, BLAS and OpenBLAS.¹³ Also, many machine learning libraries consist of multiple languages. The Keras machine learning library, is one such example. It has back-end choices which include Tensorflow, Theano and CNTK.¹⁴ With the scale of the current model, it is not a priority to implement the performant version yet. However, for modelling entire derivatives markets on a real-time basis, for instance, such speed-ups become essential.

¹²See abcEconomics for how to do this in Python: <https://github.com/ab-ce/abce>.

¹³See: <https://docs.scipy.org/doc/numpy/user/building.html#prerequisites>.

¹⁴See: <https://keras.io/backend/>.

Correctness We add the design principle, correct, to emphasise the importance of creating bug-free code that does what it should do. There are two main ways in which the correctness of the code can be asserted. First, you can make assertions, which is a statement that a predicate (i.e. a Boolean function which either outputs true or false) is always true at that point in code execution. If we encounter a bug, we usually add a new in-line assertion, to actively prevent future bugs of the same kind. For instance, we made an assertion to ensure that each bank raised enough liquidity, wherever the bank had sufficient assets that could be liquidated, to be able to reach its risk-weighted capital target ρ_i^T .¹⁵ Second, unit tests should be implemented in the code to complement assertions. The purpose is to validate that each unit of the software performs as designed. A unit is the smallest testable part of any software (it usually has one or a few inputs and usually a single output). We plan to do more work to add unit tests going forward (open source contributions are welcome).

An evident way to make the code correct is to take out bugs. Rather than relying on detailed logging in order to inform us the internal state of the system at any given time (which we found out grew to an enormous size, especially during a sensitivity analysis), we use the line number information in the error message to immediately point us in the right direction to start debugging. In absence of logging messages the code becomes more concise, so that the reader can better grasp the logical flow of the code (see Appendix [A.2.1](#) on readability).

¹⁵It should be noted that while in principle enough liquidity has been raised, due to price impact, non-repayments, and other factors, the actual liquidity raised later on may differ from the calculated liquidity raised.

A.3 Sensitivity Analysis

From Micro to Macro: A Macroprudential Overlay to the EBA 2018 Stress Test

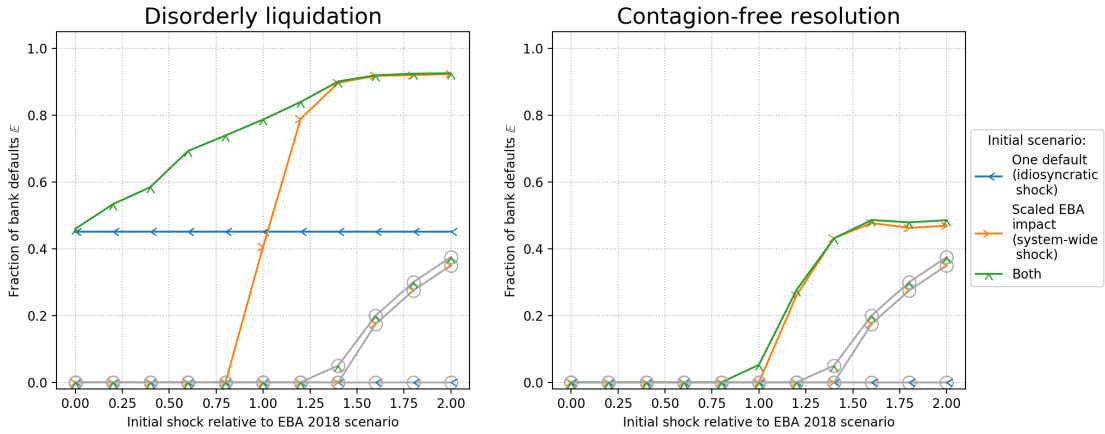


Figure A.2: Shows the same set-up as in Figure 3.3, except that we have now turned the leverage ratio λ_i off. This means that the leverage minimum λ^M , buffer λ_i^B and target λ_i^T do not apply. We observe that the financial system remains stable for a much more severe initial scenario when the leverage ratio λ_i is off compared to when it is one (in Figure 3.3). Joint with the results of Figure A.3, this indicates that instability under Basel III might be driven by a binding leverage constraint.

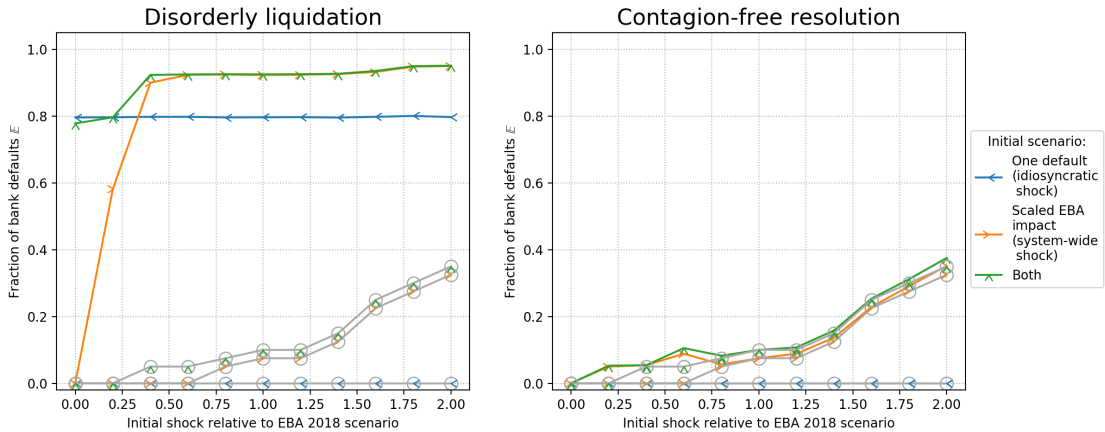


Figure A.3: Shows the same set-up as in Figure 3.3, except that we have now turned the risk-weighted capital ratio ρ_i off. This means that the risk-weighted capital ratio minimum ρ^M , buffer ρ_i^B and target ρ_i^T do not apply. We observe that this result is almost identical to the result when the risk-weighted capital ratio ρ_i is turned on in Figure 3.3. Together with the results in Figure A.2, this suggests that the risk-weighted capital ratio is relatively less binding than the leverage ratio, and is not driving the instability dynamics under Basel III.

Table A.3: Summary statistics of the leverage ratio and risk-weighted capital ratio of banks. It shows the average value (and standard deviation) of the ratios: (a) pre-distress; (b) initial-distress (value after the application of the 2018 EBA impact); (c) buffer point at which banks act; (d) combined regulatory buffer; (e) and minimum capital ratio. Furthermore, the table also shows that on average the banks' leverage ratio binds more than their risk-weighted capital ratio, both prior to and after the initial distress. The “distance-to-act” and distance-to-default are measures that express the degree to which the constraints of banks bind. These two measures are shown in this table too.

		Leverage ratio		Risk-weighted capital ratio	
Average and standard deviation raw data	Pre-distress ratio	$\bar{\lambda}^{data}$	5.5% (1.6%)	$\bar{\rho}^{data}$	15.3% (3.3%)
	Initial distress ratio (under 2018 EBA scenario)	$\bar{\lambda}^{EBA}$	4.7% (1.6%)	$\bar{\rho}^{EBA}$	11.3% (3.5%)
	Buffer ratio (point where to act to avoid getting too close to default, default setting: 50% usability of regulatory buffers)	$\bar{\lambda}^B = \lambda^M + 0.5\lambda^{CB}$	3.3% (0.05%)	$\bar{\rho}^B = \rho^M + 0.5\rho^{CB}$	6.5% (0.6%)
	Combined regulatory buffer (CB)	λ^{CB}	0.6% (\$0.1%)	ρ^{CB}	3.9% (1.1%)
	Minimum ratio	λ^M	3%	ρ^B	4.5%
Average “distance-to-act”	Prior to distress	$\lambda^{data} - \lambda^B$	2.2%	$\bar{\rho}^{data} - \bar{\rho}^B$	8.9%
	Initial distress (under 2018 EBA scenario)	$\bar{\lambda}^{EBA} - \bar{\lambda}^B$	1.4%	$\bar{\rho}^{EBA} - \bar{\rho}^B$	4.9%
Average distance-to-default	Prior to distress	$\lambda^{data} - \lambda^M$	2.5%	$\bar{\rho}^{data} - \rho^M$	10.8%
	Initial distress (under 2018 EBA scenario)	$\bar{\lambda}^{EBA} - \lambda^M$	1.7%	$\bar{\rho}^{EBA} - \rho^M$	6.8%

Amplification of Contagion Mechanisms







Figure A.4: Shows the excess systemic risk \mathbb{E} (given by the joint set of contagion mechanisms *minus* the systemic risk \mathbb{E} of the sum of the individual contagion mechanisms) for various price impacts (PI), for various sets of contagion mechanism and the same set-up as in Figure 3.6. A positive excess systemic risk \mathbb{E} means that the considered contagion mechanisms are mutually amplifying; the value gives the absolute degree of underestimation of systemic risk if the contagion mechanisms are not jointly considered. (Negative excess systemic risk is an artefact of a finitely-sized financial system, which prevents systemic risk \mathbb{E} produced by the joint set of contagion mechanisms to exceed that of the sum of the parts when the individual contagion channels already produce near maximum instability.) We observe that in absolute terms, systemic risk could be underestimated by over $\mathbb{E} \approx 55\%$ (see the *O&E&F&C* bar). Importantly, we note that the contagion mechanisms that amplify each other most in relative terms (see Figure 3.6b) may not be the same contagion mechanisms that amplify each other most in absolute terms (see Figure A.4). For instance, overlapping portfolio contagion and collateral contagion (see *O&C* at 5% price impact in 3.6b) amplify each other most in relative terms, while overlapping portfolio contagion and exposure loss contagion (see *O&E* at 5% price impact in A.4) amplify each other most in absolute terms.



(a)



(b)



(c)

Figure A.5: Shows the same set-up as in Figure 3.6, except that now we use the Basel III default settings (see Table 3.2). We observe that under Basel III, overlapping portfolio contagion ('O') and exposure loss contagion ('E'), on an individual basis, already cause the system to be unstable, so combining multiple contagion mechanisms cannot do much more harm in a finitely-sized system. As such the amplification in Plot A.5b is often smaller than one, and the excess systemic risk \mathbb{E} in plot A.5c is frequently negative.

‘Usability’ of Buffers and Contagion

Disorderly liquidation

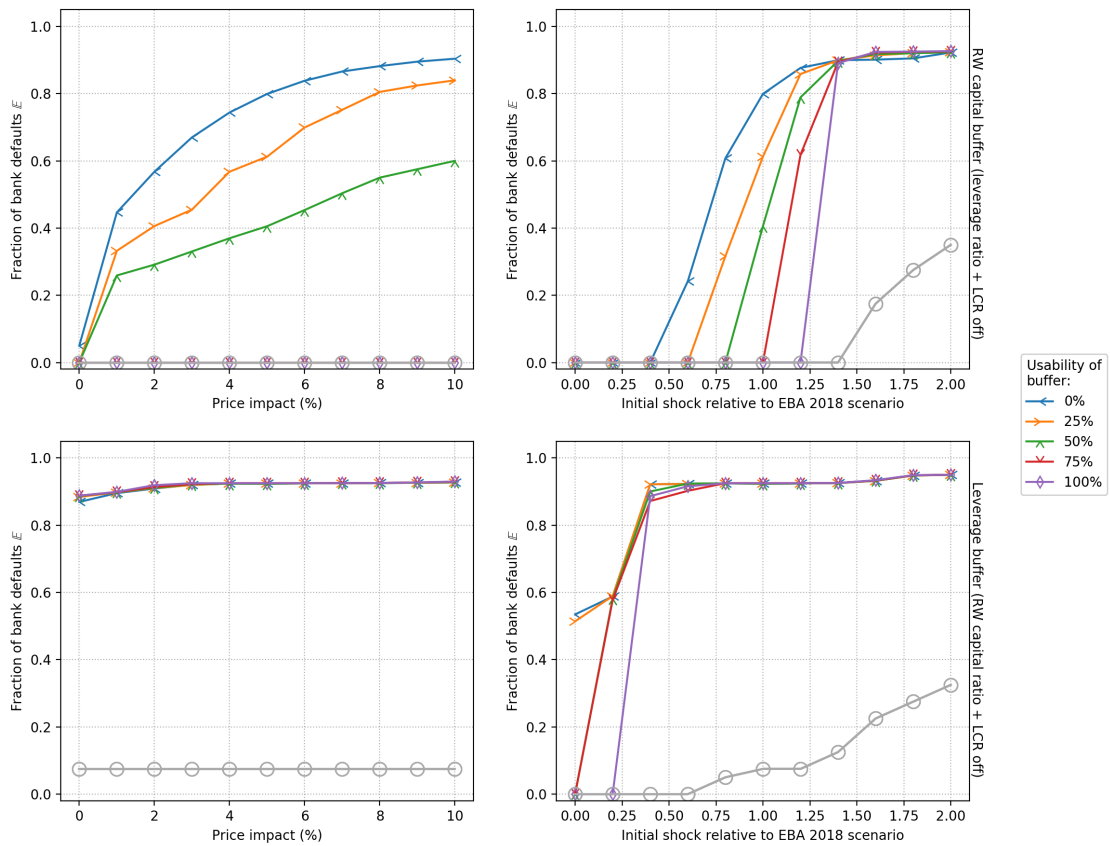


Figure A.6: Shows the same set-up as in Figure 3.7, except now for the case of ‘disorderly liquidation’ the top row shows the effect of the usability of the regulatory risk-weighted capital buffer ρ_i^{CB} only (i.e. the leverage ratio λ_i and the LCR Λ_i are turned off), and the bottom row shows the effect of the usability of the regulatory leverage buffer λ_i^{CB} only. Adding to the findings of Figure 3.7, we observe that resilience also increases in the usability of each individual regulatory capital buffer, and also holds for the case of ‘disorderly liquidation’.




Figure A.7: Shows the same set-up as in Figure 3.7, except that we now fix the usability of the regulatory leverage buffer λ_i^{CB} and the risk-weighted capital buffer ρ_i^{CB} at their default value of $u = 50\%$, and solely vary the usability of the LCR standard Λ^S . We observe that the usability of the LCR does not (or barely) affect systemic risk \mathbb{E} . This indicates that the LCR does not bind under the stress conditions imposed by the (scaled) 2018 EBA scenario.




Figure A.8: Shows systemic risk \mathbb{E} as a function of the price impact for the case where the regulatory leverage buffer λ_i^{CB} is tripled (i.e. $y^\lambda = 3$), for different excess targets above the buffer (i.e. $\rho_i^B - \lambda_i^B = \lambda_i^T - \lambda_i^B = x\%$, for $x = 0.5, \dots, 5\%$, see definitions in Section 3.5.3). The excess target above the buffer basically tells by how many percentages the bank improves its capital ratio to return to its target once it has breached its buffer value. We observe that stability decreases if banks more aggressively move away from their buffer values (for different buffer settings we also observe this effect for the right plot, but now the instability is too small to be affected by the excess target). Hence, individual stability can lead to collective instability.

Calibration of Buffers with System-Wide Stress Tests

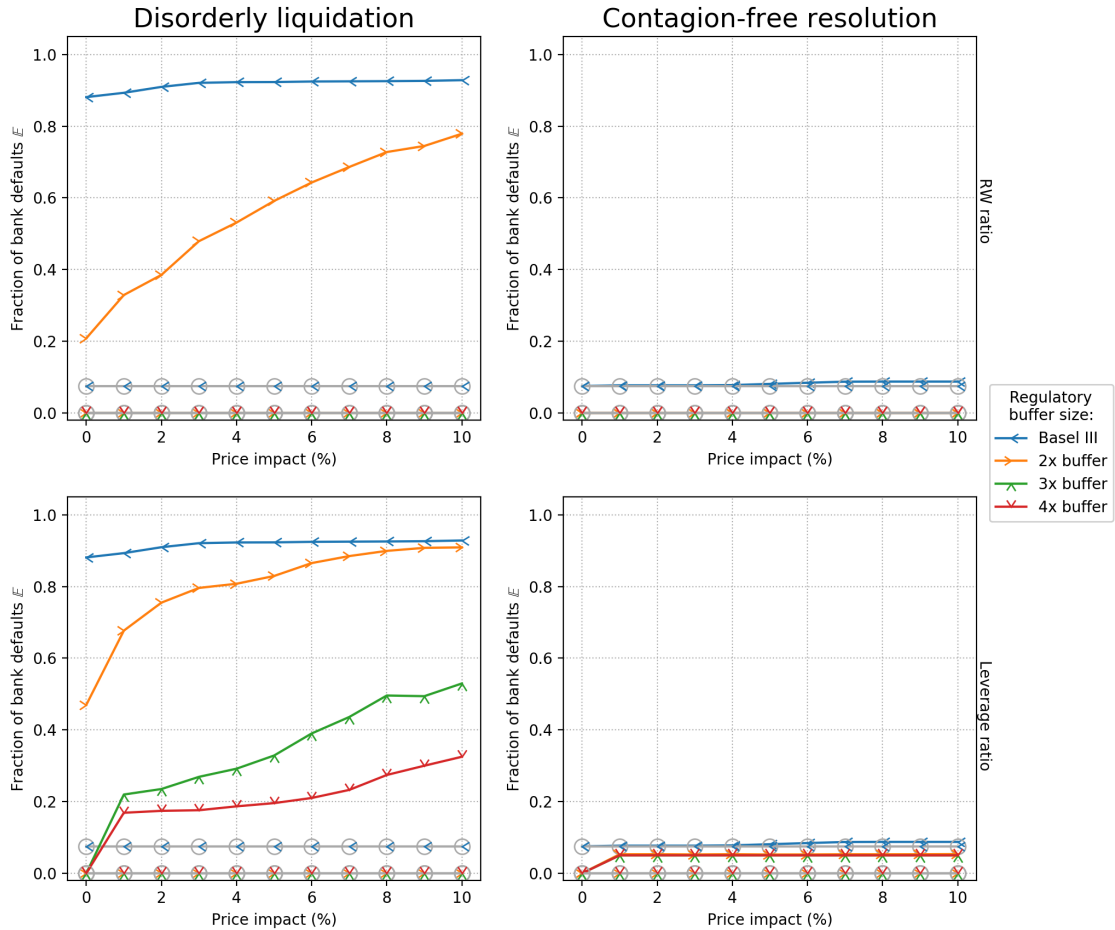


Figure A.9: Shows the same set-up as in Figure 3.8 except we now show how systemic risk \mathbb{E} decreases in the regulatory buffer sizes as a function of the price impact. Also, we add the triple buffer case. By showing systemic risk \mathbb{E} as a function of the price impact, it becomes even easier to observe that the size of regulatory capital buffers needed to confine systemic risk may be underestimated if system-wide dynamics are not taken into account (the grey-coloured lines give the necessary buffer sizes according to the microprudential stress test and the coloured lines tell the requisite the buffer sizes when system-wide effects are taken into account). Specifically, imagine that regulators believe that the initial shock size will not exceed the 2018 EBA shock ($x \in [0, 1]$), and that they wish to bound systemic risk underneath $\mathbb{E} = 10\%$ for a price impact in interval $[0\%, 10\%]$, in a regime where banks are ‘disorderly liquidated’. In such case, the microprudential stress test would find that the Basel III buffers are sufficient (the grey-blue ‘Basel III’ line at $\mathbb{E} \approx 5\%$ in the top-left panel of Figure A.9). However, when system-wide dynamics are taken into account, regulators would find that they need to more than double the risk-weighted capital buffers to achieve this (the green ‘3x buffer’ line at $\mathbb{E} = 0\%$ in the top-left panel of Figure A.9 is the first line to fall underneath $\mathbb{E} = 10\%$).

Appendix B

Systemic Implications of the Bail-In Design

B.1 Data

B.1.1 Initialisation of Balance Sheets

In Table B.1 we report how we initialised the balance sheets of banks using SNL 2017Q4 data.

Table B.1: Initialisation of Bank Balance Sheets

Balance Sheet Items	Data Identifier/Inference	Associated Data Name
Total Assets A_i	132264	Total Assets (Reported)
• Cash C_i	246025	Cash and Balances with Central Banks (Reported)
• External Assets Y_i	132214 - 224934	Total Net Loans (Reported) Net Loans to Banks (Reported)
• Interbank Assets I_i	224934	Net Loans to Banks (Reported)
• Reverse Repo R_i	224945 + 224944	Memo: Reverse Repos Incl in Customer Loans (Reported)+ Memo: Reverse Repos Incl in Bank Loans (Reported)
• Tradable Assets T_i	132191	Total Securities (Reported)
- Government Bonds T_i^1	224927 (SNL) * gov bonds EBA/ (gov + corp bonds EBA)	Total Debt Instruments (Reported)
- Corporate Bonds T_i^2	224927 (SNL) * corp bonds EBA/ (gov + corp bonds EBA)	Total Debt Instruments (Reported)
- Equities T_i^3	224928	Total Equity Instruments (Reported)
- Other Tradable Assets T_i^4	224929 + 224930 , or $\max(T_i - \sum_{k=1}^3 T_i^k, 0)$	Securities Owned: Derivative Financial Instruments (Reported) + Securities Owned: Other Investments (Reported)
• Other Assets O_i	$A_i - (C_i + Y_i + I_i + R_i + T_i)$	
Total Liabilities L_i	132367	Total Liabilities (Reported)
• Deposits D_i	132288 - 224953	
• Interbank Loans \tilde{I}_i	224953	Total Deposits from Banks (Reported)
• Repo \tilde{R}_i	224966 + 224969 + 224965	Memo: Repo Agreements in Deposits from Customers (Reported)+ Memo: Repurchase Agreements Not in Deposits (Reported) + Memo: Repo Agreements in Deposits from Customers (Reported)
• Other Liabilities \tilde{O}_i	$\max\{L_i - (D_i + \tilde{I}_i + \tilde{R}_i), 0\}$	
Equity \tilde{E}_i	$\{A_i - L_i\}$	
• CET1 Equity E_i	248877	Tier 1 Common Capital (CET1) (Reported)

B.1.2 Initialisation of Seniority Classes

In Table B.2 gives the initialisation of the seniority classes in our model using SNL data. Note that \tilde{I}_i , \tilde{O}_i , D_i and \tilde{R}_i are already defined in Table B.1. We do not have data on the amount of secured debt liabilities, hence we set x (in this table) to zero.

Table B.2: Initialisation of Seniority Classes

Seniority Class	Name	Data Identifier/Inference	Associated Data Name
1	CET1 Equity	248877	Tier 1 Common Capital (CET1) (Reported)
2	AT1 Equity	133172-248877	Tier 1 Capital (Reported) - Tier 1 Common Capital (CET1) (Reported)
3	T2 Capital	133173 - 133172	Total Capital (Reported) -Tier 1 Capital (Reported)
4	Subordinated Debt	$(225040+225030+225049)/$ $[(225040+225030+225049)+$ $(225031+225041+225050) +(x)] \times$ $(\tilde{I}_i + \tilde{O}_i)$	$(\text{Subordinated Debt Liabilities Held at Fair Value (Reported)}+)$ $\text{Subordinated Debt Liabilities Held for Trading (Reported)}+)$ $\text{Sub Debt Liabilities Held at Amortized Cost (Reported)})/$ $[(\text{Subordinated Debt Liabilities Held at Fair Value (Reported)}+)$ $\text{Subordinated Debt Liabilities Held for Trading (Reported)}+)$ $\text{Sub Debt Liabilities Held at Amortized Cost (Reported)})$ $+(\text{Senior Debt Liabilities Held for Trading (Reported)}+)$ $\text{Senior Debt Liabilities Held at Fair Value (Reported)}$ $+ \text{Senior Debt Liabilities Held at Amortized Cost (Reported)})+)$ $\text{Secured Debt Liabilities}] \times$ $(\tilde{I}_i + \tilde{O}_i)$
5	Senior Debt	$225031+225041+225050)/$ $[(225040+225030+225049)+$ $(225031+225041+225050)+(x)]$ $\times(\tilde{I}_i + \tilde{O}_i)$	$(\text{Senior Debt Liabilities Held for Trading (Reported)}+)$ $\text{Senior Debt Liabilities Held at Fair Value (Reported)}+)$ $\text{Senior Debt Liabilities Held at Amortized Cost (Reported)})/$ $[(\text{Subordinated Debt Liabilities Held at Fair Value (Reported)}+)$ $\text{Subordinated Debt Liabilities Held for Trading (Reported)}+)$ $\text{Sub Debt Liabilities Held at Amortized Cost (Reported)})+)$ $(\text{Senior Debt Liabilities Held for Trading (Reported)}+)$ $\text{Senior Debt Liabilities Held at Fair Value (Reported)}+)$ $\text{Senior Debt Liabilities Held at Amortized Cost (Reported)})+)$ $\text{Secured Debt Liabilities}] \times(\tilde{I}_i + \tilde{O}_i)$
6	Deposits	$D_i = 132288 - 224953$	Total Deposits (Reported) - Total Deposits from Banks (Reported)
7	Secured Debt	$\tilde{R}_i + \tilde{O}_i^{k^7} + \tilde{I}_i^{k^7} =$ $\tilde{R}_i + (x) / [(225040+225030+225049)+$ $(225031+225041+225050) +(x)] \times$ $(\tilde{I}_i + \tilde{O}_i)$	$\tilde{R}_i + (\text{Secured Debt Liabilities})/$ $[(\text{Subordinated Debt Liabilities Held at Fair Value (Reported)}+)$ $\text{Subordinated Debt Liabilities Held for Trading (Reported)}+)$ $\text{Sub Debt Liabilities Held at Amortized Cost (Reported)})+)$ $(\text{Senior Debt Liabilities Held for Trading (Reported)}+)$ $\text{Senior Debt Liabilities Held at Fair Value (Reported)}+)$ $\text{Senior Debt Liabilities Held at Amortized Cost (Reported)})+)$ $\text{Secured Debt Liabilities}] \times(\tilde{I}_i + \tilde{O}_i)$

Mapping between Liability Types and Seniority Classes

The hierarchy of bank i 's bail-inable debt, which runs from bail-inable debt in the first priority class $B_i^{k_1}$ to bail-inable debt in the fifth priority class $B_i^{k_5}$, consists of the following instruments (see a summary in Figure 4.3):

1. Priority class 1 consists of CET1 equity:

$$B_i^{k_1} = \tilde{E}_i := \sum_{j \in \mathcal{F}} E_{ji}; \quad (\text{B.1})$$

- where E_{ji} denotes the holding of institution $j \in \mathcal{F}$ (where \mathcal{F} is the set of financial institutions) of bank i 's CET1 equity \tilde{E}_i . The bank's CET1 equity \tilde{E}_i is approximated by

$$\tilde{E}_i^t \approx E_i^t - \Delta_i^{t0}, \quad (\text{B.2})$$

where Δ_i^{t0} is defined as the difference between book equity E_i and CET1 equity \tilde{E}_i at time zero (i.e. $\Delta_i^{t0} := E_i^{t0} - \tilde{E}_i^{t0}$). This is a reasonable approximation to capture how asset losses and liability changes effect the value of the CET1

equity \tilde{E}_i^t in a stress test. The CET1 equity \tilde{E}_i of a bank strongly relates to the book equity of a bank E_i , defined as

$$E_i =: A_i - L_i, \quad (\text{B.3})$$

but may not be equal to it. With this approximation, we assume that the difference between the equity E_i and the CET1 equity \tilde{E}_i is constant over time.

2. *Priority class 2 consists of additional tier I (AT1) equity:*

$$B_i^{k_2} =: \tilde{E}_i^{AT1} = \sum_{j \in \mathcal{F}} E_{ji}^{AT1}, \quad (\text{B.4})$$

- where E_{ji}^{AT1} denotes the holding of institution $j \in \mathcal{F}$ of bank i 's AT1 equity \tilde{E}_i^{AT1} . We note that all AT1 contracts E_{ji}^{AT1} are perpetual, else they would not be allowed to count towards a bank's AT1 equity \tilde{E}_i^{AT1} .

3. *Priority class 3 consists of tier II (T2) equity:*

$$B_i^{k_3} := \tilde{E}_i^{T2} = \sum_{j \in \mathcal{F}} \sum_{m \in \mathcal{M}} E_{ji}^{T2,m}, \quad (\text{B.5})$$

- where $E_{ji}^{T2,m}$ denotes the holding of institution $j \in \mathcal{F}$ with time to maturity $m \in \mathcal{M}$ of bank i 's T2 equity \tilde{E}_i^{T2} . The time to maturity m is defined as the difference between the maturity date T and the current time t (i.e. $m := T - t$).

4. *Priority class 4 consists of subordinate interbank contracts and other liabilities with a time to maturity of at least 7 days:*

$$B_i^{k_4} := \sum_{j \in \mathcal{B}} \sum_{m \in \mathcal{M}} (I_{ji}^{k_4 m} + O_{ji}^{k_4 m}), \quad (\text{B.6})$$

- where $I_{ji}^{k_4 m}$ and $O_{ji}^{k_4 m}$ denote the value of institution j 's holdings of bank i 's subordinated interbank contracts and other liabilities with time to maturity $m \in \mathcal{M}$.

5. *Priority class 5 consists of senior interbank contracts and other liabilities with a time to maturity of at least 7 days:*

$$B_i^{k_5} := \sum_{j \in \mathcal{B}} \sum_{m \in \mathcal{M}} (I_{ji}^{k_5 m} + O_{ji}^{k_5 m}), \quad (\text{B.7})$$

- where $I_{ji}^{k_5 m}$ and $O_{ji}^{k_5 m}$ denote the value of institution j 's holdings of bank i 's senior interbank contracts and other liabilities with time to maturity $m \in \mathcal{M}$.

B.2 Notation

Table B.3: Shows the definition of notation.

Category	Variable	Definition	
Balance sheet	E_i	The book equity of bank i .	
	\bar{E}_i	The CET1 equity of bank i .	
	$\bar{E}_i^{t,n}$	The CET1 equity of bank i in Monte Carlo run n at time t .	
	E_{ji}	The CET1 equity of bank i held by institution j .	
	E_{ji}^{km}	The CET1 equity of bank i in priority class k with time to maturity m held by institution j .	
	E_{ji}^{AT1}	The AT1 equity of bank i .	
	E_{ji}^{AT1}	The AT1 equity of bank i held by bank j .	
	E_i^{T2}	The T2 equity of bank i .	
	$E_{ji}^{T2,m}$	The T2 equity of bank i with a time to maturity m held by bank j .	
	F_i	The total funds of bank i .	
	A_i	The asset value of the bank i .	
	$A_i^{t,n}$	The asset value of bank i in Monte Carlo run n at time t .	
	L_i	The liability value of bank i .	
	C_i	The cash of bank i .	
	Y_i	The external asset value of bank i .	
	T_i	The tradable asset value of bank i .	
	I_i	The interbank asset value of bank i .	
	\bar{L}_i	The interbank liability value of bank i .	
	Bail-inable debt	I_{ij}^{km}	The interbank assets in priority class k with time to maturity m provided by bank i to bank j .
		R_i	The reverse repo value of bank i .
\bar{R}_i		The repo value of bank i .	
O_i		The value of the other assets of bank i .	
\bar{O}_i		The value of the other liabilities of bank i .	
O_{ji}^{km}		The other liabilities of bank j in priority class k with time to maturity m held by institution j .	
B_i		The bail-inable debt value of bank i .	
B_i^k		The bail-inable debt value of bank i in priority class k .	
B_{ji}^{km}		The bail-inable debt of bank i in priority class k with time to maturity m held by institution j .	
\bar{B}_i		The bail-inable debt value of bank i excluding debt in the first priority class (i.e. CET1 equity).	
Risk-weighted capital ratio	ρ_i	The risk-weighted capital ratio of bank i .	
	ρ_i^{data}	The risk-weighted capital ratio of bank i as given data.	
	$\rho_i^{t,n}$	The risk-weighted capital ratio of bank i in Monte Carlo run n at time t .	
	ρ_i^F	The failing-likely to fail (FLTF) (risk-weighted capital) ratio applicable to bank i .	
	ρ_i^T	The (risk-weighted capital ratio) recapitalisation target applicable to bank i .	
Time	t	Current time.	
	t_s	Current time in simulation (subscript is only used if necessary to avoid confusion).	
	T	Maturity of a contract.	
	\bar{T}	Length of a contract (in days).	
	m	Time to maturity ($m := T - t$).	
	τ_i	Bail-in time of bank i .	
	$\tau_i^{m,n}$	Bail-in time of bank i in Monte Carlo run n that affects contracts with a time to maturity m .	
Loss absorption & recapitalisation amount	τ_a	State of balance sheet after the loss absorption phase of the bail-in (phase a) has been completed.	
	τ_b	State of the balance sheet after the recapitalisation phase of the bail-in (phase b) has been completed.	
	l_i^τ	Necessary loss absorption amount of bank i at time τ .	
	l_i^τ	Feasible loss absorption amount of bank i at time τ .	
	\bar{b}_i	Preferred recapitalisation amount of bank i .	
b_i	Feasible recapitalisation amount of bank i .		

Table B.4: Shows the definition of notation.

Category	Variable	Definition
Haircuts	h_i^k	Total haircut applied to priority class k of bank i .
	$h_i^{k\tau_a}$	Total haircut applied to priority class k of bank i in the loss absorption phase (phase a).
	$h_i^{\tau_b k}$	Total haircut applied to priority class k of bank i in the recapitalisation phase (phase b).
	$h_{ji}^{km\tau_a}$ $h_{ji}^{km\tau_b}$	Haircut applied in phase a to a contract in priority class k with time to maturity m held by creditor j of bank i . Haircut applied in phase b to a contract in priority class k with time to maturity m held by creditor j of bank i .
Conversion rates	Δ_{ia}^k	Conversion rate (or: debt-to-equity conversion rate) applicable to contracts in priority class k of bank i .
	Δ_{ib}^k	Conversion rate that holds for the loss absorption phase applicable to contracts in priority class k of bank i .
	Δ_a	Conversion rate that holds for the recapitalisation phase applicable to contracts in priority class k of bank i .
	Δ_b	'Fair' conversion rate in the loss absorption phase (phase a).
	Δ_{ia}^k	'Fair' conversion rate in the recapitalisation phase (phase b).
	Δ_{ib}^k	'Unfair' conversion rate component that holds for phase a applicable to contracts in priority class k of bank i . 'Unfair' conversion rate component that holds for phase b applicable to contracts in priority class k of bank i .
Acquired equity share	$\epsilon_{ji}^{km\tau_b}$	The acquired share (at time τ_b) of bank i 's equity for a contract in priority class k with time to maturity m held by institution j .
	$\epsilon_{ji}^{km\tau_a}$	The 'fair' acquired share (at time τ_b) of bank i 's equity for a contract in priority class k with time to maturity m held by institution j .
	$\tilde{\epsilon}_{ji}^{km\tau_b}$	The 'unfair' part of the acquired share of bank i 's equity for a contract in priority class k with time to maturity m held by institution j .
	$\epsilon_{ji}^{k_1\tau}$	The share of an existing equity holder (i.e. priority class k_1) j of bank i 's equity at the start time of bail-in τ .
Number of shares	η_i^a	The number of shares of bank i at the start time of bail-in τ .
	η_i^a	The number of shares of bank i after the loss absorption phase (phase a) of bail-in.
	η_i^b	The number of shares of bank i after the recapitalisation phase (phase b) of bail-in.
	w	The arbitrarily chosen number of newly created shares if equity holders are wiped out.
Valuation of bail-inable debt	V_{ji}^{kmt}	The time- t value of a bail-inable debt contract of bank i in priority class k with time to maturity m held by institution j .
	I_{ji}^{kT}	The payoff at maturity T of a bail-inable debt contract of bank i in priority class k held by institution j .
Jump process of asset value	r	The risk-free rate (we set $r = 0$).
	σ_i^2	The instantaneous variance of bank i 's asset returns conditional on a Poisson event not occurring.
	W_i^Ω	Gauss-Wiener process of bank i under the risk-neutral dynamics.
	Z_i	Value of a draw from the standard normal distribution for bank i .
	λ_i	The jump intensity of bank i .
	j_i	The size of a random jump of bank i .
	\tilde{j}_i	The mean random jump size of bank i .
	μ_i^j	The mean for bank i of the log-normal distribution associated to the jump size j_i .
	σ_i^j	The standard deviation for bank i of the log-normal distribution associated to the jump size j_i .
ϕ_i	A Boolean variable for bank i that takes value one with probability λ_i and value zero with probability $1 - \lambda_i$.	
Runs	ρ_{ji}^k	The threshold value of the risk-weighted capital ratio of bank i beyond which a creditor j attempts to run.
	Λ_{ji}^{kmt}	The relative valuation loss at time t of a bail-inable contract of bank i in priority class k with time to maturity m held by institution j .
	$\tilde{\Lambda}_{ji}^{kmt}$	The VaR valuation loss at time t of a bail-inable debt contract of bank i in priority class k with time to maturity m held by institution j .
	ψ_{ji}^k $\tilde{\psi}_{ji}^k$	The threshold value of the relative valuation loss Λ_{ji}^{kmt} of bank i beyond which a creditor j attempts to run. The threshold value of the VaR valuation loss $\tilde{\Lambda}_{ji}^{kmt}$ of bank i beyond which a creditor j attempts to run.

B.3 Model Specifics

B.3.1 Alternative Ways of Dealing with Bank Failure

Bail-Out A bank $i \in \mathcal{B}$ can be recapitalised from a risk-weighted capital ratio ρ_i^t at time t to a recapitalisation target ρ_i^T (see section 4.5.1.4) by means of a bail-out. Bail-out is achieved by injecting equity in the bank. The equity injection also gives a bank an equal amount of cash. So other than in a bail-in, a bail-out also *directly* addresses any liquidity issues a bank may face. Specifically, the amount of equity a_i that a government needs to inject to recapitalise a bank at time t is given by

$$a_i = \Omega_i^t \rho_i^T - \tilde{E}_i^t, \quad (\text{B.8})$$

which follows from

$$\rho_i^T = \frac{\tilde{E}_i^t + a_i}{\Omega_i^t}. \quad (\text{B.9})$$

Equation B.9 tells that, due to the bail-out, bank i 's CET1 equity updates to $\tilde{E}_i^{t+1} = \tilde{E}_i^t + a_i$ (numerator) and its cash updates to $C_i^{t+1} = C_i^t + a_i$ (denominator). However, the

cash does not show up in the risk-weighted assets Ω_i (denominator), since it bears a zero risk-weight.

Bail-out can be performed once the bank has failed (i.e. at time $t = \tau_i$), if the loss absorption phase of bail-in was not successful (i.e. at time $t = \tau_a$), or if phase b was not successful (i.e. at time $t = \tau_b$). In the first case a bail-out rather than a bail-in is used to fully recapitalise the bank, whereas in the latter two cases bail-out is performed to further recapitalise the bank because the bail-inable debt B_i was not sufficient.

Liquidation Liquidation is performed in the same way as the liquidation procedure described in Farmer et al. (2020). Assets are liquidated (i.e. tradable assets are sold and short-term funding is pulled). Liabilities face a loss given default of a hundred percent, since in the short-time period that the stress test considers the recovery is typically zero.

B.4 Sensitivity Analysis




Figure B.1: Shows the stability impact of the ‘primary’ bail-in design, when *smaller* SIBs fail. The left x-axis now displays the failure of the 5 smallest SIBs, rather than the 5 largest SIBs as the main results do. We observe that stability is left unaffected by the choice of ‘good’ or ‘bad’ bail-in parameters for idiosyncratic failures of small European SIBs.

Comparative stability impact of failure proceedings

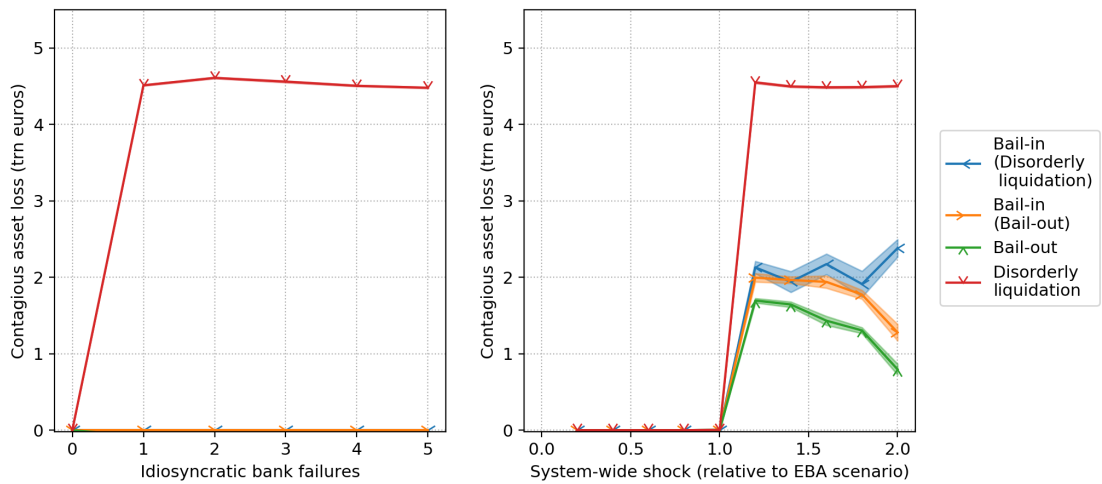


Figure B.2: Compares the stability impact of the leading modes of dealing with bank failure: disorderly liquidation, bail-out, and bail-in (disorderly liquidation if bail-in is not succesful; or bail-out if bail-in is not succesful). It does so for the case where the ‘primary’ bail-in parameters are well-chosen. And the same parameters, if applicable, apply to the other methods of dealing with SIB failure.

Comparative stability impact of failure proceedings




Figure B.3: Compares the stability impact of the leading approaches of dealing with bank failure: disorderly liquidation, bail-out, and bail-in (disorderly liquidation if bail-in is not succesful, or bail-out if bail-in is not succesful). It does so for the case where the ‘primary’ bail-in parameters are ill-chosen. And the same parameters, if applicable, apply to the other means of dealing with SIB failure.

Failing-likely-to-fail ratio




Figure B.4: Shows the stability impact of the failure threshold, for the case where ‘bad’ primary parameters are applied.

Recapitalisation target

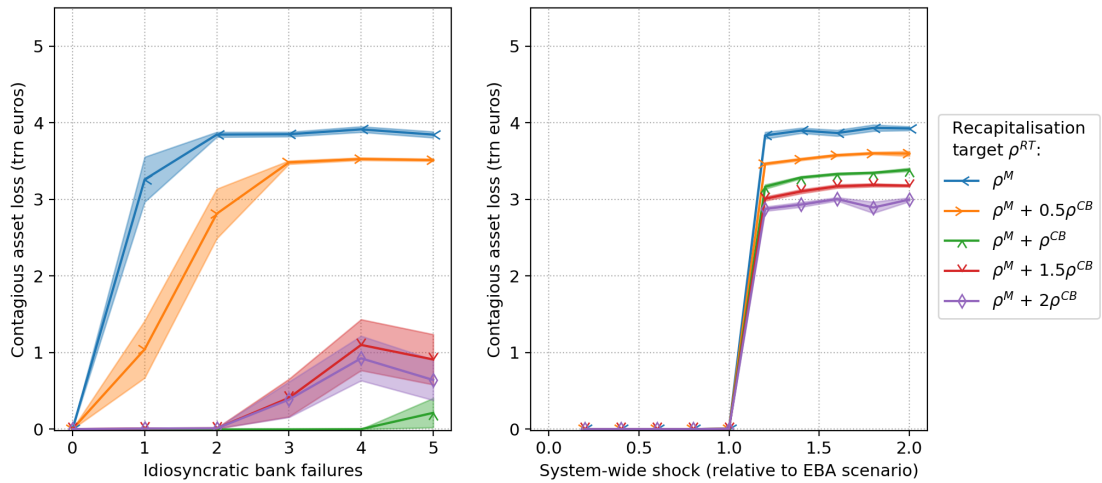


Figure B.5: Shows the stability impact of the recapitalisation target, for the case where ‘bad’ primary parameters are used.




Figure B.6: Shows the contagious amplifications of the ‘primary’ bail-in design, for the case where bad’ primary parameters are wielded.




Figure B.7: Shows the stability impact of the initial leverage of leveraged non-banks. We observe that systemic risk in the banking system spikes if the initial leverage of non-banks is higher.

Leveraged non-bank size impacts bail-in stability

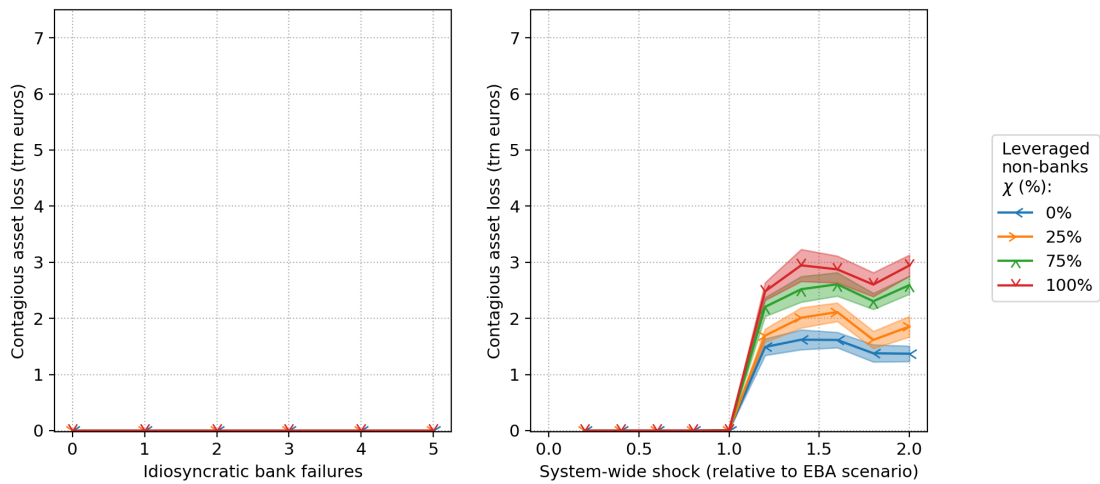


Figure B.8: Shows the stability impact of the relative size of leveraged non-banks χ vs. non-leveraged non-banks $(1 - \chi)$. We observe that banking sector fragility surges if a relative larger percentage of the bail-inable debt that is held by the non-banking sector is held by leveraged non-banks.

Asset loss to real economy

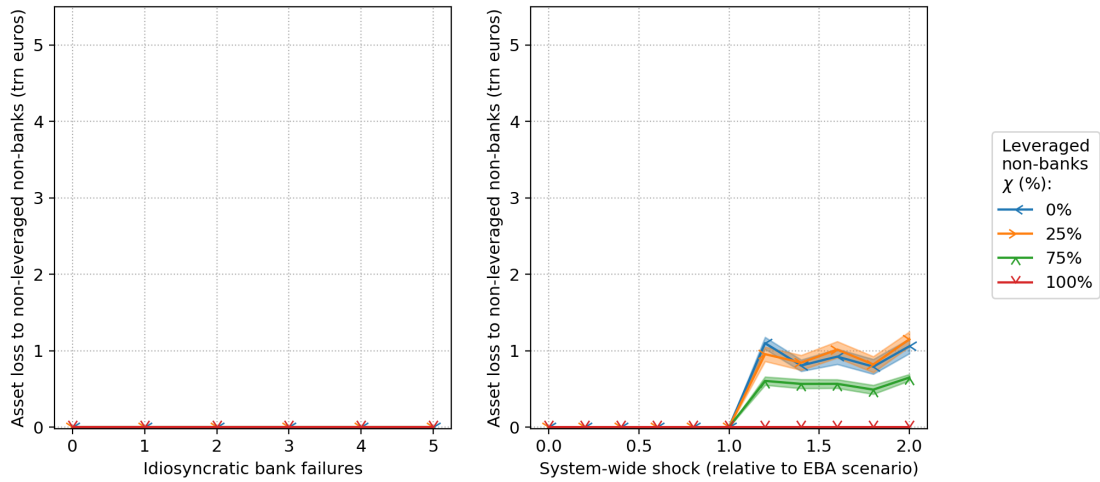


Figure B.9: Shows the asset loss to non-leveraged non-banks – a good proxy for the asset losses that the real economy may suffer – as a function of the percentage of bail-inable debt held by the non-banking sector that is possessed by non-leveraged non-banks, such as pension funds and insurance companies.

B.5 Detailed Descriptions

B.5.1 Comparing Debt Eligibility Criteria for MREL, TLAC, and Bail-In

In this section we will compare the main commonalities and differences between instruments that are eligible for: (i) bail-in B_i (these have been discussed in Section 4.5.1.1); (ii) TLAC T_i ; and (iii) MREL M_i . A complete description of the eligible instruments for bail-in B_i , TLAC T_i and MREL M_i are given in: Article 44 of the BBRD, [FSB \(2015b\)](#), and Article 45 of the BRRD and [EBA \(2016\)](#), respectively.

Commonalities The main commonalities are twofold. First, CET1 equity \tilde{E}_i , AT1 capital \tilde{E}_i^A , and tier 2 capital \tilde{E}_i^T are eligible for bail-in B_i , MREL M_i , and TLAC T_i . Second, structured products, covered deposits and non-contracts (e.g. tax liabilities) are excluded from bail-in B_i , as well as from MREL M_i and TLAC T_i .

Differences The main differences regard the treatment of: (a) the maturity; (b) derivatives; (c) cross holdings; (d) buffer requirements; (e) the denominator; and (f) subordination. We will discuss each of these in turn.

Maturity The *time to maturity* $T - t$ of instruments that are eligible for MREL M_i and TLAC T_i is at least one year, whereas instruments to institutions with a *maturity* T greater than or equal to seven days, or instruments to settlement systems with a *time to maturity* $T - t$ of at least seven days are eligible for bail-in B_i . Why does a discrepancy exist?

By making only those contracts eligible for TLAC and MREL that have a residual maturity of at least a year, regulators encourage banks to have a maturity profile of bail-inable debt B_i that is sufficiently long that the bail-inable debt amount B_i cannot suddenly collapse (e.g. if counterparties do not roll-over debt). This rule in principle allows the resolution authority to make use of the loss absorbing capacity with a residual maturity less than one year but greater than seven days. So it gives extra capacity, while encouraging stable forms of bail-inable claims B_i . Yet, in our opinion this inconsistency is not helpful. As we explained in the Section about runs 4.5.2.3, instruments that are maturity-wise not eligible for MREL M_i and TLAC T_i but are eligible for bail-in B_i may be subject to a run in anticipation of a bail-in. In the run Section 4.5.2.3 and the result Section 4.7, we will argue and show that the minimum maturity of instruments that are excluded needs to be increased. Perhaps to a level in line with the TLAC T_i and MREL M_i exclusions.

Derivatives Derivatives are excluded from MREL M_i and TLAC T_i , while derivatives are in principle included in bail-in B_i . Yet, authorities may exclude derivatives on an ad-hoc basis (such as when bailing in derivatives would give rise to contagion) (see Section 4.5.1.1).¹ Why does a disparity exist?

Authorities presumably anticipate that it may be hard to bail-in derivatives. Therefore, to ensure a bank has sufficient bail-inable claims B_i that in practise (rather than in principle) can be bailed in, regulators exclude derivatives from the measures (TLAC and MREL) that aim to ensure that a bank has sufficient loss absorbing capacity.

Cross Holdings Cross holdings of MREL instruments by banks (both G-SIBs and non-GSIBs) and cross holdings of TLAC instruments by G-SIBs are excluded (subject to some exceptions) from the MREL M_i and TLAC T_i measure respectively.²³ However, cross holdings of bail-inable debt are not excluded from a bank's bail-inable debt B_i . Yet, the authority has the power to exclude cross holdings of bail-inable debt in exceptional circumstances.⁴ Why this dissimilarity?

The MREL M_i and TLAC T_i measure exclude cross holdings of MREL and TLAC instruments to avoid the risk of contagion within the (G-SIB) banking sector. However, these measures do not exclude the risk of contagion between the banking sector and the non-banking sector and the risk of contagion within the non-banking sector.

Despite the TLAC and MREL exclusions of cross holdings from the loss absorbing requirements, cross holdings of TLAC and MREL may be bailed in. This rule is presumably in place to provide regulators with extra loss absorbing capacity if bailing in cross holdings would not risk contagion. These are the most important differences from the point of view of contagion. Three further differences remain.

Buffers As discussed in Section 4.5.1.2, CET1 equity \tilde{E}_i that is counted towards the buffer standards ρ_i^{CB} and λ_i^{CB} cannot count towards TLAC T_i or MREL M_i . However, this part of CET1 equity is subject to bail-in B_i . This difference makes sense to keep buffer usable.

The MREL and TLAC requirements are computed with respect to different denominators (as we saw in section 4.5.1.2). MREL requirements are computed with respect to

¹See: Article 44(3) of the BRRD.

²In the case of TLAC this is brought about by the requirement that TLAC cross holdings by G-SIBs are subtracted from a bank's T2 \tilde{E}_i^A capital, and if this is not sufficient from its AT1 \tilde{E}_i^T and CET1 equity \tilde{E}_i next (FSB (2016b)). Note, that this requirement is a double blow. It reduces both TLAC T_i and a bank's capital requirements.

³In the case of MREL, exceptions include that MREL cross holdings by banks below a certain threshold may count towards MREL M_i , so as to support market-making activities (EBA (2016)).

⁴See: Article 44(3) of the BRRD.

a bank's liabilities and own funds $F_i + L_i$, while TLAC requirements are determined with respect to a bank's risk exposure Ω_i . Strikingly, regulators have not specified a measure to determine the minimum amount of bail-inable debt B_i . Instead the set of instruments that are eligible for bail-in B_i is a superset of the instruments that are eligible for MREL M_i and TLAC T_i (as we saw from point (a), (b) and (c)),⁵ because the MREL M_i^M and TLAC T^M requirements implicitly set a floor to the minimum amount of bail-inable instruments B_i , which a bank must have.

Subordination To be eligible towards TLAC T_i instruments must be contractually and statutory subordinated (subject to some exceptions). Contractual subordination means that TLAC eligible instruments must be contractually subordinated to instruments that are explicitly excluded from TLAC.⁶ Statutory subordination means that a TLAC eligible instrument must be junior in the creditor hierarchy to instruments that are explicitly excluded from TLAC. The main reason why regulators have introduced subordination is to circumnavigate the no creditor worse-off condition (NCWO) safeguard (EBA (2016)). This safeguard requires that no shareholder or creditor must be left worse off from the use of resolution, else the shareholder or creditor is entitled to the payment of the difference from the resolution financing arrangements (in Section 4.5.1.2 we discuss why exclusions make a breach of the no-creditor-worse-off condition more likely, and therefore why subordination is useful).⁷ Instruments, which are eligible towards MREL M_i originally did not have a subordination requirement. However, EBA (2016) proposes that under the revised framework G-SIBs (not non-GIBs) should have a subordination requirement in line with TLAC.⁸

Bail-inable claims B_i do not have to be subordinated contractually or statutory to instruments that are explicitly excluded from bail-in as defined in Article 44(2) of the BRRD. Why is there a lack of congruence?

By requiring subordination, in order for instruments to be eligible for TLAC T_i and MREL M_i , regulators encourage banks to hold more subordinated debt. This increases the bail-inable debt amount B_i of banks that is at a low risk of breaching the NCWO principle. Why not make the bail-inable debt B_i consistent with MREL M_i and TLAC T_i regarding subordination? This would reduce investors' uncertainty regarding the instruments are likely subject to bail-in and allow them to price risk better.

⁵That is, typically it would be the case that $\hat{B}_i \geq T_i$ and $\hat{B}_i \geq M_i$

⁶The full list of contracts that are excluded from TLAC can be found in FSB (2015b).

⁷See: Article 75 of the BRRD.

⁸Specifically, under the revised framework, G-SIBs should be required to meet their MREL with subordinated instruments, at least to a level of 16% of RWAs in 2019 and 18% of RWAs in 2022 in line with the TLAC term sheet (EBA (2016)).

Closing Remarks To conclude, we have seen that instruments that are eligible for bail-in B_i substantively differ from the instruments that are eligible from MREL M_i and TLAC T_i . Also, differences in eligibility between TLAC T_i and MREL M_i instruments exist. Since both TLAC and MREL requirements serve the same purpose: to establish a minimum loss absorbing capacity for banks (which can be used in a bail-in), it makes sense to eliminate any discrepancies between TLAC T_i and MREL M_i eligible instruments. Eligibility to count towards TLAC T_i and MREL M_i is more conservative than eligibility towards bail-inable debt. In some sense this makes sense. Regulators require easily usable loss absorbing capacity but also allow less easily usable loss absorbing capacity to be used in a bail-in would the former not be sufficient in size to restore a bank to viability. On the other hand, it seems odd that if the aim is to increase the loss absorbing capacity in a bail-in, regulators – rather than setting a minimum amount of bail-inable debt requirement (as this is what actually matters in equation 4.15 and equation 4.25, which stipulate the maximum loss absorbing capacity and maximum recapitalisation capacity respectively) – define a different yet related measure. Therefore, we propose the following minimum requirement. The minimum amount of bail-inable debt B_i that banks must hold must satisfy the following equation

$$B_i \geq B_i^M \Omega_i, \quad (\text{B.10})$$

where B_i^M is some specified minimum set on the amount of risk-weighted bail-inable debt that bank i should hold. The question remains: what claims should be eligible as bail-inable debt B_i ? Either authorities could stick to the current standard as enshrined in Article 44(2). Or if authorities think that certain bail-inable debt in Article 44(2) is in practise not bail-inable, regulators should update 44(2) to make it consistent with the instruments that are eligible for TLAC.

B.5.2 Principles on the Debt-to-Equity Conversion Rate

B.5.2.1 Only ‘Fair’ Conversion Rates are Typically Compatible with P1 & P2

It turns out that it is much easier to satisfy P1 and P2 when the conversion rate Δ_i^k is split in the conversion rate applicable to haircuts in the loss absorption phase Δ_{ia}^k and to haircuts in the recapitalisation phase Δ_{ib}^k (see equation 4.33 for the relation among the phase- a - and phase- b conversion rate and the overall conversion rate Δ_i^k). Typically, the only compatible conversion rate with P1 and P2 is the ‘fair’ conversion rate (exceptions are discussed in Section 4.5.1.2). We will now explain why this is so.

For now let us assume (we will relax these in Section 4.5.1.2) that the liquidation cost

c_i^τ in a hypothetical liquidation is zero (see equation 4.56) and that there are no exclusions X_i^τ of claims in a bail-in relative to the insolvency hierarchy (see equation 4.68). In such case, the sole conversion rate in phase a and b compatible with P1 and P2 is the ‘fair’ conversion rate given by $\Delta_{ia}^k = \tilde{\Delta}_a = 0$ and $\Delta_{ib}^k = \tilde{\Delta}_b$ for $k \in \mathcal{K}$ (see equation 4.40, 4.41 and 4.42). When conversion rates in phase a and b are set ‘fairly’ the joint conversion rate Δ_i^k is also compatible with P1 and P2. These can be seen as follows.

- A ‘fair’ conversion rate $\tilde{\Delta}_a = 0$ in phase a satisfies P1: The loss absorption need \hat{l}_i in a bail-in is equal to that in a hypothetical liquidation (see equation 4.14). Hence, if conversion rates Δ_{ia}^k are set to zero in phase a of a bail-in, then creditors are not worse off than in a liquidation: the loss given default $\zeta_i^{k\tau_b}$ in each priority class k is equal (see equation ??).

Applying positive conversion rates $\Delta_{ia}^k > 0$ in phase a makes creditors who receive a haircut in phase b worse off. In part this can be understood from equation 4.36. Equation 4.36 shows that allotting shares is a zero-sum game. If creditors in phase a receive more than their ‘fair’ share (i.e. $\Delta_{ia}^k > \tilde{\Delta}_a = 0$), then creditors in phase b by implication obtain less than their fair share (i.e. $\Delta_{ib}^k < \tilde{\Delta}_b$).

- A ‘fair’ conversion rate $\tilde{\Delta}_b > 0$ in phase b satisfies P1: Creditors who are subject to haircuts in the recapitalisation phase should not suffer any losses if P1 is to be satisfied, since in liquidation losses are absorbed but no recapitalisation is done. Therefore, fair rates $\tilde{\Delta}_b$ should be applied so that any haircut $h_{ji}^{km\tau_b}$ in phase b (see equation 4.13) is replaced with an equal claim $E_{ji}^{km\tau_b} = \tilde{\epsilon}_{ji}^{km\tau_b} \tilde{E}_i^{\tau_b}$ of bank i ’s CET1 equity $\tilde{E}_i^{\tau_b}$ (see equation 4.39), resulting in no net losses.
- ‘Fair’ conversion rates in phase a and b jointly satisfy P2: P2 holds because $\tilde{\Delta}_a < \tilde{\Delta}_b$ and creditors in more junior priority classes are the first to be subjected to phase a , while creditors in more senior priority classes will likely be (only) subjected to haircuts in phase b . The hierarchy of claims is thus preserved.
- ‘Fair’ conversion rates in phase a and b imply the joint conversion rates Δ_i^k satisfy P1 and P2: P1 clearly holds for the joint conversion rates Δ_i^k , since it holds for the individual conversion rates in phase a and b . From equation

$$\Delta_i^k = \frac{\tilde{\Delta}_b h_{ji}^{km\tau_b}}{h_{ji}^{km\tau_a} + h_{ji}^{km\tau_b}}, \quad (\text{B.11})$$

which simplifies equation 4.33 for the case of ‘fair’ conversion rates, also shows clearly that P2 is satisfied: creditors in a junior priority class k more often have a

positive haircut $h_{ji}^{km\tau_a} > 0$ in phase a than creditors in a more senior priority class do, leading to a smaller or equal conversion rate Δ_i^k for more junior priority classes.

Setting conversion rates ‘fairly’ is not only easy, as well as compatible with P1 and P2, it is also desirable. Setting rates other than ‘fairly’, that is, ‘unfairly’ gives excessive losses to some (i.e. junior creditors) and net profits to others (i.e. senior creditors). This not only skews the wealth distribution and could be argued to be ‘disproportionate’⁹, it also risks exacerbating contagion (our main concern).

We will now proceed to formalise P1 so that we can discuss the case when ‘unfair’ rates are possible without breaching P1 and P2 and treat the situation where neither ‘fair’ nor ‘unfair’ rates are compatible with P1 and P2 necessitating a contribution of the resolution financing fund.

⁹A regulator in principle is allowed to convert a debt claim to a certain amount of equity greater or smaller than the original converted debt claim, as long as it is not disproportionate. See [EBA \(2017c\)](#) point (1.18).

Bibliography

- Abad, J., Aldasoro, I. n., Aymanns, C., D'Errico, M., Rousová, L. F., Hoffmann, P., Langfield, S., Neychev, M. & Roukny, T. (2016), 'Shedding light on dark markets: First insights from the new EU-wide OTC derivatives dataset', ESRB Occasional Paper Series **10**, 1–32.
- Acemoglu, D., Ozdaglar, A. & Tahbaz-Salehi, A. (2015), 'Systemic risk and stability in financial networks', The American Economic Review **105**(2), 564–608.
- Acharya, V., Engle, R. & Pierret, D. (2014), 'Testing macroprudential stress tests: The risk of regulatory risk weights', Journal of Monetary Economics **65**, 36–53.
- Acharya, V. V. & Skeie, D. (2011), 'A model of liquidity hoarding and term premia in interbank markets', Journal of Monetary Economics **58**(5), 436–447.
- Acharya, V. V. & Yorulmazer, T. (2008), 'Information contagion and bank herding', Journal of Money, Credit and Banking **40**(1), 215–231.
- Admati, A. & Hellwig, M. (2014), The bankers' new clothes: What's wrong with banking and what to do about it, Princeton University Press, Princeton NJ.
- Adrian, T. & Shin, H. S. (2010), 'Liquidity and leverage', Journal of Financial Intermediation **19**(3), 418–437.
- Afonso, G., Kovner, A. & Schoar, A. (2011), 'Stressed, not frozen: The federal funds market in the financial crisis', The Journal of Finance **66**(4), 1109–1139.
- Aikman, D., Alessandri, P., Eklund, B., Gai, P., Kapadia, S., Martin, E., Mora, N., Sterne, G. & Willison, M. (2009), 'Funding liquidity risk in a quantitative model of systemic stability'.
- Aikman, D., Chichkanov, P., Douglas, G., Georgiev, Y., Howat, J. & King, B. (2019), 'System-wide stress simulation', Staff Working Paper (809).

- Akerlof, G. A. (1970), ‘The market for “lemons”: Quality uncertainty and the market mechanism’, Quarterly Journal of Economics **84**(3), 488–500.
- Aldasoro, I., Gatti, D. D. & Faia, E. (2017), ‘Bank networks: Contagion, systemic risk and prudential policy’, Journal of Economic Behavior & Organization **142**, 164–188.
- Alessandri, P., Gai, P., Kapadia, S., Mora, N., Pühr, C. et al. (2009), ‘Towards a framework for quantifying systemic stability’, International Journal of Central Banking **5**(3), 47–81.
- Allen, F. & Gale, D. (2000), ‘Financial contagion’, Journal of Political Economy **108**(1), 1–33.
- Amini, H., Cont, R. & Minca, A. (2013), ‘Resilience to contagion in financial networks’, Mathematical Finance .
- Anand, K., Bédard-Pagé, G. & Tractlet, V. (2014), ‘Stress testing the canadian banking system: a system-wide approach’, Financial System Review **61**.
- Anand, K., Gauthier, C. & Souissi, M. (2015), Quantifying contagion risk in funding markets: A model-based stress-testing approach, Technical report.
- Anand, K., van Lelyveld, I., Banai, Á., Friedrich, S., Garratt, R., Halaj, G., Figue, J., Hansen, I., Jaramillo, S. M., Lee, H. et al. (2017), ‘The missing links: A global study on uncovering financial network structures from partial data’, Journal of Financial Stability **35**, 107–119.
- Anderson, P. W. et al. (1972), ‘More is different’, Science **177**(4047), 393–396.
- Anderson, R., Danielson, J., Baba, C., Das, U. S., Kang, H. & Segoviano, M. (2018), Macroprudential stress tests and policies: Searching for robust and implementable frameworks, Technical report, IMF Working Paper.
- Anderson, R. W. (2016), ‘Stress testing and macroprudential regulation: A transatlantic assessment’, Stress Testing and Macroprudential Regulation pp. p1–30.
- Arinaminpathy, N., Kapadia, S. & May, R. M. (2012), ‘Size and complexity in model financial systems’, Proceedings of the National Academy of Sciences **109**(45), 18338–18343.
- Armour, J., Awrey, D., Davies, P., Enriques, L., Gordon, J. N., Mayer, C. & Payne, J. (2016), Principles of Financial Regulation, Oxford University Press, Oxford.

- Arnold, M. & Jenkins, P. (2016), ‘Eba boss says recapitalisation of banks has been successful’, <https://www.ft.com/content/5496cc86-5417-11e6-befd-2fc0c26b3c60>.
- Avgouleas, E. & Goodhart, C. (2015), ‘Critical reflections on bank bail-ins’, Journal of Financial Regulation **1**(1), 3–29.
- Avgouleas, E. & Goodhart, C. (2016), A critical evaluation of bail-in as a bank recapitalisation mechanism, in ‘New International Financial System: Analyzing the cumulative impact of regulatory reform’, World Scientific Studies, New Jersey/London, pp. 267–306.
- Awrey, D. (2019), ‘Split derivatives: Inside the world’s most misunderstood contract’, Yale J. on fReg. **36**(2), 495–574.
- Aymanns, C., Caccioli, F., Farmer, J. D. & Tan, V. W. (2016), ‘Taming the Basel leverage cycle’, Journal of Financial Stability **27**, 263–227.
- Aymanns, C. & Farmer, J. D. (2015), ‘The dynamics of the leverage cycle’, Journal of Economic Dynamics and Control **50**, 155–179.
- Aymanns, C., Farmer, J. D., Kleinnijenhuis, A. M. & Wetzer, T. (2018), ‘Models of financial stability and their application in stress tests’, Handbook of Computational Economics **4**, 329–391.
- Aymanns, C., Georg, C.-P. & Golub, B. (2017), ‘Illiquidity spirals in coupled over-the-counter markets’, University of St. Gallen, School of Finance Research Paper (2018/10).
- Ban (2014), ‘Stress testing the UK banking system: 2014 results’.
- Bank of England (2015), The framework of capital requirements for UK banks, Technical report, Bank of England.
- Bank of England (2018), Stress testing the UK banking system: 2018 results, Technical report, Bank of England.
- Baranova, Y., Coen, J., Lowe, P., Noss, J. & Silvestri, L. (2017), ‘Simulating stress across the financial system: the resilience of corporate bond markets and the role of investment funds’, Bank of England Financial Stability Paper (42).
- Bardoscia, M., Barucca, P., Brinley, A. & Hill, J. (2017), ‘The decline of solvency contagion risk’, Bank of England working paper 662 .

- Bardoscia, M., Bianconi, G. & Ferrara, G. (2018), ‘Multiplex network analysis of the UK OTC derivatives market’, Bank of England Working Paper 726 .
- Battiston, S., Farmer, J. D., Flache, A., Garlaschelli, D., Haldane, A. G., Heesterbeek, H., Hommes, C., Jaeger, C., May, R. & Scheffer, M. (2016), ‘Complexity theory and financial regulation’, Science **351**(6275), 818–819.
- Battiston, S., Gatti, D. D., Gallegati, M., Greenwald, B. & Stiglitz, J. E. (2012), ‘Liasons dangereuses: Increasing connectivity, risk sharing, and systemic risk’, Journal of Economic Dynamics and Control **36**(8), 1121–1141.
- Battiston, S., Puliga, M., Kaushik, R., Tasca, P. & Caldarelli, G. (n.d.), ‘Debtrank: Too central to fail? Financial networks, the FED and systemic risk’, Scientific reports **2**.
- BCBS (2004), International convergence of capital measurement and capital standards, Technical report, Bank of International Settlements.
- BCBS (2009), Principles for sound stress testing practices and supervision, Technical report, Basel Committee on Banking Supervision, Bank of International Settlements.
- BCBS (2011), Basel III: A global regulatory framework for more resilient banks and banking systems, Technical report, Basel Committee on Banking Supervision.
- BCBS (2013), Supervisory framework for measuring and controlling large exposures, Technical report, Basel Committee Banking Supervision.
- BCBS (2016), Literature review on integration of regulatory capital and liquidity instruments, Technical report, Basel Committee on Banking Supervision.
- Behn, M., Haselmann, R. F. & Vig, V. (2016), ‘The limits of model-based regulation’, ECB Working Paper 1928 .
- Berkowitz, J. & O’Brien, J. (2002), ‘How accurate are value-at-risk models at commercial banks?’, The Journal of Finance **57**(3), 1093–1111.
- Bernanke, B. (2013), ‘Stress testing banks: What have we learned?’, <https://www.federalreserve.gov/newsevents/speech/bernanke20130408a.htm>.
- Bernanke, B. (2015), The courage to act: a memoir of the crisis and its aftermath, W.W. Norton and Company, New York.
- Bernanke, B. S. (2010), Remarks on the Squam Lake Report – fixing the financial system, Technical report, Squam Lake Conference.

- Bernanke, B. S. (2017), Why Dodd-Frank’s orderly liquidation authority should be preserved, Technical report, Brookings Institute.
- Bernard, B., Capponi, A. & Stiglitz, J. E. (2017), Bail-ins and bail-outs: Incentives, connectivity, and systemic stability, Technical report, National Bureau of Economic Research.
- BIS (2008), Addressing financial system procyclicality: a possible framework, Technical report, Bank of International Settlements.
- BIS (2009), Strengthening the resilience of the banking sector, Technical report, Bank of International Settlements.
- BIS (2010), Countercyclical capital buffer proposal, Technical report, Bank of International Settlements.
- BIS (2012), A framework for dealing with domestic systemically important banks, Technical report, Bank of International Settlements.
- BIS (2013), Basel III: The liquidity coverage ratio and liquidity risk monitoring tools, Technical report, Bank of International Settlements.
- BIS (2014), The G-SIB assessment methodology – score calculation, Technical report, Bank of International Settlements.
- BIS (2015), Revisions to the standardised approach for credit risk, Technical report, Bank for International Settlements, Basel Committee on Banking Supervision.
- BIS (2019), BIS Quarterly Review: International banking and financial market developments, Technical report, Bank of International Settlements.
- Black, F. & Scholes, M. (1973), ‘The pricing of options and corporate liabilities’, Journal of Political Economy **81**(3), 637–654.
- Blinder, A. S. (2013), ‘After the music stopped’, The Financial Crisis, the Response, and the Work Ahead. Nueva York: Penguin Books .
- BoC (2012), Understanding systemic risk in the banking sector: a macrofinancial risk assessment framework, Technical report, Bank of Canada.
- BoE (2013), A framework for stress testing the UK banking system, Technical report, Bank of England.

- BoE (2016), Stress testing the UK banking system: key elements of the 2016 stress test, Technical report, Bank of England.
- BoE (2017a), The bank of england’s approach to resolution, Technical report, Bank of England.
- BoE (2017b), Stress testing the UK banking system: 2017 results, Technical report, Bank of England.
- Bookstaber, R. (2017), The End of Theory: Financial Crises, the Failure of Economics, and the Sweep of Human Interaction, Princeton University Press, Princeton.
- Bookstaber, R., Cetina, J., Feldberg, G., Flood, M. & Glasserman, P. (2014), ‘Stress tests to promote financial stability: Assessing progress and looking to the future’, Journal of Risk Management in Financial Institutions **7**(1), 16–25.
- Bookstaber, R. & Kirman, A. (2018), Modeling a heterogeneous world, in C. Hommes & B. LeBaron, eds, ‘Handbook of Computational Economics’, Vol. 4, Elsevier, pp. 1769–795.
- Bookstaber, R. M. (2012), ‘Using agent-based models for analyzing threats to financial stability’, **3**.
- Bookstaber, R. M. & Paddrik, M. E. (2015), ‘An agent-based model for crisis liquidity dynamics’.
- Bookstaber, R., Paddrik, M. & Tivnan, B. (2014), ‘An agent-based model for financial vulnerability’, Office of Financial Research Working Paper Series **14**(05).
- Boss, M., Krenn, G., Pühr, C., Schwaiger, M. S. et al. (2007), ‘Stress testing the exposure of Austrian Banks in Central and Eastern Europe’.
- Bouchaud, J.-P. (2010), ‘Price impact’, Encyclopedia of Quantitative Finance .
- Bouchaud, J.-P., Farmer, J. D. & Lillo, F. (2008), ‘How markets slowly digest changes in supply and demand’, arXiv preprint arXiv:0809.0822 .
- Brinkhoff, J., Langfield, S. & Weeken, O. (2018), ‘From the horse’s mouth: surveying responses to stress by banks and insurers’, ESRB Occasional Paper Series (15).
- Brooke, M., Bush, O., Edwards, R., Ellis, J., Francis, B., Harimohan, R., Neiss, K. & Siegert, C. (2017), ‘Measuring the macroeconomic costs and benefits of higher UK bank capital requirements’, Financial Stability Paper (35).

- Brown, C. O. & Dinç, I. S. (2011), ‘Too many to fail? evidence of regulatory forbearance when the banking sector is weak’, The Review of Financial Studies **24**(4), 1378–1405.
- Brunnermeier, M., Crockett, A., Goodhart, C. A., Persaud, A., Shin, H. S. et al. (2009), The fundamental principles of financial regulation, Vol. 11, ICMB, International Center for Monetary and Banking Studies.
- Brunnermeier, M. K. (2008), Deciphering the liquidity and credit crunch 2007-08, Technical report, National Bureau of Economic Research.
- Brunnermeier, M. K. & Pedersen, L. H. (2009), ‘Market liquidity and funding liquidity’, Review of Financial studies **22**(6), 2201–2238.
- Brunnermeier, M. K. & Sannikov, Y. (2014), ‘A macroeconomic model with a financial sector’, The American Economic Review **104**(2), 379–421.
- Bulow, J. & Klemperer, P. (2013), ‘Market-based bank capital regulation’, Rock Center for Corporate Governance at Stanford University Working Paper (151), 13–3.
- Burrows, O., Learmonth, D., McKeown, J. & Williams, R. (2012), ‘RAMSI: a top-down stress-testing model developed at the Bank of England’.
- Burrows, O., Low, K. & Cumming, F. (2015), ‘Mapping the UK financial system’.
- Caccioli, F., Bouchaud, J.-P. & Farmer, J. D. (2012), ‘Impact-adjusted valuation and the criticality of leverage’, Risk **25**(12), 74–77.
- Caccioli, F., Catanach, T. A. & Farmer, J. D. (2012), ‘Heterogeneity, correlations and financial contagion’, Advances in Complex Systems **15**(supp02), 1250058.
- Caccioli, F., Farmer, J. D., Foti, N. & Rockmore, D. (2013), ‘How interbank lending amplifies overlapping portfolio contagion: A case study of the Austrian banking network’, arXiv preprint arXiv:1306.3704 .
- Caccioli, F., Farmer, J. D., Foti, N. & Rockmore, D. (2015), ‘Overlapping portfolios, contagion, and financial stability’, Journal of Economic Dynamics and Control **51**, 50–63.
- Caccioli, F., Shrestha, M., Moore, C. & Farmer, J. D. (2014), ‘Stability analysis of financial contagion due to overlapping portfolios’, Journal of Banking & Finance **46**, 233–245.

- Candelon, B. & Sy, A. N. (2015), How Did Markets React to Stress Tests?, International Monetary Fund Working Paper 15/75.
- Capgemini (2014), Basel III: Comparison of standardized and advanced approaches, Technical report, Capgemini.
- Capponi, A., Chen, P.-C. & Yao, D. D. (2015), ‘Liability concentration and systemic losses in financial networks’, Operations Research **64**(5), 1121–1134.
- Cecchetti, S. & Anil, K. (2018), What binds? Interactions between bank capital and liquidity regulations, in P. Hartmann, H. Huang & D. Schoenmaker, eds, ‘The Changing Fortunes of Central Banking’, Cambridge University Press, Cambridge, pp. 192–202.
- Cetina, J., Lelyveld, I. & Anand, K. (2015), ‘Making supervisory stress tests more macroprudential: Considering liquidity and solvency interactions and systemic risk’, BCBS Working Paper .
- Chen, N., Glasserman, P., Nouri, B. & Pelger, M. (2013), ‘Cocos, bail-in, and tail risk’, Office of Financial Research Working Paper (0004).
- Chennells, L. & Wingfield, V. (2015), ‘Bank failure and bail-in: an introduction’, Bank of England Quarterly Bulletin 2015 Q3 .
- Cifuentes, R., Ferrucci, G. & Shin, H. S. (2005), ‘Liquidity risk and contagion’, Journal of the European Economic Association **3**(2-3), 556–566.
- Coen, J., Lepore, C. & Schaanning, E. (2019), ‘Taking regulation seriously: Fire sales under solvency and liquidity constraints’, Bank of England working paper .
- Conlon, T. & Cotter, J. (2014), ‘Anatomy of a bail-in’, Journal of Financial Stability **15**, 257–263.
- Cont, R., Moussa, A. & Santos, E. B. (2010), Network structure and systemic risk in banking systems, in ‘Handbook of Systemic Risk’, Cambridge University Press, Cambridge, pp. 327–367.
- Cont, R. & Schaanning, E. (2019), ‘Monitoring indirect contagion’, Journal of Banking & Finance **104**, 85–102.
- Cont, R. & Schaanning, E. F. (2017), ‘Fire sales, indirect contagion and systemic stress testing’, Norges Bank Working Paper, 2017/2 .

- Cont, R. & Wagalath, L. (2013), ‘Running for the exit: distressed selling and endogenous correlation in financial markets’, Mathematical Finance **23**(4), 718–741.
- Cont, R. & Wagalath, L. (2016), ‘Fire sale forensics: measuring endogenous risks’, Mathematical Finance **26**(4), 835–866.
URL: <http://dx.doi.org/10.1111/mafi.12071>
- Copeland, A., Martin, A. & Walker, M. (2014), ‘Repo runs: Evidence from the tri-party repo market’, The Journal of Finance **69**(6), 2343–2380.
- Coval, J. & Stafford, E. (2007), ‘Asset fire sales (and purchases) in equity markets’, Journal of Financial Economics **86**(2), 479–512.
- Czech, R. & Roberts-Sklar, M. (2017), ‘Investor behaviour and reaching for yield: evidence from the sterling corporate bond market’, Staff Working Paper 685 .
- Danielsson, J. & Shin, H. S. (2003), ‘Endogenous risk’, Modern risk management: A history pp. 297–316.
- De Haan, L. & van den End, J. W. (2013), ‘Bank liquidity, the maturity ladder, and regulation’, Journal of Banking & Finance **37**(10), 3930–3950.
- De Spiegeleer, J., Schoutens, W. & Van Hulle, C. (2014), The handbook of hybrid securities: convertible bonds, CoCo bonds, and bail-in, John Wiley & Sons, Hoboken.
- Dees, S. & Henry, J. (2017), ‘Stress-test analytics for macroprudential purposes: Introducing stamp€’, ECB, STAMP€: stress-test analytics for macroprudential purposes in the Euro area pp. 13–30.
- Diamond, D. W. & Dybvig, P. H. (1983), ‘Bank runs, deposit insurance, and liquidity’, Journal of Political Economy **91**(3), 401–419.
- Dieci, R. & He, T. (2018), Heterogeneous agent models in finance, in C. Hommes & B. LeBaron, eds, ‘Handbook of Computational Economics, Volume 4, Heterogeneous Agent Models’, Elsevier, Amsterdam.
- Drehmann, M., Borio, C. E., Gambacorta, L., Jimenez, G. & Trucharte, C. (2010), ‘Countercyclical capital buffers: exploring options’.
- Duarte, F. & Eisenbach, T. M. (2015), ‘Fire-sale spillovers and systemic risk’, FRB of New York Staff Report 645 .

- E Santos, E. B., Cont, R. et al. (2010), The Brazilian interbank network structure and systemic risk, Technical report.
- EBA (2015), EBA final draft regulatory technical standards - on criteria for determining the minimum requirement for own funds and eligible liabilities under Directive 2014/59/EU, Technical report, European Banking Authority.
- EBA (2016), Final report on MREL, Technical report, European Banking Authority.
- EBA (2017a), ‘Cebs stress testing results’, <http://www.eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing/2009>.
- EBA (2017b), ‘Eu-wide stress testing’, <http://www.eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing>.
- EBA (2017c), Final guidelines on the rate of conversion, Technical report, European Banking Authority.
- Ebner, A. (2018), ‘The financial stability aspects of the EU-wide stress test’, Journal of Financial Regulation **4**(2), 326–336.
- ECB (2015), Report on financial structures, Technical report, European Central Bank.
- ECB (2017), Report on financial structures, Technical report, European Central Bank.
- Economist, T. (2018), ‘Python is becoming the world’s most popular coding language’, <https://www.economist.com/graphic-detail/2018/07/26/python-is-becoming-the-worlds-most-popular-coding-language>.
- Eichengreen, B. & Rühl, C. (2001), ‘The bail-in problem: systematic goals, ad hoc means’, Economic Systems **25**(1), 3–32.
- Eisenberg, L. & Noe, T. H. (2001), ‘Systemic risk in financial systems’, Management Science **47**(2), 236–249.
- Elliott, M., Golub, B. & Jackson, M. O. (2014), ‘Financial networks and contagion’, American economic review **104**(10), 3115–3153.
- Enriques, L., Romano, A. & Wetzler, T. (2019), ‘Network-sensitive financial regulation’, Journal of Corporation Law (forthcoming) .
- ESRB (2016), Adverse macro-financial scenario for the eba 2016 eu-wide bank stress testing exercise, Technical report, European Systemic Risk Board.

- Farmer, J. D. (2002), ‘Market force, ecology and evolution’, Industrial and Corporate Change **11**(5), 895–953.
- Farmer, J. D. (2012), Economics needs to treat the economy as a complex system, in ‘Paper for the INET Conference. Rethinking Economics and Politics’, Vol. 14.
- Farmer, J. D., Kleinnijenhuis, A. M., Nahai-Williamson, P. & Wetzer, T. (2020), ‘Foundations of system-wide stress testing’.
- Farmer, J. D. & Lo, A. W. (1999), ‘Frontiers of finance: Evolution and efficient markets’, Proceedings of the National Academy of Sciences **96**(18), 9991–9992.
- FCA (2015), Hedge fund survey, Technical report, Financial Conduct Authority.
- FED (2010), ‘AIG - goldman sachs collateral call timeline’, https://fraser.stlouisfed.org/files/docs/historical/fct/fcic/fcic_docs_aig_goldman_timeline_20100630.pdf.
- FED (2016), Dodd-frank act stress test 2016: Supervisory stress test methodology and results, Technical report, Board of Governors of the Federal Reserve.
- FED (2017a), ‘Stress tests and capital planning: Comprehensive capital analysis and review’, <https://www.federalreserve.gov/supervisionreg/ccar.htm>.
- FED (2017b), ‘Stress tests and capital planning: Dodd-frank act stress tests’, <https://www.federalreserve.gov/supervisionreg/ccar.htm>.
- Fender, I., Lewrick, U. et al. (2015), ‘Calibrating the leverage ratio’, BIS Quarterly Review pp. 43–58.
- Fique, J. (2017), ‘The macrofinancial risk assessment framework (MFRAF)’.
- Foglia, A. (2008), ‘Stress testing credit risk: a survey of authorities’ approaches’, Bank of Italy Occasional Paper **37**.
- French, K., Baily, M., Campbell, J., Cochrane, J., Diamond, D., Duffie, D., Kashyap, A., Mishkin, F., Rajan, R., Scharfstein, D. et al. (2010), ‘The Squam Lake report: fixing the financial system’, Journal of Applied Corporate Finance **22**(3), 8–21.
- FSB (2009a), The financial crisis and information gaps, Technical report, Financial Stability Board.

- FSB (2009b), Progress since the pittsburgh summit in implementing the G20 recommendations for strengthening financial stability, Technical report, Financial Stability Board.
- FSB (2013), Progress and next steps towards ending too-big-to-fail (TBTF), Technical report, Financial Stability Board.
- FSB (2015a), Global shadow banking monitoring report, Technical report.
- FSB (2015b), ‘Principles on loss-absorbing and recapitalisation capacity of G-SIBs in resolution’, Total Loss-absorbing Capacity (TLAC) Term Sheet, Basel (www.fsb.org/wp-content/uploads/TLAC-Principles-and-Term-Sheet-for-publication-final.pdf)
- FSB (2016a), Guiding principles on the temporary funding needed to support the orderly resolution of a global systemically important bank (‘G-SIB’), Technical report, Financial Stability Board.
- FSB (2016b), Standard TLAC holdings, Technical report, Financial Stability Board.
- FSB (2017), Basel III: Finalising post-crisis reforms, Technical report, Financial Stability Board.
- FSB, IMF, B. (2009c), Guidance to assess the systemic importance of financial institutions, markets and instruments: Initial considerations, Technical report, Financial Stability Board, International Monetary Fund, Bank of International Settlements.
- Gai, P., Haldane, A. & Kapadia, S. (2011), ‘Complexity, concentration and contagion’, Journal of Monetary Economics **58**(5), 453–470.
- Gai, P. & Kapadia, S. (2010), Contagion in financial networks, in ‘Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences’, The Royal Society.
- Gandy, A. & Veraart, L. A. (2017), ‘A Bayesian methodology for systemic risk assessment in financial networks’, Management Science **63**(12), 4428–4446.
- Gauthier, C., Lehar, A. & Souissi, M. (2012), ‘Macroprudential capital requirements and systemic risk’, Journal of Financial Intermediation **21**(4), 594–618.
- Gauthier, C., Souissi, M., Liu, X. et al. (2014), ‘Introducing funding liquidity risk in a macro stress-testing framework’, International Journal of Central Banking **10**(4), 105–141.

- Geithner, T. F. (2014), ‘Stress test: Reflections on the financial crisis’, Business Economics **49**(3), 201–203.
- Georg, C.-P. (2013), ‘The effect of the interbank network structure on contagion and common shocks’, Journal of Banking & Finance **37**(7), 2216–2228.
- Glasserman, P., Tangirala, G. et al. (2015), Are the Federal Reserve’s stress test results predictable?, Technical report.
- Glasserman, P. & Young, H. P. (2015), ‘How likely is contagion in financial networks?’, Journal of Banking & Finance **50**, 383–399.
- Goodhart, C. (2013), ‘Ratio controls need reconsideration’, Journal of Financial Stability **9**(3), 445–450.
- Goodhart, C. & Avgouleas, E. (2016), Critical reflections on bank bail-ins, in ‘The Reform of International Economic Governance’, Routledge, pp. 75–99.
- Goodhart, C., Farmer, J. D. & Kleinnijenhuis, A. M. (2020), ‘Systemic implications of the bail-in design’.
- Goodhart, C. et al. (2008), ‘Liquidity risk management’, Banque de France Financial Stability Review **11**, 39–44.
- Gorton, G. B. & Metrick, A. (2009), Haircuts, Technical report, National Bureau of Economic Research.
- Gorton, G. & Metrick, A. (2012), ‘Securitized banking and the run on repo’, Journal of Financial economics **104**(3), 425–451.
- Gorton, G. & Muir, T. (2016), Mobile collateral versus immobile collateral, Technical report, National Bureau of Economic Research.
- Goyal, S. (2018), Economic and social networks, in C. Hommes & B. LeBaron, eds, ‘Handbook of Computational Economics, Volume 4, Heterogeneous Agent Models’, Elsevier.
- Greenwood, R., Landier, A. & Thesmar, D. (2015), ‘Vulnerable banks’, Journal of Financial Economics **115**(3), 471–485.
- Greenwood, R., Stein, J. C., Hanson, S. G. & Sunderam, A. (2017), ‘Strengthening and streamlining bank capital regulation’, Brookings Papers on Economic Activity **2017**(2), 479–565.

- Groendahl, B. (2015), ‘Leverage ratio for banks can rise as high as 5%, BIS says’, <https://www.bloomberg.com/news/articles/2015-12-06/leverage-ratio-for-banks-can-be-raised-as-high-as-5-bis-says>.
- Halaj, G. (2018), ‘Agent-based model of system-wide implications of funding risk’.
- Halaj, G. & Kok, C. (2013), ‘Assessing interbank contagion using simulated networks’, Computational Management Science **10**(2-3), 157–186.
- He, Z. & Xiong, W. (2012), ‘Dynamic debt runs’, Review of Financial Studies **25**(6), 1799–1843.
- Heider, F., Hoerova, M. & Holthausen, C. (2009), ‘Liquidity hoarding and interbank market spreads: The role of counterparty risk’.
- Hirtle, B., Kovner, A. & Zeller, S. (2016), ‘Are stress tests still informative’, <http://libertystreeteconomics.newyorkfed.org/2016/04/are-stress-tests-still-informative.html>.
- Hommes, C. H. (2006), ‘Heterogeneous agent models in economics and finance’, Handbook of computational economics **2**, 1109–1186.
- Hüser, A.-C., Halaj, G., Kok, C., Perales, C. & van der Kraaij, A. (2017), ‘The systemic implications of bail-in: a multi-layered network approach’, Journal of Financial Stability **38**, 81–97.
- Hüser, A.-C. & Kok, C. (2019), ‘Mapping bank securities across euro area sectors: comparing funding and exposure networks’.
- IMF (2014), Canada financial sector stability assessment, Technical report, International Monetary Fund.
- IMF (2019), ‘Financial sector assessment program (fsap)’, <https://www.imf.org/external/np/fsap/fssa.aspx>.
- Iori, G. & Mantegna, R. (2018), Complex financial networks, in C. Hommes & B. LeBaron, eds, ‘Handbook of Computational Economics, Volume 4, Heterogeneous Agent Models’, Elsevier, Amsterdam.
- IOSCO (2017), Report on the fourth IOSCO hedge funds survey, Technical report, International Organization of Securities Commissions.

- Iyer, R. & Peydro, J.-L. (2011), ‘Interbank contagion at work: Evidence from a natural experiment’, Review of Financial Studies **24**(4), 1337–1377.
- Judge, K. & Berner, R. (2019), ‘The data standardization challenge’, Columbia Law School Scholarship Repository, Working Paper .
- Kapadia, S., Drehmann, M., Elliott, J. & Sterne, G. (2013), Liquidity risk, cash flow constraints, and systemic feedbacks, in ‘Quantifying Systemic Risk’, University of Chicago Press, pp. 29–61.
- Khandani, A. E. & Lo, A. W. (2011), ‘What happened to the quants in august 2007? Evidence from factors and transactions data’, Journal of Financial Markets **14**(1), 1–46.
- Kiyotaki, N. & Moore, J. (1997), ‘Credit cycles’, Journal of political economy **105**(2), 211–248.
- Klimek, P., Poledna, S., Farmer, J. D. & Thurner, S. (2015), ‘To bail-out or to bail-in? Answers from an agent-based model’, Journal of Economic Dynamics and Control **50**, 144–154.
- Kok, C. & Montagna, M. (2013), Multi-layered interbank model for assessing systemic risk, Technical report, European Central Bank.
- Kyle, A. S. (1985), ‘Continuous auctions and insider trading’, Econometrica: Journal of the Econometric Society **53**(6), 1315–1335.
- LeBaron, B. (2006), ‘Agent-based computational finance’, Handbook of computational economics **2**, 1187–1233.
- Lo, A. W. (2017), Adaptive markets: Financial evolution at the speed of thought, Princeton University Press, Princeton, Princeton.
- Martin, A., Skeie, D. & Von Thadden, E.-L. (2014), ‘Repo runs’, Review of Financial Studies **27**(4), 957–989.
- May, R. M. & Arinaminpathy, N. (2010), ‘Systemic risk: the dynamics of model banking systems’, Journal of the Royal Society Interface **7**(46), 823–838.
- McAndrews, J., Morgan, D. P., Santos, J. A. & Yorulmazer, T. (2014), ‘What makes large bank failures so messy and what to do about it?’, pp. 229–244.
- Mehrling, P., Pozsar, Z., Sweeney, J. & Neilson, D. H. (2013), ‘Bagehot was a shadow banker: shadow banking, central banking, and the future of global finance’, Central Banking, and the Future of Global Finance .

- Merton, R. C. (1974), ‘On the pricing of corporate debt: The risk structure of interest rates’, Journal of Finance **29**(2), 449–470.
- Merton, R. C. (1976), ‘Option pricing when underlying stock returns are discontinuous’, Journal of Financial Economics **3**(1-2), 125–144.
- Moretti, M., Stolz, S. M. & Swinburne, M. (2008), Stress Testing at the IMF, IMF Working Paper WP/08/206.
- Morris, S. & Shin, H. S. (2001), ‘Global games: theory and applications’.
- Noonan, L., Binham, C. & Shotter, J. (2016), ‘Deutsche bank received special treatment in eu stress tests’, <https://www.ft.com/content/44768ea8-8c71-11e6-8aa5-f79f5696c731>.
- OFR (2014), Office of financial research: 2014 annual report, Technical report, Office of Financial Research.
- Ong, L. L. & Pazarbasioglu, C. (2014), ‘Credibility and crisis stress testing’, International Journal of Financial Studies **2**(1), 15–81.
- Paddrik, M. E. & Young, P. (2017), ‘How safe are central counterparties in derivatives markets?’, LSE Research Online .
- Paddrik, M., Rajan, S. & Young, H. P. (2016), ‘Contagion in the CDS market’, Office of Financial Research Working Paper pp. 16–12.
- Paulin, J., Calinescu, A. & Wooldridge, M. (2018), ‘Understanding flash crash contagion and systemic risk: A micro-macro agent-based approach’, arXiv preprint arXiv:1805.08454 .
- Pennacchi, G. (2010), ‘A structural model of contingent bank capital’, FRB of Cleveland Working Paper 10-04 .
- Persaud, A. (2014), ‘Why bail-in securities are fool’s gold’, Available at SSRN 2643702 .
- Pistor, K. (2013), ‘A legal theory of finance’, Journal of Comparative Economics **41**(2), 315–330.
- Poledna, S., Molina-Borboa, J. L., Martínez-Jaramillo, S., Van Der Leij, M. & Thurner, S. (2015), ‘The multi-layer network nature of systemic risk and its implications for the costs of financial crises’, Journal of Financial Stability **20**, 70–81.

- Pozsar, Z. (2013), ‘Shadow banking and the global financial ecosystem’, VoxEU, November **6**.
- Pozsar, Z., Adrian, T., Ashcraft, A. B. & Boesky, H. (2010), ‘Shadow banking’.
- Pozsar, Z. & Singh, M. (2011), ‘The nonbank-bank nexus and the shadow banking system’, IMF Working Paper 11/289 .
- Prudential Regulation Authority (2017), Constulation paper CP15/17: The minimum requirement for own funds and eligible liabilities (MREL) – buffers, Technical report, Bank of England.
- Quagliariello, M. (2009), Stress-testing the Banking System: Methodologies and Applications, Cambridge University Press.
- Rachev, S. T., Menn, C. & Fabozzi, F. J. (2005), Fat-tailed and skewed asset return distributions: implications for risk management, portfolio selection, and option pricing, Vol. 139, John Wiley & Sons, Hoboken.
- Rogers, L. C. & Veraart, L. A. (2013), ‘Failure and rescue in an interbank network’, Management Science **59**(4), 882–898.
- Rutledge, V., Moore, M., Dobler, M., Bossu, W., Jassaud, N. & Zhou, J. (2012), ‘From bail-out to bail-in: mandatory debt restructuring of systemic financial institutions’.
- Sarin, N. & Summers, L. H. (2016), ‘Have big banks gotten safer?’, Brookings Papers .
- Sarin, N. & Summers, L. H. (2020, forthcoming in Handbook of Financial Stress Testing), On Market-Based Approaches to the Valuation of Capital, Cambridge University Press.
- Schäfer, A., Schnabel, I. & Weder, B. (2016), ‘Bail-in expectations for European banks: Actions speak louder than words’.
- Scheule, H. & Roesh, D. (2008), Stress Testing for Financial Institutions: Applications, regulations and techniques, Risk Books Publications, London.
- Schnabel, I. & Shin, H. S. (2004), ‘Liquidity and contagion: The crisis of 1763’, Journal of the European Economic Association **2**(6), 929–968.
- Schuermann, T. (2014), ‘Stress testing banks’, International Journal of Forecasting **30**(3), 717–728.
- Shleifer, A. & Vishny, R. (2011), ‘Fire sales in finance and macroeconomics’, Journal of Economic Perspectives **25**(1), 29–48.

- Shleifer, A. & Vishny, R. W. (1992), ‘Liquidation values and debt capacity: A market equilibrium approach’, The Journal of Finance **47**(4), 1343–1366.
- Siddique, A. & Hasan, I. (2012), Stress Testing: Approaches, Methods and Applications, Risk Books, London.
- Sommer, J. H. (2014), ‘Why bail-in? and how!’, Economic Policy Review, Forthcoming .
- SRB (2017), Minimum requirement for own funds and eligible liabilities (mrel), Technical report, Single Resolution Board.
- Tarullo, D. K. (2014), Stress testing after five years: a speech at the Federal Reserve third annual stress test modeling symposium, Boston, Massachusetts, June 25, 2014, Technical report, Board of Governors of the Federal Reserve System (US).
- Tarullo, D. K. (2016), Next steps in the evolution of stress testing, in ‘Remarks at the Yale University School of Management Leaders Forum (September 26, 2016)’.
- Turner, S. & Poledna, S. (2013), ‘Debt-tranparency: Controlling systemic risk in financial networks’, arXiv preprint arXiv:1301.6115 .
- Upper, C. & Worms, A. (2004), ‘Estimating bilateral exposures in the German interbank market: Is there a danger of contagion?’, European Economic Review **48**(4), 827–849.
- Wackerbeck, P., Crijns, J. & Karsten, C. (2016), ‘Stress testing: From regulatory burden to strategic capability’, <https://www.strategyand.pwc.com/reports/stress-testing>.
- WBG (2017), Understanding bank recovery and resolution in the EU - A guidebook to the BRRD, Technical report, World Bank Group.
- Wiersema, G., Kleinnijenhuis, A. M., Wetzter, T. & Farmer, J. D. (2019), ‘Scenario-free analysis of financial stability with interacting contagion channels’, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3408533 .
- Zenios, S. A. (2016), ‘Fairness and reflexivity in the cyprus bail-in’, Empirica **43**(3), 579–606.