



BANK OF ENGLAND

Staff Working Paper No. 861

Foundations of system-wide financial stress testing with heterogeneous institutions

J Doyne Farmer, Alissa M Kleinnijenhuis,
Paul Nahai-Williamson and Thom Wetzer

May 2020

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee.



BANK OF ENGLAND

Staff Working Paper No. 861

Foundations of system-wide financial stress testing with heterogeneous institutions

J Doyne Farmer,⁽¹⁾ Alissa M Kleinnijenhuis,⁽²⁾ Paul Nahai-Williamson⁽³⁾ and Thom Wetzer⁽⁴⁾

Abstract

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.

Using this framework, we implement a system-wide stress test for the European financial system. We obtain three key findings. First, the financial system may be stable or unstable for a given microprudential stress test outcome, depending on the system's shock-amplifying tendency. Second, the 'usability' of banks' capital buffers (the willingness of banks to use buffers to absorb losses) is of great consequence to systemic resilience. Third, there is a risk that the size of capital buffers needed to limit systemic risk could be severely underestimated if calibrated in the absence of system-wide approaches.

Key words: Systemic risk, stress testing, financial contagion, financial institutions, capital requirements, macroprudential policy.

JEL classification: G17, G21, G23, G28, C63.

(1) Institute for New Economic Thinking, University of Oxford, UK. Email: doyne.farmer@inet.ox.ac.uk

(2) Institute for New Economic Thinking, University of Oxford, UK and MIT Sloan School of Management, Massachusetts Institute of Technology, US. Email: alissa.kleinnijenhuis@maths.ox.ac.uk

(3) Bank of England. Email: paul.nahai-williamson@bankofengland.co.uk

(4) Institute for New Economic Thinking and Faculty of Law, University of Oxford, UK. Email: thom.wetzer@law.ox.ac.uk

We thank David Aikman, Mustafa G. Aydogan, Christoph Aymanns, Adam Brinley Codd, Stephen Burgess, 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. Rudi Tahin provided outstanding research assistance. We dedicate this paper to Rafa Baptista (in memoriam), who made invaluable contributions at the early stages of this project. The usual disclaimers apply.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH

Email enquiries@bankofengland.co.uk

© Bank of England 2020

ISSN 1749-9135 (on-line)

Contents

1	Introduction	1
2	Relevant Literature	4
2.1	Modelling System-Wide Stress Dynamics	4
2.2	Stress Tests and Prudential Regulation	6
3	A Structural Framework for System-Wide Stress Tests with Heterogeneous Institutions	7
3.1	Five Building Blocks to Represent Financial Systems	9
3.2	Contagion and Amplification	13
3.3	Key Design Choices for Simulation Software	13
4	A System-Wide Stress Test Model for the European Financial System	15
4.1	Financial Institutions and their Constraints	15
4.2	Financial Contracts, Markets & Contagion Mechanisms	20
4.3	Behaviour	25
5	Policy experiments and Results	30
5.1	Default Parameters and Visualisation	31
5.2	From Micro to Macro: A Macroprudential Overlay for the EBA 2018 Stress Test	33
5.3	Contagious Feedback Loops Between Banks and Non-Banks	36
5.4	Amplification of Contagion Mechanisms	38
5.5	‘Usability’ of Buffers and Contagion	40
5.6	Calibration of Buffers with System-Wide Stress Tests	41
6	Policy Implications and Conclusion	43

A	Technical Appendix: Foundation of System-Wide Stress Testing	54
A.1	Technicalities of Model Implementation	54
A.2	System-Wide Stress Testing Software	67
A.3	Supplementary Figures and Tables	72

1 Introduction

In a highly connected financial system, seemingly localised 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 and 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 understand the various financial institutions involved, their interconnections, and their interactions under stress. This challenge is compounded by the fact that the financial system constantly changes, not least in response to new regulation. Critically, as conditions change and financial market participants respond to the new regulations, the interacting effects of the suite of regulatory reforms will become apparent, and new risks are likely to emerge. The challenge for regulators is to constantly evaluate these risks to the resilience of individual institutions and the financial system as a whole.

Few regulatory instruments embody this challenge as clearly as financial stress tests (Aymanns et al. (2018)). In 2009, the US Federal Reserve’s Supervisory Capital Assessment Program (SCAP) led the way to what ultimately resulted in a major post-crisis increase in regulatory stress testing. Microprudential stress tests like the SCAP focus on modelling the first order impact of a defined macroeconomic scenario on banks’ balance sheets, and do not consider the potential for banks’ actions in the stress to cause spillovers or to amplify shocks. They can provide useful information to inform systemic risk analysis, both on institutions’ behaviour in a stress, and on their balance sheet exposures. This is not limited to banks, and regulators around the world have seized on that to subject non-banks - including central clearing parties, insurers, and pension funds - to microprudential stress tests. Microprudential stress tests have been credited not only with restoring confidence in the financial system during the financial crisis, but also with enabling its successful recapitalisation (Bernanke (2013)).

Nevertheless, their narrow focus on the resilience of individual institutions to pre-defined shocks implies that microprudential stress tests are not designed to capture systemic risk. Simply put, “the system is not the sum of its parts” (Brazier (2017)); to understand how shocks can propagate and amplify, regulators need stress tests that cap-

ture the effect of endogenous shock amplification. This realisation has motivated a push to develop macroprudential stress tests – including by incorporating amplification risks into existing regulatory stress testing frameworks, as at the Bank of England¹. At first, such efforts focused on the interaction between similar institutions, primarily banks (e.g. [Bookstaber, Cetina, Feldberg, Flood and Glasserman \(2014\)](#), [Cetina et al. \(2015\)](#)), but researchers and regulators increasingly recognise that this is not enough: as the example above illustrates, endogenous shock amplification arises not only from the interaction between banks, but also involves interactions with non-banks (e.g. [Anderson et al. \(2018\)](#)). Credit intermediation channels, for example, run through the non-banking sector, and non-bank financial institutions are generally strongly interconnected with banks (see e.g. [Brazier \(2017\)](#)). In tandem with developments in macroprudential stress tests for banks, regulators therefore increasingly look for ways to develop macroprudential stress tests that are truly system-wide.

Macroprudential stress testing of the wider financial system is still in its infancy ([Aymanns et al. \(2018\)](#), [Anderson et al. \(2018\)](#)). Pioneering work in this field has focused on particular markets and the interactions among representative sectors ([Aikman et al. \(2019\)](#)), representing an important but incomplete advance in the understanding of system-wide resilience. The main obstacles to further progress are associated with the difficulty in designing a modelling framework that (1) comprehensively captures amplification of solvency and liquidity shocks and (2) takes account of the heterogeneity of institutions and their responses to these shocks given the constraints they face ([Anderson et al. \(2018\)](#)). Even if this modelling challenge is solved, scaling and flexibly adjusting such system-wide models in data-driven ways to account for the characteristics of different (subsets of) financial systems presents novel computational challenges that microprudential stress tests avoid.

In this paper, we address those challenges and propose a structural 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 Euro-

¹See [Bank of England \(2015a\)](#) and [Churm and Nahai-Williamson \(2019\)](#) for more details on the Bank of England’s approach to incorporating amplification risks in its concurrent stress testing exercise.

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

pean financial system. This stress test model captures solvency and liquidity channels, 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 overlay’ on top of the regular micro-prudential European Banking Authority (EBA) stress test from 2018 and compare the stress test results.⁵

This comparison yields three main findings, which are robust to extensive sensitivity and robustness checks.⁶ First, depending on the shock-amplifying tendency of the financial system, the system may be stable or unstable for a given microprudential stress test outcome.⁷ 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 set of initial shocks may be endogenously amplified to starkly different degrees depending on the characteristics of different financial systems.

We show that the outcome of a system-wide stress test depends on which (interacting) contagion channels are taken into account. We confirm the result that interacting contagion channels can produce significantly higher rates of bank failure (by as much as a factor of 5) than the sum of failures when each acts in isolation (Caccioli et al. (2013), Kok and Montagna (2013), Poledna et al. (2015), Hüser and Kok (2019), Wiersema et al. (2019)). Our model can serve as a tool to evaluate which set of amplification mechanisms is most destabilising under different conditions. 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, contagion via asset sales becomes more dominant and amplifies other channels. We also show that including heterogeneous financial institutions, and in particular non-banks, changes the magnitude of

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, investment funds, and hedge funds - and allow for heterogeneity within these classes.

⁵We do not model the potential for banks to respond to stress by reducing their lending to the real economy. In part, this reflects the fact that financial market dynamics are likely to operate over faster timescales than changes to real economy lending behaviour. It is also consistent with the fact that regulatory banking sector concurrent stress tests generally either enforce static balance sheets; or, as in the case of the Bank of England’s stress test, require that banks continue to meet credit demand from the real economy in the scenario (see e.g. Bank of England (2015a)). Calibrating capital buffers based on this approach – and releasing them if a stress materialises – means that capital buffers should not constrain banks’ ability to lend in a real economic downturn. In line with this approach, if capital buffers were sized taking into account financial market amplification risks, as proposed in this paper, those buffers should be sufficient both to address contagion risks and to support banks’ continued lending to the real economy.

⁶Our 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.

⁷Exogenous shocks are amplified if the systemic risk measure including endogenous shocks is higher than without. In line with Gai and Kapadia (2010), Gai et al. (2011), Paulin et al. (2018), we use as our systemic risk measure the average fraction of defaults in a systemic event – one in which at least 5% of the banking system defaults.

systemic risk.

Our second finding is that banks’ willingness to draw on their capital buffers to absorb losses - the ‘usability’ of capital buffers - significantly affects the shock-amplifying tendency of a financial system. If banks take actions to avoid using their buffers in response to an adverse shock, this can generate pro-cyclical dynamics that substantially increase system-wide losses. In light of this result, regulators should be mindful of how the design and enforcement of regulatory buffers may affect their ‘usability’ in times of financial stress ([Goodhart et al. \(2008\)](#), [Goodhart \(2013\)](#)).

Finally, we find that microprudential stress tests that omit endogenous amplification mechanisms 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 the results of banking sector stress tests to calibrate the discretionary time-varying capital requirements under Pillar II of the Basel capital adequacy framework and, in the United Kingdom, to inform the calibration of the countercyclical capital buffer.⁸ Our findings suggest that system-wide stress tests can meaningfully complement microprudential stress tests when calibrating capital buffers.

The paper proceeds as follows. Section 2 specifies our contribution to the literature. Section 3 sets out the foundations of our structural framework for system-wide stress tests, and in Section 4 we use this framework to develop the model for the system-wide stress test of the European financial system. Section 5 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 6.

2 Relevant Literature

2.1 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. \(2012\)](#), [Fique \(2017\)](#), [Kok and Montagna \(2013\)](#), [Dees and Henry \(2017\)](#), [Aikman et al. \(2019\)](#)), and have been joined by academics (e.g. [Cont and Schaanning \(2017\)](#)). However important these contributions may be, they do not propose a structural approach to modelling system-wide dynamics. On that front, we make three contributions.

⁸As discussed above, the Bank of England’s concurrent stress test results already incorporate the risks associated with several potential amplification mechanisms, including solvency contagion risk, the risk of spillovers due to sales of commonly held tradable assets, and the interaction between deteriorating solvency and increasing funding costs [Bank of England \(2017b\)](#).

First, we outline a structural framework that allows for the systematic modelling of interacting contagion mechanisms. We distinguish two modes of contagion - node and contract amplification (see Section 3.2) - which we use to study four specific contagion mechanisms (Section 4.2.4): 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 and Montagna \(2013\)](#) consider the first three contagion mechanisms, [Caccioli et al. \(2013\)](#) and [Hüser and Kok \(2019\)](#) explore the first two, [Poledna et al. \(2015\)](#) investigate funding contagion for different contracts, and [Brunnermeier and 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 and 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 structural 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 3.1.1, we do so by characterising each institution on the basis of: the contracts (or, if the data are 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 structural 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 and 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 and 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 (see e.g. [Luna and Hardy \(2019\)](#)), the importance of this modeling innovation to regulators is likely to grow. This is especially true because the capacity of existing models to capture heterogeneity remains limited ([Halaj \(2018\)](#), [Baranova et al. \(2017\)](#), [Aikman et al. \(2019\)](#)).

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 and Eisenbach \(2015\)](#), [Aymanns et al. \(2016\)](#), [Caccioli et al. \(2014\)](#)) and, moreover, that financial institutions face an increasingly complex array of interacting and overlapping regulatory constraints (see e.g. [Armour et al. \(2016\)](#)). Despite that reality, existing con-

tagion models typically model either the leverage ratio⁹ or the risk-weighted capital ratio (see e.g. [Kok and Montagna \(2013\)](#), [Cifuentes et al. \(2005\)](#)), and they rarely implement the Basel III liquidity constraints (see e.g. [De Haan and 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.

In our structural 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 4.3.1). This approach is consistent with the empirical findings of [Adrian and Shin \(2010\)](#) and has intuitive appeal. Consistent with the findings of [Coen et al. \(2019\)](#) on optimal asset selling behaviour under different binding constraints, 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 and paying off liabilities can reduce the leverage ratio but will not improve the risk-weighted capital ratio.

2.2 Stress Tests and Prudential Regulation

The development of a structural framework allows us to develop a system-wide stress test of the European financial system.¹⁰ 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 micro-prudential 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 and Schaanning \(2017\)](#)). Moreover, various authors have applied a ‘macroprudential overlay’ to regulatory stress tests (see e.g. [Burrows et al. \(2012\)](#),

⁹Most consider leverage targeting (e.g. [Greenwood et al. \(2015\)](#), [Duarte and Eisenbach \(2015\)](#)); and some consider the case of a distinct buffer value and target value for the leverage ratio constraint, in which once the buffer is breached, banks deleverage to return above it to the target value ([Cont and Schaanning \(2017\)](#), [Bookstaber, Paddrik and Tivnan \(2014\)](#)).

¹⁰We stress that this model implementation by no means exhausts the options offered by the structural framework; rather, it showcases the possibilities of the framework for system-wide stress testing and studying policy questions.

Dees and Henry (2017), Paddrik et al. (2016), Paddrik and 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 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 a factor of five) than suggested by the sum of failures when they act in isolation. This suggests that purely microprudential stress tests alone will overstate resilience in at least some scenarios and may 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 extended strict capital *requirements* with *buffers* – which institutions can (temporarily) draw on without breaching their regulatory obligations – so that institutions can absorb shocks and refrain from taking procyclical 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 explicitly consider how usability of capital (or liquidity) would affect resilience (e.g. Battiston, Puliga, Kaushik, Tasca and Caldarelli (2012), Greenwood et al. (2015), Cont and Schaanning (2017), Duarte and Eisenbach (2015)). Aymanns and Farmer (2015) show that higher capital requirements may be *destabilising*. Using the ability of the structural 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 is likely to be underestimated if stress tests that omit such mechanisms are used for calibration.

3 A Structural Framework for System-Wide Stress Tests with Heterogeneous Institutions

In this section, we outline a structural 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*

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.	

Table 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.

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 each contractual type, that represents the financial system. Studying the topology of this network can already yield valuable insights about systemic risk (see e.g. [Battiston, Gatti, Gallegati, Greenwald and Stiglitz \(2012\)](#)), 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 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, which can host large-scale data-driven models.

Section 3.1 outlines the five building blocks used to represent financial systems in greater detail. In Section 3.2, we discuss the endogenous (amplifying) dynamics that can arise in this structural framework, and how we conceptualise them. We conclude by highlighting some important design principles of the software that we have developed to host these stress-test models in Section 3.3.

3.1 Five Building Blocks to Represent Financial Systems

Our structural 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 and Shin \(2003\)](#). We discuss these five building blocks in turn.

3.1.1 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 structural 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 function¹¹, and (3) the set of (contingent) *liquidity* obligations, including the contract’s ‘liquidity function’ and its inputs.¹²

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 they hold. 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 accounting¹³ 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.¹⁴

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.

¹¹This information tells us how the price of the contract is determined, the contractual maturity, and whether a contract is secured or not, etc.

¹²For 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 4.2.1). 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 15).

¹³Our structural framework supports various accounting standards, which are made available in the online repository of the Economic Simulation Library.

¹⁴In Section 3.2 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.

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 are not available, the 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.1.2.2](#).

3.1.2 Markets Contracts are formed in financial markets, and it is there that their price is determined by interacting market participants. Different types of contracts are traded in 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 and 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.¹⁵ In our structural 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.

3.1.3 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, regulatory¹⁶, market¹⁷, or contractual¹⁸ constraints, or their internal risk limits¹⁹. These rules and constraints limit the time-dependent set of actions available to each financial institution. They could include rules for operating under normal conditions – 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.

Importantly, these constraints can act both to trigger action during balance sheet

¹⁵The 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. \(2017\)](#)).

¹⁶Examples of regulatory constraints include minimum leverage and risk-based capital ratios for banks

¹⁷Market-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 and Tivnan \(2014\)](#)). Such limits may be stricter than those imposed by regulators.

¹⁸Contractual constraints include obligations to exchange margin or to repay liabilities at maturity.

¹⁹Internal-risk limits are institution-specific limits, such as value-at-risk (VaR) limits ([Berkowitz and O’Brien \(2002\)](#)), which are typically set by the risk managers of the institution.

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 a number of (prudential) regulatory constraints that are of no relevance to other institution types such as 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-dependent²⁰, which in turn affects firm-specific behaviour and ultimately system-level dynamics.

3.1.4 Behaviour Behaviour is central to understanding systemic risk; it is also the most challenging aspect to pin down and model (Farmer and Lo (1999), Farmer (2002), Lo (2017), Aymanns et al. (2018)). To fully capture the build-up and crystallisation of systemic risk, an ultimate understanding of behaviour under both normal and stress conditions is important. Understanding how agents optimise their balance sheets subject to the constraints described above to meet their business targets (e.g., 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 be forced into when under severe stress is key to modelling the dynamics that could occur when risks start to crystallise.

Behaviour in our structural 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 focus in particular on how the constraints that institutions face can limit their options and can force certain behaviours. These behaviours will be institution specific, but generally speaking, they will relate to issues of solvency/profitability and/or liquidity.

²⁰Regulators and contractual counterparties may, for example, loosen constraints during crisis when they fear that rigid enforcement might lead to default, see e.g. Awrey (2019), Pistor (2013). We do not consider such dynamics here.

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, they 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. In addition to this, 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 face²¹; 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 built using our structural framework are, however, explicitly conditional on the behavioural assumptions made. We therefore set up our framework 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 example about the types of behaviour that are most destabilising and should therefore be avoided.

²¹That does not mean that these assumptions are completely uninformed: market surveys, for example, can be helpful, as can empirical analysis of institutional trading patterns where such data exist. Detailed modelling of specific markets can also be used to identify the sets of institutional behaviour that produce the observed dynamics, see for example [Braun-Munzinger et al. \(2016\)](#).

3.2 Contagion and Amplification

A financial system modelled 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 or contagion via 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 or contagion, 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, mortgage loan defaults can result in valuation shocks to mortgage-backed securities; and valuation shocks to tradable assets can lead to margin calls (liquidity shocks) on repurchase agreements. We include the latter mechanism in our implementation of a system-wide stress test of the European financial system (Section 4.2.4).

3.3 Key Design Choices for Simulation Software

To operationalise the approach outlined above, we have developed an object-oriented modelling software (a simulation engine with an 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). Unless otherwise stated, by embedding the structural 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 and Berner \(2019\)](#)) to run advanced, data-driven system-wide stress testing models at scale. The fully documented simulation engine, and its accompanying software library, is available at <https://github.com/ox-inet-resilience/resilience>.

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 structural 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 e.g. Figue (2017)). Although this takes away systematic biases, it does not prevent the simulation from featuring 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, i.e., 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 the sequence of the acting institutions in each time step.²⁴ We have included more information about the mailbox system in Appendix A.2.2.

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 re-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 timestep. The practical effect would be that this institution responds more slowly to market developments. We do not explore this option here.

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>.

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

Using our structural 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 structural framework. We stress, however, that this framework can be, and is indeed 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, investment funds and hedge funds. Because we take an existing banking sector stress test as our starting point, and have relatively good institution-level data for banks, we model them at the institution-level. We then add investment funds and hedge funds to our model in a more stylised way, due to data limitations. The banks in our model are directly connected via unsecured interbank lending and borrowing, and are indirectly connected via common holdings of tradable assets. Investment funds and hedge funds also hold these tradable assets, and hedge funds are also directly connected with banks via repo funding. In this section, we discuss the setup of the model by outlining how we 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, and in Appendix A.1.4.2 we describe the behavioural assumptions in greater detail. We include a table of notation in Appendix A.1.1.

4.1 Financial Institutions and their Constraints

4.1.1 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*.²⁵ The stylised balance sheet of banks $i \in \mathcal{B}$, where \mathcal{B} is the set of banks, is depicted in Figure 1.

²⁵Due to data limitations (e.g. missing fields), we exclude a handful of banks and end up with a total of 42 banks.

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

Figure 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.

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

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 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.²⁶ 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.²⁷

Banks are required to meet the minimum CET1 capital ratios of 4.5% and minimum leverage ratio requirements of 3%²⁸ at all times:

²⁶We 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.

²⁷When 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. See Appendix A.1.4.2 for details.

²⁸We note that UK banks must meet a leverage ratio of 3.25%, with the leverage exposure measure

$$\rho_i := \frac{\tilde{E}_i}{\Omega_i} = \frac{\tilde{E}_i}{\sum_{p \in \mathcal{P}} \omega_p A_{ip}} \geq \rho^M = 4.5\%. \quad (1)$$

and

$$\lambda_i := \frac{\tilde{E}_i^{T1}}{\hat{A}_i} \geq \lambda^M = 3\%. \quad (2)$$

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 and Montagna (2013), we assume that banks that breach minimum capital requirements will fail and are either liquidated or resolved (see Section 4.3 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)). 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,²⁹ 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

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

²⁹See 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.”

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)).

Given this complication, we investigate the impact that banks’ willingness to use these combined buffers (see Section 4.3.2 and 5.5), 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

$$\rho_i^{CB} := \rho_i^{CCoB} + \rho_i^{CCyB} + \max\{\rho_i^{GSIB}, \rho_i^{DSIB}, \rho_i^{SR}\}; \quad (3)$$

$$\lambda_i^{CB} := \frac{1}{2}\rho_i^{G-SIB}. \quad (4)$$

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 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 authorities 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,³⁰ the domestically important bank surcharge (D-SIB) $\rho_i^{D-SIB} \in [0\%, 2\%]$,³¹ and the systemic risk buffer $\rho_i^{SR} \geq 1\%$.³² 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 .³³ 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 and 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 equation 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.1.2.4). In normal times, a bank is expected to have an LCR Λ_i that complies with the LCR buffer standards $\Lambda^S = 100\%$ (BIS (2013)):

$$\Lambda_i := \frac{Q_i}{\Theta_i} = \frac{Q_i}{\Theta_i^O - \min\{\Theta_i^I, 0.75 \cdot \Theta_i^O\}} \geq \Lambda^S = 100\%. \quad (5)$$

4.1.2 Investment Funds We extend our model to include four representative investment funds - 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 investment fund $i \in \mathcal{M}$, where \mathcal{M} is the set of investment funds, is shown in Figure 2a. The assets A_i of an investment fund 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) investment funds face is that they must

³⁰The 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 $\rho_i^{G-SIB} = 0$.

³¹The 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)).

³²The 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.

³³Banks are also subject to a Net Stable Funding Ratio (NSFR). Because Cecchetti and 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.

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

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

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 2: Balance sheets of non-banks.

fulfil redemption requests from investors. Empirical evidence shows that investment funds 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 and Stafford (2007), Baranova et al. (2017)). In line with the empirical evidence of Coval and Stafford (2007), we assume that the investment fund investors redeem shares proportional to the relative loss of their NAV in our simulations. Investment funds have an obligation to pay back these shares at their prevailing NAV - we set out how they do so in section 4.3.

4.1.3 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%³⁴ (based on the FCA (2015) survey).³⁵ A hedge fund's asset holdings are calibrated to data from the *ECB Statistical Data Warehouse* (see initialisation details in Appendix A.1.2.1). The stylised balance sheet of a hedge fund is displayed in Figure 2b. 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.

4.2 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.

³⁴Or, equivalently, 2.3 if leverage is defined as assets A_i over book equity E_i .

³⁵As explained in more detail in Appendix A.1.2.1, we do not consider synthetic leverage (that attained by derivatives, for instance).

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.³⁶

As discussed in Section 3, the structural 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 scenario, we take an approach consistent with the relevant literature (see e.g. Caccioli et al. (2013), Kok and 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.³⁷

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

$$T_i = \sum_{a \in \mathcal{A}} T_{ia} = \sum_{a \in \mathcal{A}} \sum_{m=1}^{M_a} T_{iam} = \sum_{a \in \mathcal{A}} \sum_{m=1}^{M_a} s_{iam} p_{am}, \quad (6)$$

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} .

³⁶We 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 Bank of England (2018a)), suggesting that including this missing channel would not significantly influence 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 would 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) would be likely to be limited.

³⁷We 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.

In line with [Cont and 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 4.2.4.2). The overlapping portfolio network is reconstructed using the random network method employed in [Kok and Montagna \(2013\)](#) (see Appendix A.1.2.2).

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.³⁸ 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

$$p_{am}^t = p_{am}^{t_0} \max\{1 - \beta_{am} f_{am}^t; 0\}, \quad (7)$$

and is floored so that it never falls below zero. In equation 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.

4.2.2 Interbank Contracts In line with [Amini et al. \(2013\)](#) and [Hałaj and 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 and Kok \(2013\)](#) and employed by [Kok and 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).

Following [Kok and Montagna \(2013\)](#), we assume that interbank contracts are overnight contracts (i.e. mature every day) and are automatically rolled over, unless they are ex-

³⁸A more advanced study of price dynamics could for example be facilitated by modelling the limit order book in exchange-traded markets ([Paulin et al. \(2018\)](#)), and/or structurally modelling the market pricing mechanism for dealer-intermediated markets as in [Baranova et al. \(2017\)](#)). Even though we do not model limit order books here, our structural framework supports such modelling.

plicitly *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 precautionary liquidity hoarding (Acharya and 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).

4.2.3 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 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, see details in Appendix A.1.4.1) – 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. 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 (see details in Appendix A.1.4.2). 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.

When an institution j has offsetting reverse repo R_{ji} and repo contracts R_{kj} , and rehypothecation³⁹ is allowed, it has a ‘matched book’.⁴⁰ In such cases, equal and opposite margin calls will be due on the repo and reverse repo contracts when the price of collateral changes, and the institution can simply pass on the collateral it received in the

³⁹In this context, we loosely define rehypothecation as passing on collateral received as part of a reverse repo to secure (other) repo funding.

⁴⁰We call a repo and reverse repo position ‘offsetting’ if the notional, haircuts, and the set of collateral are the same for the reverse repo and repo contract.

reverse repo contract R_{ji} to the repo contract R_{kj} (or vice versa). Since (1) banks in the European system largely have offsetting total reverse repo and repo positions (according to the balance sheet data we use), and (2) we allow for rehypothecation, banks in our model are not significantly exposed to liquidity risk *unless* there are delays in the delivery of collateral (Gorton and Muir (2016)), which our model does not allow for.⁴¹ We assume that the banks in our system play the role of intermediary when providing repo funding to hedge funds (as their reverse repo R_j and repo books are broadly offsetting \tilde{R}_j , see Appendix A.1.2.1), whereas each hedge fund i is exposed to margin calls M_{ji} . We give a more detailed description of margin calls associated to repurchase agreements in Appendix A.1.4.1.

4.2.4 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.

4.2.4.1 Exposure Loss Contagion Exposure loss contagion is a form of node amplification (see Section 3.2). It occurs when liquidation following the default of bank $j \in \mathcal{B}$ leads to further contagion, which is 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.

4.2.4.2 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

⁴¹Even in the matched book case, there may be a maturity mismatch between repo and reverse repo positions that could lead to temporary liquidity pressures. We do not model this effect, but note that the setup of timesteps in our framework supports explicit modelling of contractual maturities, and so the ability to capture this risk if desired.

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 and Eisenbach (2015)).

Traditionally, many models have generally motivated delevering by using a binding leverage constraint λ_i (see e.g. Cont and Schaanning (2017), Greenwood et al. (2015)), though the authors of a recent paper Coen et al. (2019) model asset sales in the face of risk-weighted capital ratio and liquidity coverage ratios in addition. 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.

4.2.4.3 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 3.2. Hence, overlapping portfolio contagion and collateral contagion can reinforce each other, similar to the margin-price spiral Brunnermeier and Pedersen (2009) identified.

4.2.4.4 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.

4.3 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 situa-

tions when multiple constraints bind simultaneously, and also discuss the ‘pecking orders’ banks use to decide how to meet their constraints. Finally, we turn to investment funds and hedge funds, and describe how we model their behaviour.⁴²

We again stress that our structural 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, investment fund, 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.

4.3.1 Banks’ Behaviour: 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 and Tivnan \(2014\)](#), [Cont and Schaanning \(2017\)](#)). Such behaviour is consistent with the empirical findings of [Adrian and 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 [4.1.1.1](#) and [6](#)). 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 (i.e., if they behave as if buffers are requirements), their management buffer and target may exceed regulatory buffer standards.

4.3.2 Usability of Regulatory Buffers and Targets To assess the impact of regulatory buffer usability on systemic risk (see Section [5.5](#)), we introduce a parameter u that determines what fraction of each regulatory buffer (risk-based capital buffer, lever-

⁴²Our 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 structural 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.

age 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

$$\rho^B = \rho^M + (1 - u)(y^\rho \rho_i^{CB}), \quad (8)$$

$$\lambda^B = \lambda^M + (1 - u)(y^\lambda \lambda_i^{CB}), \quad (9)$$

$$\Lambda^B = (1 - u)\Lambda^S, \quad (10)$$

where y^λ and y^ρ determine 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 2.

4.3.3 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.⁴³ 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 (details on banks' behaviour to meet constraints are given in Appendix A.1.4.2).⁴⁴

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

⁴³In line with the literature, we focus on actions that banks can take that can impact on financial markets and other financial institutions; but note that banks may also have actions available to them that could increase their ability to absorb stress. These actions – such as cost-cutting, restricting dividends, and drawing on central bank liquidity facilities – may be available to banks when projecting the impact of microprudential stress tests and so can be reflected in the initial shock we input to the model, so we do not focus on them here. The addition of central bank liquidity facilities in particular would be a natural and important extension to this work; but given our focus on solvency shocks in this paper, and the fact that solvency constraints drive our results, the exclusion of liquidity support from the current model does not have a bearing on our findings.

⁴⁴For each pecking order, we assume that assets are liquidated proportionally within an asset class.

holdings (see details in Appendix [A.1.4.1](#) and [A.1.4.2](#)).

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).

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 (see details in Appendix [A.1.4.2](#)).

Defending the leverage ratio - Banks first delever by using cash to proportionally pay back liabilities L_i (see details in Appendix [A.1.4.2](#)).⁴⁵ 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.

Defending the Liquidity Coverage Ratio (LCR) - 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 (see details in Appendix [A.1.4.2](#)).

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 ([Federal Reserve Bank of St. Louis \(2010\)](#), [McDonald and Paulson \(2014\)](#), [Brown and Dinċ \(2011\)](#)). 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,

⁴⁵We note that in the UK, central bank reserves do not contribute to the leverage ratio, so this option would not be available.

may have no impact on the risk-weighted capital ratio (e.g., if zero risk-weighted assets are liquidated and the raised cash is 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 least binding (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 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 $\Lambda_i \geq \Lambda_i^B$.

4.3.4 Bank Failure Contagion models often assume that banks are liquidated on default (see e.g. [Kok and 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*.⁴⁶

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 and Montagna \(2013\)](#)), and unsecured creditors take losses (see Section 4.2.2 and 4.2.4.1) while secured creditors take title to repo collateral (see Section A.1.4.2). In the other edge case, *contagion-free resolution*, every defaulted bank is resolved without any contagion: the bank simply becomes inactive.⁴⁷ In reality, the consequences of bank failure would be between these two extremes (see e.g. [Bank of England \(2017a\)](#), [Klimek et al. \(2015\)](#), [Hüser et al. \(2017\)](#), [Chennells and 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.

⁴⁶We define failure as a breach of minimum capital requirements or illiquidity.

⁴⁷This 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.

4.3.5 Investment fund behaviour Each investment fund $i \in \mathcal{M}$ has the obligation to pay back shares at their prevailing net asset value when they are redeemed by investors. To do so, an investment fund first uses cash. If this is insufficient to meet redemptions, they liquidate tradable assets in a ‘vertical slice’ (i.e. proportional to their asset holdings). We present more details in the Appendix [A.1.4.3](#).

4.3.6 Hedge fund behaviour The behaviour of a hedge fund $i \in \mathcal{H}$ is driven by (1) the obligation to meet margin calls (see a specification in Appendix [A.1.4.2](#) and [A.1.4.1](#)), (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 [4.3.3](#)): 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 $\lambda_i = \frac{E_i}{A_i}$ ⁴⁸ falls underneath its buffer value λ_i^B to return to its leverage target λ_i^T .⁴⁹ As explained in [Bookstaber, Paddrik and 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.

5 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 baseline, we use the European Banking Authority (EBA) 2018 stress test results.⁵⁰ The EBA stress test was conducted with static balance sheets and did not model second-round effects that could arise as a consequence of banks’ responses ([Ebner \(2018\)](#)), but surveys involving participants in previous EBA stress tests have suggested that these effects could be sizable ([Brinkhoff et al. \(2018\)](#)).⁵¹ 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. We then use this output as an *input* for our system-

⁴⁸Note the leverage ratio λ_i of hedge funds has a different definition from the leverage ratio of banks (see equation [2](#)).

⁴⁹The 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 [2](#).

⁵⁰See the 2018 EBA stress test outputs here: <https://eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing/2018/results>.

⁵¹Participants 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.

wide stress testing model. Our system-wide model then accounts for potential contagion mechanisms that might amplify this initial shock.⁵²

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 more financial system-wide perspective by including (representative) non-bank financial institutions, and show that this leads to further - albeit limited⁵³ - 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.

5.1 Default Parameters and Visualisation

Unless otherwise stated, all experiments use the same default parameters, which are given in Table 2.⁵⁴ 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 and 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

⁵²In general, regulatory stress test scenarios feature severe shocks to economic and financial variables such as GDP, unemployment, house prices and stock prices (see e.g. ESRB (2018)). The severity of these scenarios is in general calibrated to historical stresses, and can also reflect current levels of risk across the financial and economic system (see e.g. Bank of England (2015a) for a discussion of how the Bank of England increases scenario severity when underlying risks are increasing). As such, scenarios implicitly capture the effects of higher-order contagion mechanisms to some extent; but do not disentangle 'initial shocks' from the dynamics that amplify them into crises. In principle, an adverse scenario in a system-wide stress test should not be made overly severe, to avoid double counting contagion effects; instead an initial shock should be amplified by the system to result in a crisis outcome. In recognition of this point, we consider below how endogenous amplification effects can amplify relatively mild initial shocks - and how, under certain conditions, this leads to outcomes more severe than the first order impact of significantly larger initial shocks.

⁵³The limited impact of including non-banks may in part be due to the fact that we do not account for balance sheet or investment model heterogeneity in hedge funds; and by the limited size of the assets held by these institutions in Europe.

⁵⁴Each 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.

banking system defaults (see Appendix A.1.3.1 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 4.2.1 and 4.2.2 for the reconstruction methods). 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 and Kapadia (2010)); per x-y-axis point all else is kept constant. The shaded areas around the result lines in the figures, plot the error bars.

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 lines⁵⁵). Since the microprudential stress test does not capture contagion defaults, the grey 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 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 4.2.1 for details).

Of course, as is common for models of contagion dynamics, the *magnitude* of systemic losses generated in these experiments is sensitive to a number of assumptions and parameters.⁵⁶ Accordingly, and in line with use-cases of models using similar techniques (see e.g. Haldane and Turrell (2018)), 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.

⁵⁵The 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.

⁵⁶Cont and 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 2: Default settings for the figures in the result Section 5.

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 & investment funds turned off.	This choice is motivated by data quality. Since our initialisation of investment funds 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 is may lead to an underestimation of systemic risk (see e.g. Kok and 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 and Shin (2004) , Cifuentes et al. (2005) , Gai and 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.

5.2 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 (default parameters) and 4 (tripled risk-weighted capital buffers ρ_i^{CB}) show the 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. Two key findings emerge: 1) microprudential stress tests alone are insufficient to assess financial stability; and 2) whether banks fail in a disorderly or managed way has a significant impact on financial stability. In addition, we find that 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 (shown by Figures 10 and 11 in Appendix A.3).⁵⁷

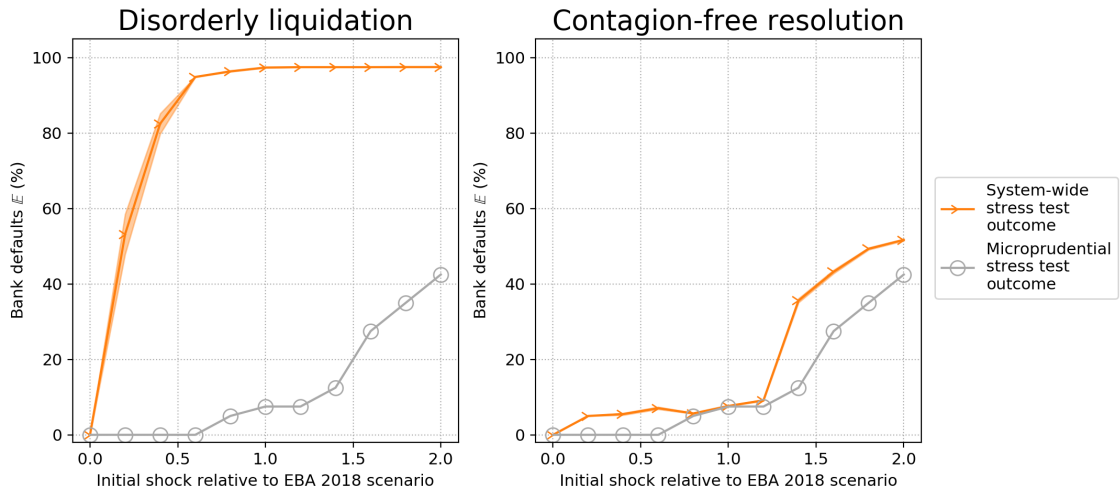


Figure 3: This figure shows systemic risk \mathbb{E} as a function of the scaled impact of the 2018 EBA scenario. The coloured lines show the system-wide stress test outcome and the grey lines show the (scaled 2018 EBA) microprudential stress test outcome. These results illustrate that, for a given microprudential 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*,⁵⁸ which implies that microprudential stress tests alone are insufficient to assess

⁵⁷A 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, more recently, of 3.25% with central bank reserves excluded from the leverage exposure measure – and an additional countercyclical leverage buffer (CCLB) for several years (Bank of England (2015b)); 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 (2018b)). When we remove the leverage ratio constraint in our model, contagion reduces substantially, but the qualitative findings remain robust.

⁵⁸The shock-amplifying tendency of the system depends, among others, on the ‘usability’ (see Section 5.4) and size (see Section 5.6) of regulatory buffers as well as the resolution regime (this section).

financial stability. For example, comparing the orange lines in the left and right panels of Figure 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 $x = 0.5$ – demonstrating that the microprudential stress test outcome provides incomplete insight into the system’s stability. Comparing the two panels of Figure 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 4.3 on the banks’ behaviour). In this case (the right plot), the excess systemic risk (orange line) above the initial impact (grey line) is solely generated by such ‘pre-default contagion’.

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 and Figure 4 (where risk-weighted capital buffers are tripled) shows that the system becomes more shock absorbing when capital buffers are increased. Increasing capital buffers not only delays the onset of the non-linear jump towards instability in the case of ‘disorderly liquidation’ (from $x = 0$ in Figure 3 to $x = 1.2$ in Figure 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.⁵⁹

Second, it is clear from Figure 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.

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 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,

⁵⁹In 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 4. If banks 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.

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) results in a systemic risk level of $\mathbb{E} \approx 10\%$. This is well below the degree of instability we observe when systemic amplification mechanisms are included, even for mild initial shocks, eg of $x = 0.25$, for which $\mathbb{E} \approx 50\%$. So in such a case, the underestimation of systemic risk in the microprudential stress test would be 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*. We illustrate the relative importance of the leverage ratio in Figures 10 and 11 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 5 in Appendix A.3).⁶⁰

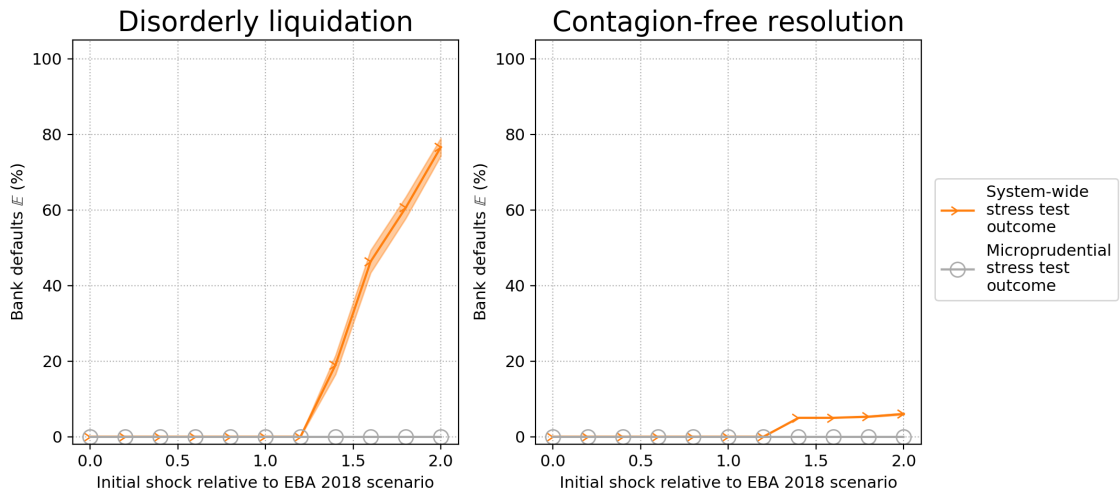


Figure 4: This figure has the same set-up as Figure 3, except here we have tripled the combined risk-weighted capital buffer ρ_i^{CB} (i.e. we have set $y^\rho = 3$ in equation 8). Tripling the buffer reduces the shock-amplifying tendency of the financial system and delays the non-linear divergence of 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’.

5.3 Contagious Feedback Loops Between Banks and Non-Banks

Our baseline model only includes banks and their interactions (see Table 2). In the second policy experiment, we add investment funds and hedge funds to our financial system to

⁶⁰The magnitude of amplification under the leverage ratio constraint is also a function of the size of the management buffer banks seek to maintain over minimum requirements. Reducing the buffer by adjusting banks’ leverage ratio target reduces the magnitude of systemic risk amplification in some cases, as illustrated in Figure 16 in Appendix A.3; but the qualitative results hold.

show that *expanding the types of financial institutions included in the stress test changes the expected financial stability outcomes*. In Figure 5, we show how systemic risk outcomes change when we add non-banks, first separately and then together, for the ‘disorderly liquidation’ case. We also 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. Including investment funds captures the risk that the price impacts of banks’ asset sales affect the performance of investment funds. That, in turn, could prompt shareholder redemptions, which may force investment funds to sell into the market to meet them, further reinforcing fire sale dynamics. To the best of our knowledge, we are the first to model the interactions of the above-mentioned contagious feedback loops among heterogeneous banks, hedge funds, and investment funds in a realistic financial system setting.

We find that *including hedge funds and investment funds in our system-wide stress test increases systemic risk modestly*. Accordingly, excluding hedge funds and investment funds 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 investment funds 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.⁶¹ 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 and Roberts-Sklar \(2017\)](#) who also note that while investment funds 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

⁶¹Hedge 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 investment funds 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.

shocks in order to obtain a holistic picture of systemic stability.

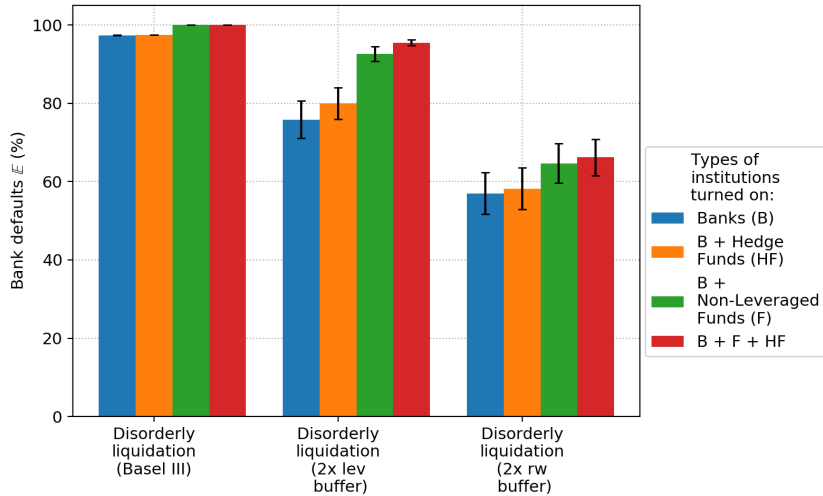


Figure 5: This figure shows the average fraction of bank defaults in a systemic event \mathbb{E} for stress tests including different constellations of institutions (i.e. only banks (B); B and investment funds (F); B and hedge funds (HF); B, F and HF) and for different regulatory regimes (i.e. for the ‘disorderly liquidation’, for Basel III settings or for Basel III settings with the leverage buffer λ_i^{CB} and risk-weighted capital buffer ρ_i^{CB} doubled in turn). 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.

5.4 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 4.2.4, 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 that the structural 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 6 where, for instance, the label *O&E* means that only overlapping portfolio contagion and exposure loss contagion are included (‘turned on’).

Figure 6a shows that contagion mechanisms are mutually amplifying.⁶² For example, if we assume a price impact of 5%, we find that systemic risk due to exposure loss

⁶²To 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.

contagion is moderate in size (i.e. around $\mathbb{E} \approx 20\%$), and instability due to overlapping portfolio contagion, 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 $\mathbb{E} \approx 80\%$). Figure 6b, which shows a direct measure of this amplification, illustrates this finding. Focusing on the 5% price impact point in Figure 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 five when all contagion mechanisms are considered.⁶³ Based on these findings, it is clear that *modelling contagion mechanisms in isolation may lead to an underestimation of systemic risk by a factor as large as 5.5*. 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 6b. By comparing the height of the bars for the different price impacts in Figure 6b (and 12 in Appendix A.3), 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.⁶⁴ This result clearly shows that market illiquidity can act as a powerful amplifier of other contagion mechanisms. To the best of our knowledge, we are also the first to highlight that the degree of amplification is heterogeneous for different sets of contagion mechanisms and in market illiquidity.

⁶³The same results are shown in absolute terms in Figure 12 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.

⁶⁴In Figure 6b, we cap the sum of amplification due to individual mechanisms at $\mathbb{E} = 100\%$.

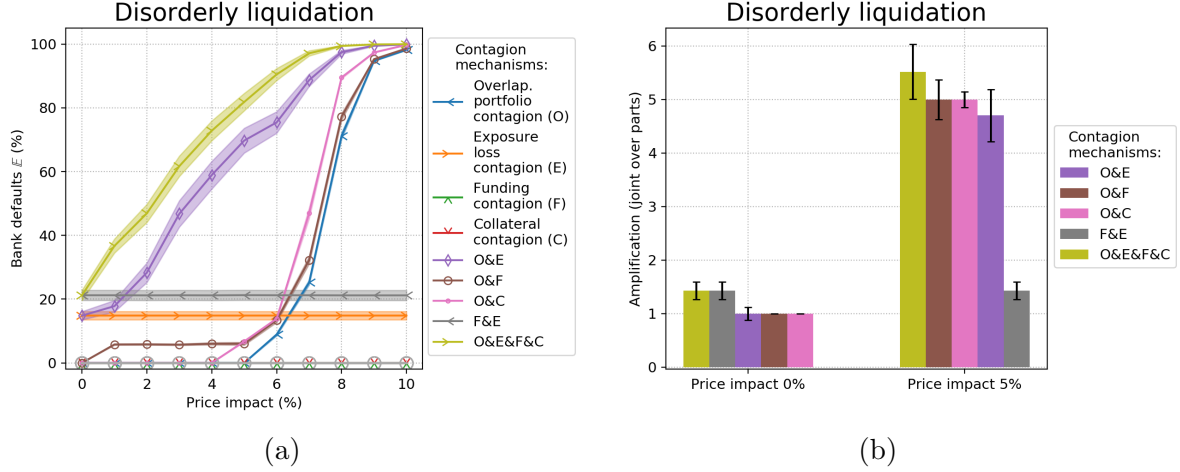


Figure 6: This figure shows the amplification among contagion mechanisms (overlapping portfolio contagion, exposure loss contagion, funding contagion and collateral contagion) for the case of ‘disorderly liquidation’ where the leverage buffer λ_i^{CB} is made two-and-a-half times larger (i.e. $y^\lambda = 2.5$). 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 6a shows systemic risk \mathbb{E} as a function of the price impact for various combinations of contagion mechanisms. Plot 6b also elucidates Plot 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 (capped at 100%). 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.

5.5 ‘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 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 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 $u = 100\%$.⁶⁵ As far as we are aware, 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.

⁶⁵This result also holds for the case of ‘disorderly liquidation’, see Figure 14 in Appendix A.3. In our experiments, the usability of the liquid asset buffer is not important for systemic risk (see Figure 15 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.

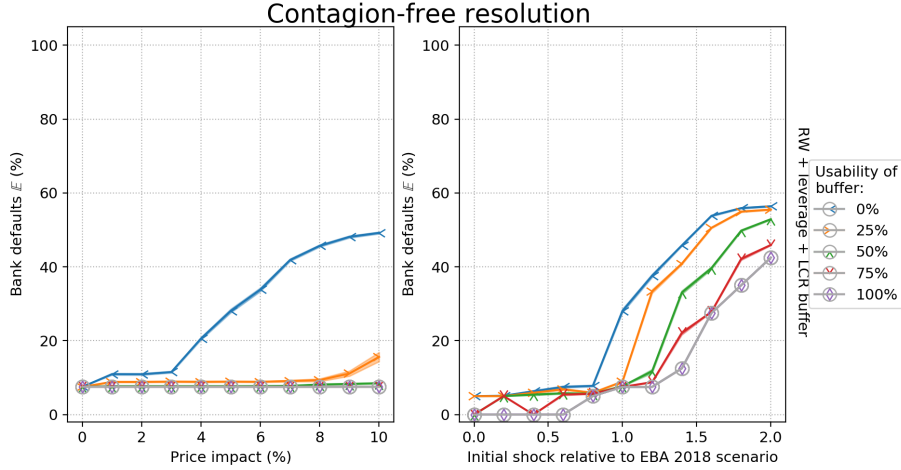


Figure 7: This figure shows systemic risk \mathbb{E} for the case of ‘contagion-free resolution’ as a function of the price impact (left plot) or as a function of the scaled 2018 EBA scenario (right plot) for different levels of ‘usability’ of capital buffers. If banks consider $u = 25\%$ of their regulatory 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 4.3.2 for implementation details). Resilience increases in the ‘usability’ of regulatory buffers. This result holds irrespective of the market liquidity or the stress scenario.

5.6 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 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.⁶⁶

⁶⁶We 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 2.

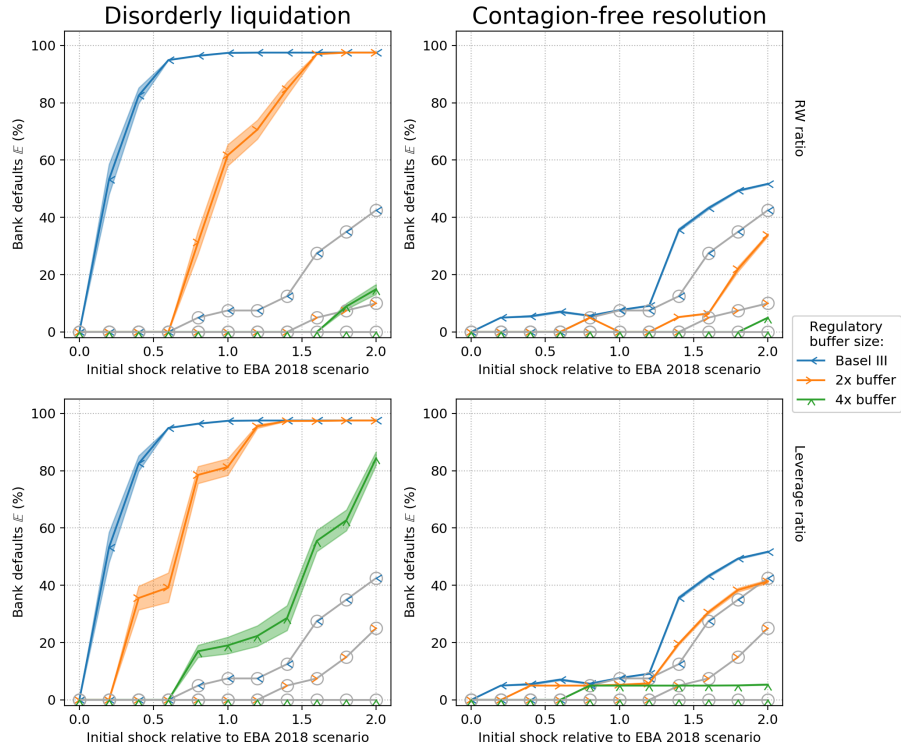


Figure 8: This figure shows systemic risk \mathbb{E} for both the case of ‘disorderly liquidation’ and ‘contagion-free resolution’ as a function of the (scaled) 2018 EBA scenario for different regulatory buffer sizes. The top row shows the effect of doubling and quadrupling the regulatory risk-weighted 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 purely microprudential stress tests (grey-coloured lines) rather than system-wide stress tests (coloured lines) to calibrate regulatory buffers are at risk of overestimating resilience.

Figure 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 17 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.⁶⁷ It is

⁶⁷Figure 17 in Appendix A.3 further illustrates this finding. It shows that while the microprudential stress test estimates that Basel III buffers can effectively mitigate systemic risk to a level under $\mathbb{E} = 10\%$, when system-wide dynamics are taken into account, in the edge case of disorderly liquidation buffers need to be more than doubled to achieve the same outcome.

also 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. We also note that increasing regulatory buffers tends to bring down contagion defaults more than it reduces initial defaults (see Figure 17 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.

6 Policy Implications and Conclusion

In a highly connected financial system, seemingly localised shocks can be propagated and amplified to take on systemic importance. While this is widely recognised, this reality is only partly and inconsistently reflected in the design of banking system stress tests, which are not yet system-wide in scope, and only partly – if at all – combine multiple interacting contagion and amplification mechanisms as well as the behavioural responses of heterogeneous financial institutions to shocks. We have outlined a structural framework for the development of system-wide financial stress tests with multiple interacting contagion and amplification channels and heterogeneous financial institutions. We have explained how 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 representative 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 potential hidden weaknesses in the resilience of the financial system. This raises the concern that current stress test results that do not incorporate such effects, or do so only partially, risk overstating systemic resilience. 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, provides grounds for careful consideration of these risks. Because our findings demonstrate that financial stability is critically defined by amplification dynamics, such stress tests can be meaningfully complemented by macroprudential overlays (e.g. [Dees and Henry \(2017\)](#), [Fique \(2017\)](#)) and by considering contagion dynamics in regulatory exercises (e.g. [Bank of England \(2017b\)](#)).

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 can arise when banks take action to avoid using their buffer capacity – actions that are individually rational but collectively destabilising. This should motivate careful consideration on the part of regulators when setting stress test hurdle rates⁶⁸, and around factors that may influence banks’ willingness to use their capital buffers.

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. Using microprudential stress tests to calibrate capital requirements can therefore be meaningfully complemented by the use of system-wide stress tests. The findings of such exercises could for example be used to calibrate capital requirements under Pillar II of the Basel supervisory framework ([BCBS \(2009\)](#)), and could also inform the calibration of the countercyclical capital buffer (see e.g. [Bank of England \(2017b\)](#)).⁶⁹ Indeed, the incorporation of feedback and amplification effects in regulatory stress tests that are used to inform capital-setting is arguably a step in this direction (see e.g. [Bank of England \(2017b\)](#)).

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 parallel 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

⁶⁸See for example a discussion of this point in [Bank of England \(2013\)](#).

⁶⁹The 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. \(2019\)](#), and with the application of additional capital buffer standards to globally systemically important banks.

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 – and 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 structural 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. With the introduction of our structural framework that can enable regulators to build and use large-scale, data-driven models, the importance of data availability at a granular level grows further – particularly given the importance of calibration.⁷⁰ And, finally, although our structural framework marks an important step forwards, it is itself incomplete. Integration of derivatives markets, for example, presents a key modelling challenge.⁷¹

⁷⁰More data (in the format described in Section 3.1.1) 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 are required.

⁷¹So far derivatives markets have only, partially, been stress tested on a stand-alone basis (see e.g. [Bardoscia et al. \(2018\)](#), [Paddrik and Young \(2017\)](#), [Paddrik et al. \(2016\)](#)). Their role in not only transmitting but in hedging risk is an important component to capture in system-wide modelling.

References

- Abad, J., Aldasoro, I., Aymanns, C., D’Errico, M., Rousová, L. F., Hoffmann, P., Langfield, S., Neychev, M. and 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.
- Acharya, V. V. and Skeie, D. (2011), ‘A model of liquidity hoarding and term premia in inter-bank markets’, Journal of Monetary Economics **58**(5), 436–447.
- Adrian, T. and Shin, H. S. (2010), ‘Liquidity and leverage’, Journal of financial intermediation **19**(3), 418–437.
- Aikman, D., Chichkanov, P., Douglas, G., Georgiev, Y., Howat, J. and King, B. (2019), ‘System-wide stress simulation’, Staff Working Paper (809).
- Aldasoro, I., Gatti, D. D. and Faia, E. (2017), ‘Bank networks: Contagion, systemic risk and prudential policy’, Journal of Economic Behavior & Organization **142**, 164–188.
- Amini, H., Cont, R. and Minca, A. (2013), ‘Resilience to contagion in financial networks’, Mathematical finance .
- Anderson, R., Dannielson, J., Baba, C., Das, U. S., Kang, H. and Segoviano, M. (2018), ‘Macroprudential stress tests and policies: Searching for robust and implementable frameworks’, Technical report, IMF Working Paper.
- Armour, J., Awrey, D., Davies, P., Enriques, L., Gordon, J. N., Mayer, C. and Payne, J. (2016), Principles of Financial Regulation, Oxford University Press.
- Awrey, D. (2019), ‘Split derivatives: Inside the world’s most misunderstood contract’, Yale J. on Reg. **36**, 495.
- Aymanns, C., Caccioli, F., Farmer, J. D. and Tan, V. W. (2016), ‘Taming the basel leverage cycle’, Journal of financial stability **27**, 263–277.
- Aymanns, C. and 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. and Wetzler, T. (2018), ‘Models of financial stability and their application in stress tests’, Handbook of Computational Economics **4**, 329–391.
- Bank of England (2013), A framework for stress testing the UK banking system, Technical report, Bank of England.
- Bank of England (2015a), The Bank of England’s approach to stress testing the UK banking system, Technical report, Bank of England.

- Bank of England (2015b), The framework of capital requirements for UK banks, Technical report, Bank of England.
- Bank of England (2017a), The Bank of England’s approach to resolution, Technical report, Bank of England.
- Bank of England (2017b), Stress testing the UK banking system: 2017 results, Technical report, Bank of England.
- Bank of England (2018a), Financial Stability Report, November 2018, Technical report, Bank of England.
- Bank of England (2018b), Stress testing the UK banking system: 2018 results, Technical report, Bank of England.
- Baranova, Y., Coen, J., Lowe, P., Noss, J. and Silvestri, L. (2017), ‘Simulating stress across the financial system: the resilience of corporate bond markets and the role of investment funds’, Financial Stability Paper (42).
- Bardoscia, M., Barucca, P., Brinley, A. and Hill, J. (2017), ‘The decline of solvency contagion risk’, Staff Working Paper (662).
- Bardoscia, M., Bianconi, G. and Ferrara, G. (2018), ‘Multiplex network analysis of the UK OTC derivatives market’.
- Battiston, S., Gatti, D. D., Gallegati, M., Greenwald, B. and Stiglitz, J. E. (2012), ‘Liaisons 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. and Caldarelli, G. (2012), ‘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), BaselIII: A global regulatory framework for more resilient banks and banking systems, Technical report, Basel Committee on Banking Supervision.
- BCBS (2016), Literature review on integration of regulatory capital and liquidity instruments, Technical report, Basel Committee on Banking Supervision.
- Berkowitz, J. and 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>.

- 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.
- Bookstaber, R., Cetina, J., Feldberg, G., Flood, M. and 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., Paddrik, M. and Tivnan, B. (2014), ‘An agent-based model for financial vulnerability’, Office of Financial Research Working Paper Series **14**(05).
- Bouchaud, J.-P. (2010), ‘Price impact’, Encyclopedia of quantitative finance .
- Braun-Munzinger, K., Liu, Z. and Turrell, A. (2016), An agent-based model of dynamics in corporate bond trading, Technical report, Bank of England.
- Brazier, A. (2017), How to: MACROPRU. 5 principles for macroprudential policy, Technical report, Bank of England.
- Brinkhoff, J., Langfield, S. and 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. and Siegert, C. (2017), ‘Measuring the macroeconomic costs and benefits of higher UK bank capital requirements’, Financial Stability Paper (35).
- Brown, C. O. and 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, Internat. Center for Monetary and Banking Studies.
- Brunnermeier, M. K. and Pedersen, L. H. (2009), ‘Market liquidity and funding liquidity’, Review of Financial studies **22**(6), 2201–2238.

- Burrows, O., Learmonth, D., McKeown, J. and Williams, R. (2012), ‘RAMSI: a top-down stress-testing model developed at the bank of england’.
- Burrows, O., Low, K. and Cumming, F. (2015), ‘Mapping the UK financial system’.
- Caccioli, F., Farmer, J. D., Foti, N. and 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. and Rockmore, D. (2015), ‘Overlapping portfolios, contagion, and financial stability’, Journal of Economic Dynamics and Control **51**, 50–63.
- Caccioli, F., Shrestha, M., Moore, C. and Farmer, J. D. (2014), ‘Stability analysis of financial contagion due to overlapping portfolios’, Journal of Banking & Finance **46**, 233–245.
- Cecchetti, S. and Anil, K. (2018), ‘12 what binds? interactions between bank capital and liquidity regulations’, The Changing Fortunes of Central Banking p. 192.
- Cetina, J., Lelyveld, I. and Anand, K. (2015), ‘Making supervisory stress tests more macroprudential: Considering liquidity and solvency interactions and systemic risk’, BCBS Working Paper .
- Chennells, L. and Wingfield, V. (2015), ‘Bank failure and bail-in: an introduction’.
- Churm, R. and Nahai-Williamson, P. (2019), Four years of concurrent stress testing at the Bank of England: Developing the macroprudential perspective, in A. Siddique, I. Hasan and D. Lynch, eds, ‘Stress Testing (2nd Edition), Approaches, Methods and Applications’, Risk Books.
- Cifuentes, R., Ferrucci, G. and Shin, H. S. (2005), ‘Liquidity risk and contagion’, Journal of the European Economic Association **3**(2-3), 556–566.
- Coen, J., Lepore, C. and Schaanning, E. (2019), ‘Taking regulation seriously: fire sales under solvency and liquidity constraints’, Staff Working Paper (793).
- Cont, R., Moussa, A. et al. (2010), ‘Network structure and systemic risk in banking systems’.
- Cont, R. and Schaanning, E. (2017), ‘Fire sales, indirect contagion and systemic stress testing’.
- Copeland, A., Martin, A. and Walker, M. (2014), ‘Repo runs: Evidence from the tri-party repo market’, The Journal of Finance **69**(6), 2343–2380.
- Coval, J. and Stafford, E. (2007), ‘Asset fire sales (and purchases) in equity markets’, Journal of Financial Economics **86**(2), 479–512.

- Czech, R. and Roberts-Sklar, M. (2017), ‘Investor behaviour and reaching for yield: evidence from the sterling corporate bond market’, Staff Working Paper (685).
- Danielsson, J. and Shin, H. S. (2003), ‘Endogenous risk’, Modern risk management: A history pp. 297–316.
- Dawid, H., Harting, P., van der Hoog, S. and Neugart, M. (2016), ‘A heterogeneous agent macroeconomic model for policy evaluation: improving transparency and reproducibility’.
- De Haan, L. and van den End, J. W. (2013), ‘Bank liquidity, the maturity ladder, and regulation’, Journal of Banking & Finance **37**(10), 3930–3950.
- Dees, S. and Henry, J. (2017), ‘Stress-test analytics for macroprudential purposes: Introducing STAMPe’, SATELLITE MODELS p. 13.
- Drehmann, M., Borio, C. E., Gambacorta, L., Jimenez, G. and Trucharte, C. (2010), ‘Countercyclical capital buffers: exploring options’.
- Duarte, F. and Eisenbach, T. M. (2015), ‘Fire-sale spillovers and systemic risk’.
- Ebner, A. (2018), ‘The financial stability aspects of the EU-wide stress test’, Journal of Financial Regulation **4**(2), 326–336.
- 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>.
- Enriques, L., Romano, A. and Wetzer, T. (2019), ‘Network-sensitive financial regulation’, The Journal of Corporation Law (forthcoming) .
- ESRB (2018), Adverse macro-financial scenario for the 2018 EU-wide banking sector stress test, 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. and 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.
- Federal Reserve Bank of St. Louis (2010), ‘AIG - Goldman Sachs collateral call timeline’, https://fraser.stlouisfed.org/files/docs/historical/fct/fcic/fcic_docs_aig_goldman_timeline_20100630.pdf.

- Fique, J. (2017), ‘The macrofinancial risk assessment framework (mfrac)’, Bank of Canada Technical Report No. 111 .
- FSB (2017), Basel III: Finalising post-crisis reforms, Technical report, Financial Stability Board.
- Gai, P., Haldane, A. and Kapadia, S. (2011), ‘Complexity, concentration and contagion’, Journal of Monetary Economics **58**(5), 453–470.
- Gai, P. and Kapadia, S. (2010), Contagion in financial networks, in ‘Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences’, The Royal Society, p. rspa20090410.
- Goodhart, C. (2013), ‘Ratio controls need reconsideration’, Journal of Financial Stability **9**(3), 445–450.
- Goodhart, C. et al. (2008), ‘Liquidity risk management’, Banque de France Financial Stability Review **11**, 39–44.
- Gorton, G. B. and Metrick, A. (2009), Haircuts, Technical report, National Bureau of Economic Research.
- Gorton, G. and Metrick, A. (2012), ‘Securitized banking and the run on repo’, Journal of Financial Economics **104**(3), 425–451.
- Gorton, G. and Muir, T. (2016), Mobile collateral versus immobile collateral, Technical report, National Bureau of Economic Research.
- Greenwood, R., Landier, A. and Thesmar, D. (2015), ‘Vulnerable banks’, Journal of Financial Economics **115**(3), 471–485.
- Greenwood, R., Stein, J. C., Hanson, S. G. and Sunderam, A. (2017), ‘Strengthening and streamlining bank capital regulation’, Brookings Papers on Economic Activity **2017**(2), 479–565.
- Halaj, G. (2018), ‘Agent-based model of system-wide implications of funding risk’.
- Halaj, G. and Kok, C. (2013), ‘Assessing interbank contagion using simulated networks’, Computational Management Science **10**(2-3), 157–186.
- Haldane, A. G. and Turrell, A. E. (2018), ‘An interdisciplinary model for macroeconomics’, Oxford Review of Economic Policy **34**(1-2), 219–251.
- Heider, F., Hoerova, M. and Holthausen, C. (2009), ‘Liquidity hoarding and interbank market spreads: The role of counterparty risk’.
- Hüser, A.-C., Halaj, G., Kok, C., Perales, C. and van der Kraaij, A. (2017), ‘The systemic implications of bail-in: a multi-layered network approach’, Journal of Financial Stability .

- Hüser, A.-C. and Kok, C. (2019), ‘Mapping bank securities across euro area sectors: comparing funding and exposure networks’.
- IOSCO (2017), Report on the fourth IOSCO hedge funds survey, Technical report, International Organization of Securities Commissions.
- Judge, K. and Berner, R. (2019), ‘The data standardization challenge’.
- Klimek, P., Poledna, S., Farmer, J. D. and 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. and Montagna, M. (2013), Multi-layered interbank model for assessing systemic risk, Technical report, European Central Bank.
- Lo, A. W. (2017), Adaptive markets: Financial evolution at the speed of thought, Princeton University Press.
- Luna, P. G. and Hardy, B. (2019), ‘Non-bank counterparties in international banking’, BIS Quarterly Review pp. 15–33.
- McDonald, R. and Paulson, A. (2014), AIG in hindsight, Technical report, Federal Reserve Bank of Chicago.
- Paddrik, M. E. and Young, P. (2017), ‘How safe are central counterparties in derivatives markets?’.
- Paddrik, M., Rajan, S. and Young, H. P. (2016), ‘Contagion in the CDS market’, Office of Financial Research Working Paper pp. 16–12.
- Paulin, J., Calinescu, A. and Wooldridge, M. (2018), ‘Understanding flash crash contagion and systemic risk: A micro-macro agent-based approach’, arXiv preprint arXiv:1805.08454 .
- 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. and 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. and Singh, M. (2011), ‘The nonbank-bank nexus and the shadow banking system’, IMF Working Papers pp. 1–18.
- Prudential Regulation Authority (2017), Constulation paper CP15/17: The minimum requirement for own funds and eligible liabilities (MREL) – buffers, Technical report, Bank of England.

Schnabel, I. and Shin, H. S. (2004), ‘Liquidity and contagion: The crisis of 1763’, Journal of the European Economic Association **2**(6), 929–968.

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

Wiersema, G., Kleinnijenhuis, A. M., Wetzer, T. and Farmer, J. D. (2019), ‘Inherent instability: Scenario-free analysis of financial systems with interacting contagion channels’, INET Working Paper Series .

A Technical Appendix: Foundation of System-Wide Stress Testing

A.1 Technicalities of Model Implementation

A.1.1 Notation Table 3 and 4 give the definition of the variables used in this paper.

Table 3: 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 .	
Repurchase Agreements	R_i	Reverse repo contract of institution i .	
	\tilde{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 .	
	O_i	Other assets of institution i .	
Markets	\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 4: Shows the definition of notation (this table is a continuation of Table 3).

Category	Subcategory (if any)	Variable	Definition
Con- straints	Risk- weighted capital ratio	ρ_i	Risk-weighted (rw) capital ratio of bank i .
		ρ^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.	
\bar{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		λ_i	Leverage ratio of bank i .
		λ^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.
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^f	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.
NAV		$\tilde{\omega}_p$	Inflow rate associated to assets of type p .
		$\tilde{\omega}_l$	Outflow rate associated to assets of type l .
Behaviour		η_i	Net asset value (NAV) of investment fund i .
		χ_i^t	Relative loss in NAV at time t of investment fund i in comparison with time t_0 .
		$u\%$	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.
Systemic risk measure		\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 .
		\mathbb{E}	Average extent of a systemic event (average fraction of bank defaults in a systemic event).
		\mathbb{P}	Probability of a systemic event.
Sets		\mathcal{S}	Set of simulation runs in which a systemic event (if at least 5% of banks default) occurs.
		$f_{\mathcal{D}}(n)$	Fraction of bank defaults in case of a systemic event in simulation run n .
		\mathcal{N}	Number of simulation runs.
		\mathcal{F}	Set of financial institutions.
		\mathcal{B}	Set of banks.
Sets		\mathcal{M}	Set of investment funds.
		\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{I}	Set of banks that defaulted due to the adverse scenario (set of initially defaulted banks).

A.1.2 Further Details on Initialisation

A.1.2.1 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 4.1.3). 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 15), which may trigger the hedge fund to engage in fire sales (see Appendix A.1.4.2).

Given the above survey information and using *ECB Statistical Warehouse Data* on the aggregate hedge fund size and its aggregate asset allocation,⁷² 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 $i \in \mathcal{H}$,⁷³ 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 .⁷⁴ As we have data on the reverse repo R_i position of each bank $i \in \mathcal{B}$ (see Section 4.1.1), the total estimated size of the hedge fund sector in Europe (from the *ECB Statistical Warehouse Data*), and the

⁷²See *ECB Statistical Warehouse*: <https://sdw.ecb.europa.eu/browse.do?node=9691340>.

⁷³In 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.

⁷⁴As 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. Second, 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. \(2018\)](#).

leverage λ_i of each hedge fund $i \in \mathcal{H}$ (FCA (2015)), we can derive the asset value A_i and repo size \tilde{R}_i ⁷⁵ 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.

A.1.2.2 Further Details on Financial Contracts

Tradable Assets As discussed in Section 4.2.1, 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.1.2.1) 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 and 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 and 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 and Montagna (2013).

A.1.2.3 Markets To estimate the price impact (see Section 4.2.1), we set the denominator of the cumulative fraction of net asset sales f_{am}^t , which appears in the price impact function (see equation 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.1.2.2). 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 $p_{am}^{t_0} = 1$.

A.1.2.4 Constraints Here we discuss the regulatory parameters that are associated to the Basel III regulatory capital requirements and buffer standards discussed in Section 4.1.1.1. Let us start with explaining how the risk-weights ω_p in the risk-weighted capital requirement ρ_i (see equation 1) are set. In line with the Basel III standard-

⁷⁵Namely, 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 $i \in \mathcal{B}$.

ised 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 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 $\{0.05, 1, 1, 0.5\}$.⁷⁶ 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 $\{0, 0.5, 0, 1, 1, 0\}$.⁷⁷ 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) are publicly listed.⁷⁸

A.1.2.5 Behaviour No data available, as discussed in Section 3.1.4.

A.1.3 Default Configuration

A.1.3.1 y-axis: Systemic Risk Measure In line with, but a generalisation upon [Schnabel and Shin \(2004\)](#), [Cifuentes et al. \(2005\)](#), [Gai and Kapadia \(2010\)](#), [Caccioli](#)

⁷⁶If repo contracts \tilde{R}_i are secured with HQLA assets the outflow rate is zero instead of one.

⁷⁷Again, if reverse repo contracts R_i are secured with HQLA assets the inflow rate is zero instead of one.

⁷⁸See 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/other-systemically-important-institutions-o-siis-/2017>. See the list 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.

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

$$\mathbb{E} := \frac{1}{|\mathcal{S}|} \sum_{n \in \mathcal{S}} f_{\mathcal{D}}(n), \quad (11)$$

where the terms of equation 11 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

$$\mathcal{S} := \{n \in [1, N] : \mathbb{1}_{SE}(n) = 1\}, \quad (12)$$

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

$$f_{\mathcal{D}}(n) = \frac{1}{|\mathcal{D}|} \sum_{i \in \mathcal{D}} \mathbb{1}_{\mathcal{D}}(i, n), \quad (13)$$

where we set $\mathcal{D} = \mathcal{B}$ (\mathcal{B} is the set of banks) or $\mathcal{D} = \mathcal{B} \setminus \mathcal{I}$ (\mathcal{I} is the set of initial defaults, so $\mathcal{B} \setminus \mathcal{I}$ 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 $f_{\mathcal{B} \setminus \mathcal{I}}(n)$ 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,

$$\mathbb{1}_{\mathcal{D}}(i, n) = \begin{cases} 1, & \text{if institution } i \in \mathcal{D} (= \mathcal{B}, \mathcal{B} \setminus \mathcal{I}) \text{ defaults in run } n, \\ 0, & \text{otherwise.} \end{cases}$$

For completeness, the probability of a systemic event \mathbb{P} is given by

$$\mathbb{P} := \frac{1}{N} \sum_{n=1}^N \mathbb{1}_{SE}(n), \quad (14)$$

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.1.2.2).

To interpret our results (see Section 5) 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 11 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 $\mathcal{D} = \mathcal{B} \setminus \mathcal{I}$. When $\mathcal{D} = \mathcal{B} \setminus \mathcal{I}$, the average extent of a systemic event could also be called ‘the average extent of contagion (in a cascade)’, as Schnabel and Shin (2004), Cifuentes et al. (2005), Gai and Kapadia (2010), Caccioli et al. (2014), Paulin et al. (2018) refer to the systemic risk measure \mathbb{E} .

A.1.4 Detailed Model Specification

A.1.4.1 Specification of Contracts

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

We assume that an individual repo contract R_{ji} can only be collateralised by one type of collateral s_{ijam}^e and be supplemented by cash collateral C_{ij}^e , which receives no haircut. s_{ijam}^e denotes the units of tradable assets m of type $a \in \mathcal{A}$ posted as collateral by institution i to institution j , and C_{ij}^e stands for the amount of cash collateral posted by institution i to institution j . The 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.

Whenever the price p_{am} of the asset collateral s_{ijam}^e falls, the ‘haircutted collateral’

(the value of the collateral after the application of the haircuts) may no longer be sufficient to cover the size of the repo loan R_{ji} . When this happens institution $i \in \mathcal{F}$ (where \mathcal{F} is the set of financial institutions) receives a margin call M_{ji}^t from institution $j \in \mathcal{F}$, where the margin call is defined as

$$M_{ji}^t := R_{ji}^t - (1 - h_{am}^t) s_{ijam}^{e,t-1} p_{am}^t + C_{ij}^{e,t-1} \quad (15)$$

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

obliging institution $i \in \mathcal{F}$ to place extra ‘unencumbered’ (u) asset collateral s_{iam}^u or cash collateral C_i^u ⁷⁹ 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 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 (see details in Appendix A.1.4.2). Since an institution i could have multiple repo contracts R_{ji} , it may 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 matched 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 and Muir (2016)), which we do not capture. In our model (see Section 4.1), 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.1.2.1), so is 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} .

⁷⁹We 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: $C_i = C_i^u + \sum_{j \in \mathcal{F}} C_{ij}^e$.

In line with [Bookstaber, Paddrik and Tivnan \(2014\)](#), we assume that a bank provides reverse repo funding to hedge funds (see details in [Appendix A.1.2.1](#)), 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 4.2.3](#)) 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 $h_{a_4,m}^t = 4\%$, 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 and Metrick \(2009\)](#) and shown by [Brunnermeier and Pedersen \(2009\)](#) to be an additional driver of margin calls-induced liquidations.

A.1.4.2 Specification of Behaviour Here we provide further details on our implementation of the behavioural building block (see [Section 4.3](#)) focussing on the ways in which banks (and non-banks) can act to alleviate their binding constraints.

Alleviating a Binding Risk-Weighted Capital Ratio As explained in [Section 4.3.1](#), 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 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 discussed in [Section 4.3.3](#), we assume that the bank returns to target by reducing the highest risk-weighted assets first, as this is the most effective way to quickly get back to the capital ratio target ρ_i^T .⁸⁰ 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

$$r_{ip_4} = \frac{1}{\omega_{p_4}} \left[\sum_{p \in \mathcal{P}} \omega_p A_{ip} - \frac{\tilde{E}_i}{\rho_i^T} \right], \quad (16)$$

and follows from

⁸⁰We 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. \(2019\)](#).

$$\rho_i^T = \frac{\tilde{E}_i}{\omega_{p_4}(A_{ip_4} - r_{ip_4}) + \sum_{p \in \mathcal{P} \setminus p_4} \omega_p A_{ip}}. \quad (17)$$

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

$$r_{ip_6} = \frac{1}{\omega_{p_6}} \left[\sum_{p \in \mathcal{P}} \omega_p A_{ip} - \sum_{p=p_4} \omega_p A_{ip}^u - \frac{\tilde{E}_i}{\rho_i^T} \mathbb{1}_{\{\hat{r}_{ip_4} < r_{ip_4}\}} \right] \quad (18)$$

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, 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

$$r_{ip_8} = \frac{1}{\omega_{p_8}} \left[\sum_{p \in \mathcal{P}} \omega_p A_{ip} - \sum_{p=p_4, p_6, p_5, p_7} \omega_p A_{ip}^u - \frac{\tilde{E}_i}{\rho_i^T} \mathbb{1}_{\{\hat{r}_{ip_x} < r_{ip_x}, \text{ for } x = 4, 6, 5, 7\}} \right] \quad (19)$$

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

$$\frac{\tilde{E}_i}{\sum_{p \in \mathcal{P}} \omega_p A_{ip} - \sum_{p=p_4, p_6, p_5, p_7} \omega_p A_{ip}^u} < \rho_i^T. \quad (20)$$

Alleviating a Binding Leverage Ratio As explained in Section 4.3.1, a bank $i \in \mathcal{B}$ returns to its leverage target λ_i^T whenever its leverage ratio λ_i (defined in equation 2 as the bank's CET1 equity \tilde{E}_i ⁸¹ over its asset exposure \hat{A}_i ⁸² falls below its buffer value

⁸¹In 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 is not equal to it due to a variety of regulatory deductions. With this approximation, we assume that the difference between the equity E_i and the CET1 equity \tilde{E}_i is constant over time.

⁸²As 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

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

$$d_i = [\hat{A}_i \frac{1}{\lambda_i^T} - \tilde{E}_i] \mathbb{1}_{\{\lambda^M < \lambda_i \leq \lambda_i^B\}}, \quad (21)$$

which follows from

$$\lambda_i^T = \frac{(A_i - d_i) - (L_i - d_i) - \Delta_i^{t_0}}{A_i - d_i} \mathbb{1}_{\{\lambda^M < \lambda_i \leq \lambda_i^B\}} = \frac{\tilde{E}_i^t}{A_i - d_i} \mathbb{1}_{\{\lambda^M < \lambda_i \leq \lambda_i^B\}}. \quad (22)$$

A bank $i \in \mathcal{B}$ delevers d_i amount by liquidating d_i amount of assets according to the ‘leverage pecking order’ given in Section 4.3.3, and by using the cash raised to proportionally pay back liabilities.

Alleviating a Binding LCR As explained in Section 4.3.1, a bank $i \in \mathcal{B}$ returns to a LCR target Λ_i^T if its LCR Λ_i (defined in equation 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 numerator of the LCR Λ_i), rather than reducing net outflows Θ_i) (i.e. 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

$$q_i = \{\Lambda_i^T \Theta_i - Q_i\} \mathbb{1}_{\{\Lambda_i < \Lambda_i^B\}} \quad (23)$$

which follows from

$$\Lambda_i^T = \frac{Q_i + q_i}{\Theta_i} \mathbb{1}_{\{\Lambda_i < \Lambda_i^B\}}. \quad (24)$$

A bank $i \in \mathcal{B}$ liquidates q_i amount of non-HQLA assets according to the ‘LCR pecking order’ given in Section 4.3.3

Fulfilling a Margin Call We now discuss how a repo party $j \in \mathcal{F}$ meets a margin call M_{ji}^t (defined in equation 15) issued by a reverse repo part $i \in \mathcal{F}$. We start with explaining the case where the margin call is positive (i.e. $M_{ji}^t > 0$). As explained in Section 4.2.3, a margin call may be met with the same type of collateral that is already

leverage exposure at time zero $\hat{A}_i^{t_0}$ is given by data as $\hat{A}_i^{t_0} = \frac{\tilde{E}_i^{t_0}}{\lambda_i^{t_0}}$.

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

$$s_{ijam}^{e,t,E} = \min\left\{\frac{M_{ji}^t}{(1 - h_{am}^t)p_{am}^t}; s_{iam}^{u,t-1}\right\} \mathbb{1}_{\{M_{ji}^t > 0\}}, \quad (25)$$

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

$$C_{ij}^{e,t,E} = \min\{\max\{M_{ji}^t - s_{ijam}^{e,t,E}(1 - h_{am}^t)p_{am}^t; 0\}; C_i^{u,t-1}\} \mathbb{1}_{\{M_{ji}^t > 0\}}, \quad (26)$$

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

$$l_i^t = \max\{M_{ji}^t - s_{ijam}^{e,t,E}(1 - h_{am}^t)p_{am}^t - C_{ij}^{e,t,E}, 0\} \mathbb{1}_{\{M_{ji}^t > 0\}}.^{83} \quad (27)$$

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 4.2.3). 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 and 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

$$C_{ji}^{e,t,R} = \min\{C_{ij}^{e,t-1}; |M_{ji}^t|\} \mathbb{1}_{\{M_{ji}^t < 0\}}. \quad (28)$$

⁸³Institution $i \in \mathcal{F}$ may also liquidate slightly more assets than l_i^t to take any potential liquidation cost into account.

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

$$s_{ijam}^{e,t,R} = \frac{\max\{|M_{ji}^t| - C_{ij}^{e,t,R}, 0\}}{1 - h_{am}^t} \mathbb{1}_{\{M_{ji}^t < 0\}}. \quad (29)$$

A.1.4.3 Meeting Share Redemptions Open-ended investment funds are subject to share redemptions by their investors. [Coval and Stafford \(2007\)](#) empirically showed that investment funds $i \in \mathcal{M}$ (specifically equity funds) tend to experience investment inflows or outflows (i.e. redemptions) based on their performance as measured by an investment fund i 's net asset value (NAV) η_i . The NAV of an investment fund $i \in \mathcal{M}$ (see Section 4.1.2 for a balance sheet description) is given by

$$\eta_i = \frac{A_i - L_i}{\sigma_i} = \frac{E_i}{\sigma_i}. \quad (30)$$

The performance of an investment fund $i \in \mathcal{M}$ 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 investment fund $i \in \mathcal{M}$ as

$$\chi_i^t = \frac{\eta_i^{t_0} - \eta_i^t}{\eta_i^t}. \quad (31)$$

In line with the empirical evidence of [Coval and Stafford \(2007\)](#), we simply assume that the investment fund investors redeem shares proportional to the relative loss of their NAV χ_i^t .⁸⁴ That is, the cumulative fraction of the original number of investment funds shares $\sigma_i^{t_0}$ that is withdrawn up to time t , f_i^t is given by

$$f_i^t = \chi_i^t \quad (32)$$

The investment fund has the obligation to pay back the shares that are redeemed at their prevailing NAV η_i . If the investment fund does not have enough cash C_i to do so it must liquidate tradable asset T_i . We assume that it does so proportional to its holdings. This can give rise to a contagious ‘firesale loop’, as discussed in Section 5.3.

⁸⁴We model the redemptions pressures in a simplistic way to be able to make a point that investment funds can affect banking sector stability (see Section 5.3), but given our stress test framework we could easily implement investment funds in a more involved way (e.g. along the lines of [Baranova et al. \(2017\)](#)).

A.2 System-Wide Stress Testing Software

We have 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 we 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 now discuss in more detail the design principles for robust system-wide stress testing code noted in Section 3.3.

A.2.1.1 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.⁸⁵ 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 A.2.2).⁸⁶ Further, to give a broad overview of the structure of the code Figure 9 displays the class diagram of the code.

⁸⁵Alternatively, 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. Dawid et al. (2016)).

⁸⁶The link to the Economic Simulation Library is given by: <https://github.com/ox-inet-resilience/py-distilledESL>.

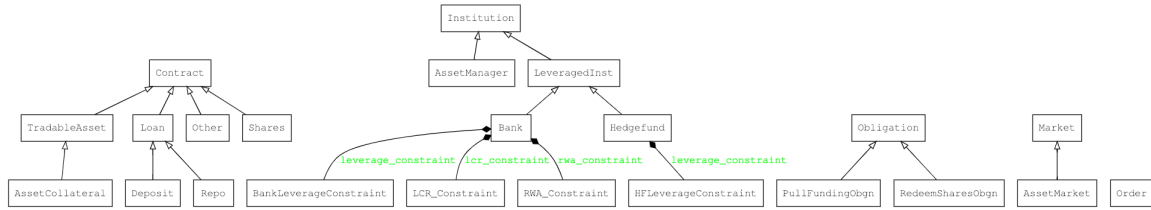


Figure 9: A system-wide stress test consists of five building blocks and so does its code. We have three main classes: First, *institutions*. The different institution types (e.g. bank, investment funds, and hedge funds) inherent from the parent class institutions. Since (regulatory) *constraints* are typically institution-type specific these inherent 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.

A.2.1.2 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.

A.2.1.3 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 3.1) 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 component 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 3.1.4) 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 and 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.

A.2.1.4 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.2.1.3 on modularity).

A.2.1.5 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.2.1.4). 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.⁸⁷

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 institution 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.⁸⁸ 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 once (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

⁸⁷Amazon 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>

⁸⁸See abcEconomics for how to do this in Python: <https://github.com/ab-ce/abce>.

subset of Numpy library has various performant back-end choices, such as LAPACK, ATLAS, BLAS and OpenBLAS.⁸⁹ 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.⁹⁰ 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.

A.2.1.6 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 .⁹¹ 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.2.1.4 on readability).

A.2.2 Synchronicity in Financial Markets using a Messaging-Mailbox System

The code underpinning a system-wide stress test must retain the concurrency of financial markets; in financial markets institutions may act (nearly) simultaneously. Code where institutions act sequentially, gives certain institutions a model-induced systematic advantage that would not exist in reality. To partially address such an artificial advantage, the order in which institutions act could be randomly shuffled around (see eg Figue (2017)). Although this takes away the time-independent systematic advantage that pertains to

⁸⁹See: <https://docs.scipy.org/doc/numpy/user/building.html#prerequisites>.

⁹⁰See: <https://keras.io/backend/>.

⁹¹It 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.

certain institutions, it does not take away the artificial advantage granted to certain institutions in a given time step. Hence, more rigorous methods to solve order dependence are needed. One method is to create parallel computer code to run a system-wide stress test. As parallel code tends to be prone to errors, in our illustrative implementation of the five building blocks for stress testing we introduce a simpler way. Institutions are given a ‘mailbox’. Whenever an institution acts (e.g. pulls funding, gives a margin call) the notification of the action first ends up in the ‘unread mailbox’ of a counterparty. Only after every institution in a given time step has acted will the ‘message’ move to the ‘read mailbox’. As such all interactions among institutions in a given time step manifest at once, as if every institution that acted in the time step did so simultaneously. Likewise, 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 the sequence of the acting institutions in each time step.⁹²

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.⁹³ When we use the random shuffling, we find that the standard deviation of the average extent of systemic event \mathbb{E} (see Appendix A.1.3.1) 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,

A.3 Supplementary Figures and Tables

A.3.1 From Micro to Macro: A Macroprudential Overlay to the EBA 2018 Stress Test Figures 10 and 11 show the impact of the initial adverse shock on systemic risk, in the cases where we relax the leverage ratio and risk-weighted capital ratio constraints respectively. Table 5 gives the summary statistics of the capital ratios of banks that partook in the 2018 EBA stress test.

⁹²An 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>.

⁹³See: https://github.com/ox-inet-resilience/firesale_stresstest/blob/master/other_simulations/random_shuffling.py.

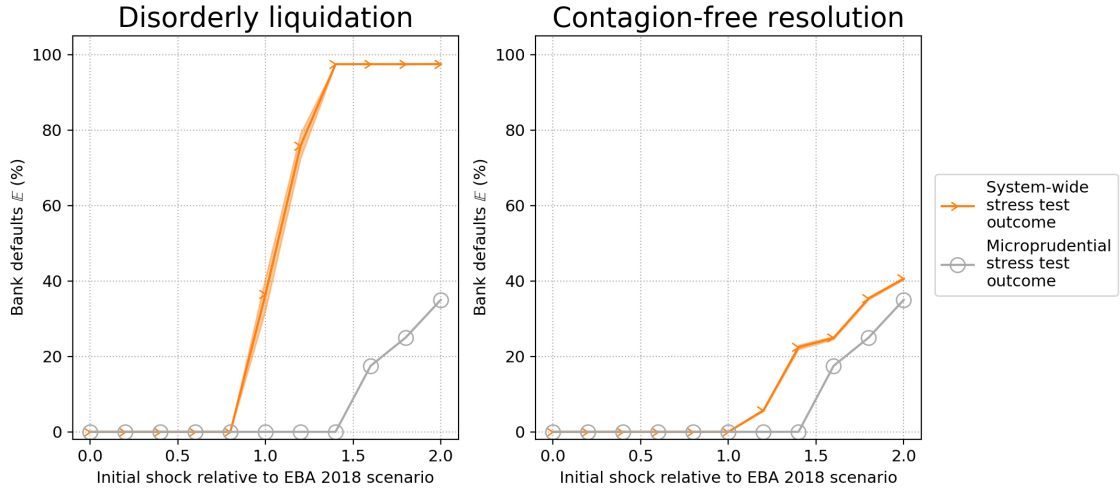


Figure 10: Shows the same set-up as in Figure 3, except that we have now turned the leverage ratio λ_i off. This means that the leverage minimum λ_i^M , buffer λ_i^B and target λ_i^T do not apply. We observe that the financial system remains stable for a much more severe initial scenario when the leverage ratio λ_i is off compared to when it is on (in Figure 3). Together with the results of Figure 11, this indicates that the Basel III leverage constraint would be more binding for European banks than the risk-weighted capital ratio, and so a greater driver of potential instability.

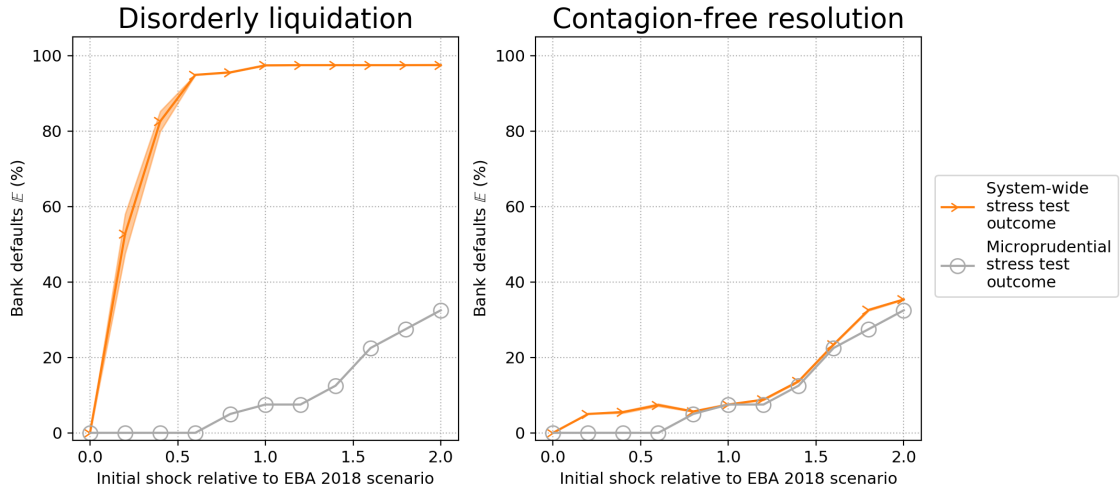


Figure 11: Shows the same set-up as in Figure 3, except that we have now turned the risk-weighted capital ratio ρ_i off. This means that the risk-weighted capital ratio minimum ρ_i^M , buffer ρ_i^B and target ρ_i^T do not apply. We observe that this result is almost identical to the result when the risk-weighted capital ratio ρ_i is turned on in Figure 3; reinforcing the point that the risk-weighted capital ratio requirement is relatively less binding than the Basel III leverage ratio requirement for European banks.

		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	$\bar{\lambda}^{data} - \bar{\lambda}^B$	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	$\bar{\lambda}^{data} - \lambda^M$	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%

Table 5: 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.

A.3.2 Amplification of Contagion Mechanisms Figure 12 and Figure 13 show the amplification among contagion mechanisms.

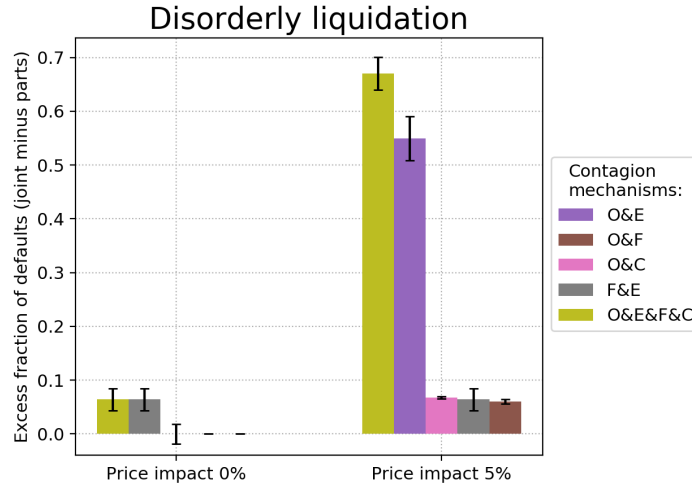
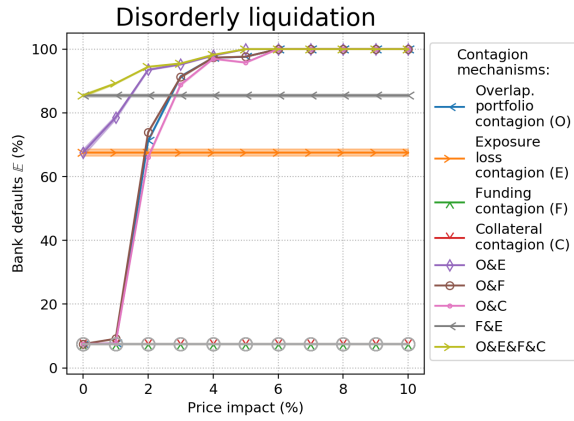
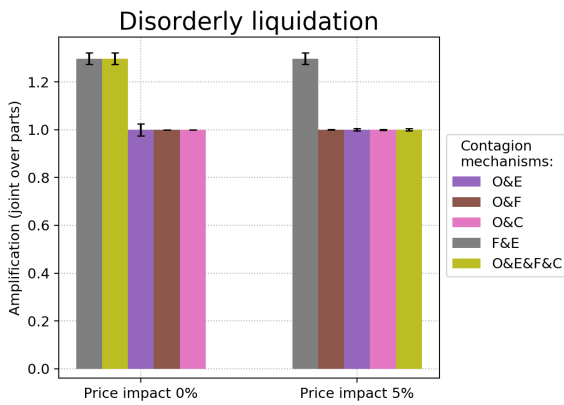


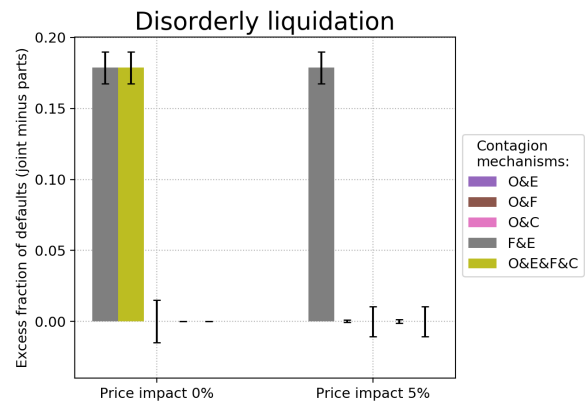
Figure 12: Shows the excess systemic risk \mathbb{E} (given by the joint set of contagion mechanisms *minus* the systemic risk \mathbb{E} of the sum of the individual contagion mechanisms) for various price impacts (PI), for various sets of contagion mechanism and the same set-up as in Figure 6. A positive 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 $\mathbb{E} \approx 65\%$ (see the *O&E&F&C* bar). Importantly, we note that the contagion mechanisms that amplify each other most in relative terms (see Figure 6b) may not be the same contagion mechanisms that amplify each other most in absolute terms (see Figure 12). For instance, overlapping portfolio contagion and collateral contagion (see *O&C* at 5% price impact in 6b) amplify each other most in relative terms, while overlapping portfolio contagion and exposure loss contagion (see *O&E* at 5% price impact in 12) amplify each other most in absolute terms.



(a)



(b)



(c)

Figure 13: Shows the same set-up as in Figure 6, except that now we use the Basel III default settings (see Table 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 13b is often smaller than one, and the excess systemic risk \mathbb{E} in plot 13c is frequently negative.

A.3.3 ‘Usability’ of Buffers and Contagion Figures 14 and 15 show the impact of the ‘usability’ of buffers on financial stability. Figure 16 shows how the level of capital ratio target affects financial stability.

Disorderly liquidation

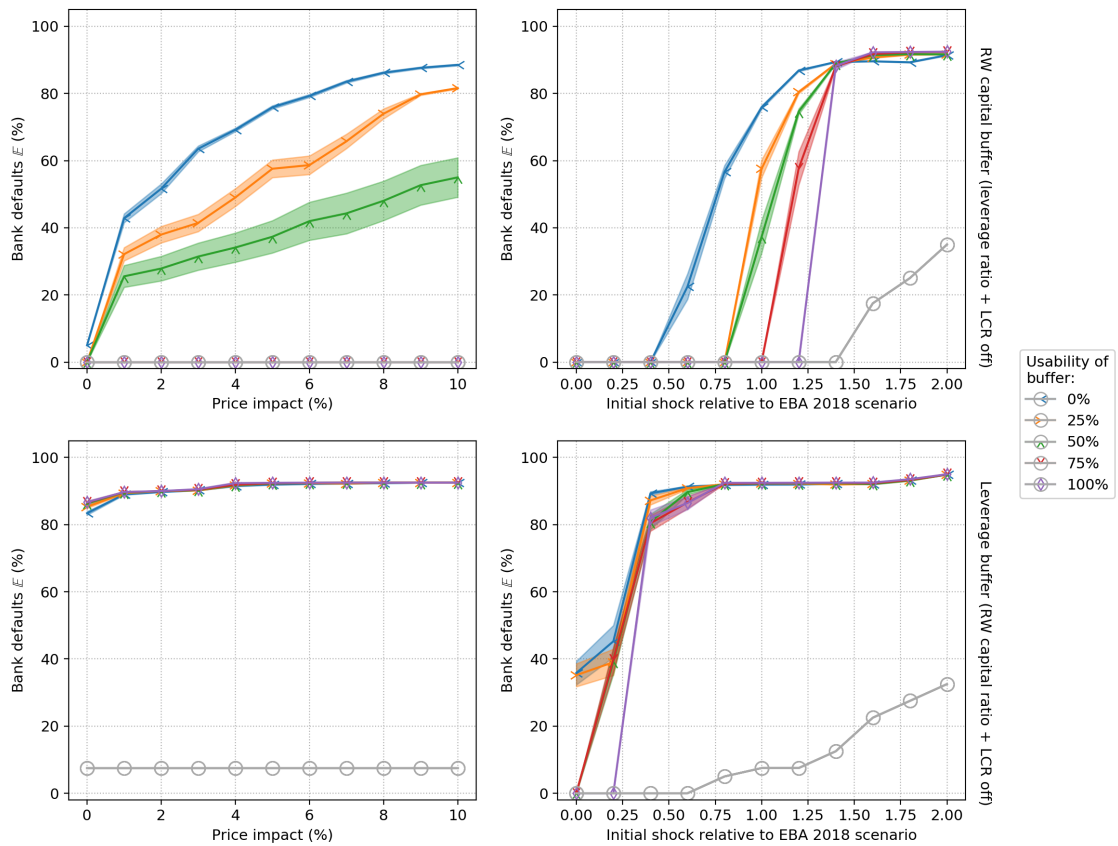


Figure 14: Shows the same set-up as in Figure 7, except now for the case of ‘disorderly liquidation’ the top row shows the effect of the usability of the regulatory risk-weighted capital buffer ρ_i^{CB} only (i.e. the leverage ratio λ_i and the LCR Λ_i are turned off), and the bottom row shows the effect of the usability of the regulatory leverage buffer λ_i^{CB} only. Adding to the findings of Figure 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’.

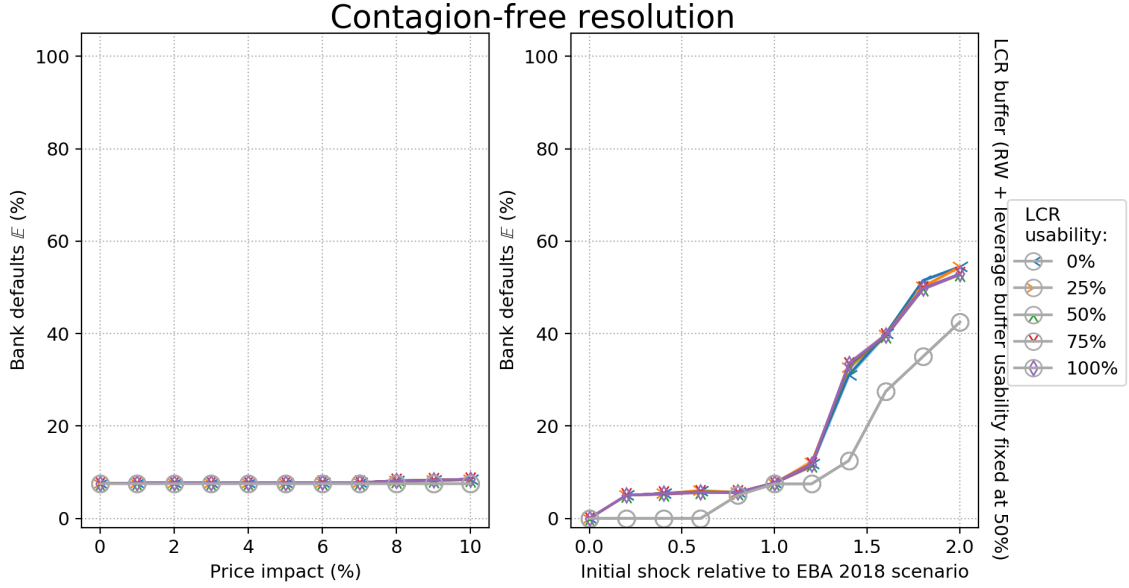


Figure 15: Shows the same set-up as in Figure 7, except that we now fix the usability of the regulatory leverage buffer λ_i^{CB} and the risk-weighted capital buffer ρ_i^{CB} at their default value of $u = 50\%$, and solely vary the usability of the LCR standard Λ^S . We observe that the usability of 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.

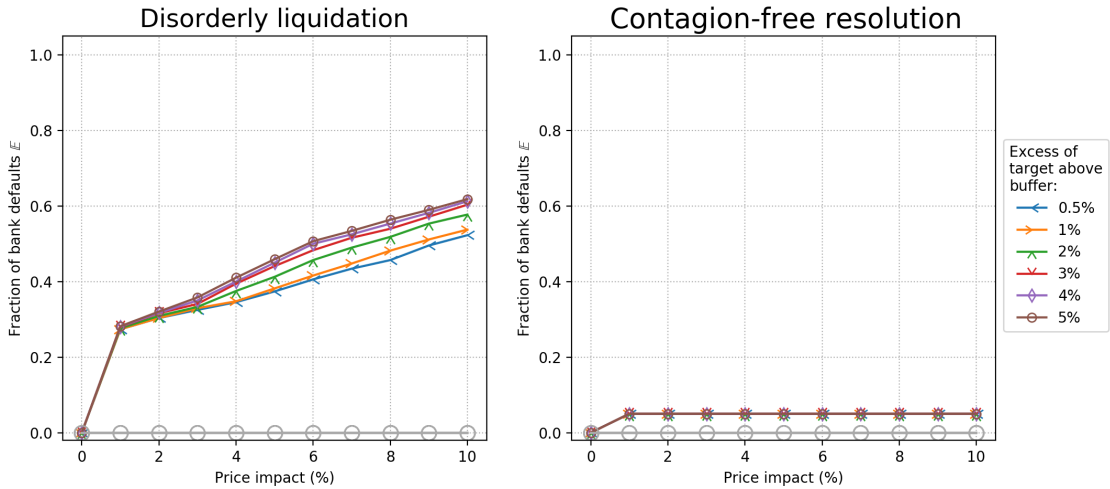


Figure 16: Shows systemic risk \mathbb{E} as a function of the price impact for the case where the regulatory leverage buffer λ_i^{CB} is tripled (i.e. $y^\lambda = 3$), for different excess targets above the buffer (i.e. $\rho_i^B - \rho_i^B = \lambda_i^T - \lambda_i^B = x\%$, for $x = 0.5, \dots, 5\%$, see definitions in Section 4.3.2). The excess target above the buffer indicates by how many percentage points the bank seeks to improve 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 in the case of ‘disorderly liquidation’. In the case of ‘contagion-free’ resolution, under these settings the instability is too small to be affected by an increase in the target - though for different buffer sizes we obtain the same qualitative finding. Hence, individual stability can lead to collective instability.

A.3.4 Calibration of Buffers with System-Wide Stress Tests Figure 17 shows the impact of the regulatory buffer size on systemic risk.

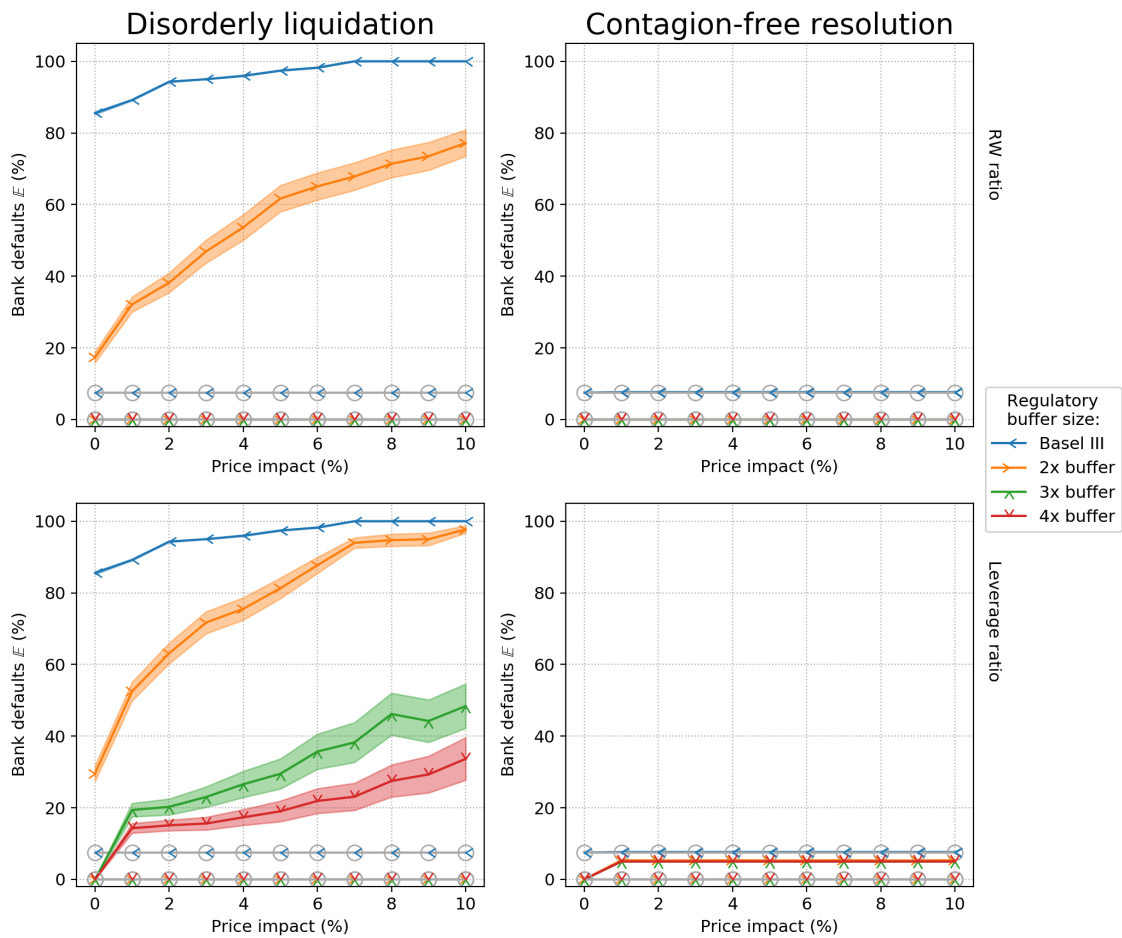


Figure 17: Shows the same set-up as in Figure 8 except we now show how systemic risk \mathbb{E} decreases in the regulatory buffer sizes as a function of the price impact. Also, we add the triple buffer case. By showing systemic risk \mathbb{E} as a function of the price impact, it becomes even easier 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 provide information about necessary buffer sizes according to the microprudential stress test and the coloured lines provide information about necessary 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 17). 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 17 is the first line to fall underneath $\mathbb{E} = 10\%$).