DISCUSSION PAPER SERIES

DP16509
(v. 2)

Systemic Implications of the Bail-In Design

Alissa Kleinnijenhuis, Charles A Goodhart and Doyne Farmer

MONETARY ECONOMICS AND FLUCTUATIONS

ISSN 0265-8003

Systemic Implications of the Bail-In Design

Alissa Kleinnijenhuis, Charles A Goodhart and Doyne Farmer

Discussion Paper DP16509
First Published 03 September 2021
This Revision 06 December 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Systemic Implications of the Bail-In Design

Abstract

The 2007-2008 financial crisis forced governments to choose between the unattractive alternatives of either bailing out a systemically important bank (SIB) or allowing it to fail in a disorderly manner. Bail-in has been put forward as an alternative that potentially addresses the too-big-to-fail problem and contagion risk simultaneously. Though its efficacy has been demonstrated for smaller idiosyncratic SIB failures, its ability to maintain stability in cases of large SIB failures and system-wide crises remains untested. This paper's novelty is to assess the financial-stability implications of bail-in design, explicitly accounting for the multi-layered networked nature of the financial system. We present a model of the European financial system that captures five prevailing contagion channels. We demonstrate that it is essential to understand the interaction of multiple contagion mechanisms and that financial institutions other than banks play an important role. Our results indicate that stability hinges on the bank-specific and structural bail-in design. On the one hand, a well-designed bail-in buttresses financial resilience, but on the other hand, an ill-designed bail-in tends to exacerbate financial distress, especially in system-wide crises and when there are large SIB failures. Our analysis suggests that the current bail-in design may be in the region of instability. While policy makers can fix this, the political economy incentives make this unlikely.

JEL Classification: N/A

Keywords: Too Big To Fail, Resolution, bail-in, financial crisis, contagion, Financial Networks, default, systemically important banks, bail-in debt pricing, political economy

Alissa Kleinnijenhuis - amklein@stanford.edu
Stanford Institute for Economic Policy Research, Stanford University,

Charles A Goodhart - caegoodhart@aol.com
LSE and CEPR

Doyne Farmer - doyne.farmer@inet.ox.ac.uk
University of Oxford

Acknowledgements
We thank Tobias Adrian and Thomas Huertes for their valuable feedback on the draft. We thank Viral Acharya, Ben Bernanke, John Cochrane, Rama Cont, Marie Donay, Darrell Duffie, Mark Egan, Grzegorz Halaj, Andreas Lehnart, Robert Merton, Marina Moretti, Jean-Charles Rochet, Emil Verner, Jeremy Stein, John Taylor, Luitgard Veraart, Michael Vincent, and Jianping Zhou for their helpful discussions. We thank the participants of seminars at the International Monetary Fund, Federal Reserve Board of Governors, European Central Bank, European Commission, MIT Sloan School of Management, Mathematical and Computational Finance Group at the University of Oxford, Institute of New Economic Thinking at the Oxford Martin School, Stanford Graduate School of Business and UK Economic & Social Research Council, as well as the staff from the Single Resolution Board and the European Systemic Risk Board. We thank the participants of the 10th Labex ReFi-NYU-SAFE/LawFin Law & Banking/- Finance

Conference and the 4th Law & Macroeconomics Conference. We are grateful for Rudy Tanin's excellent research assistance. All errors remain ours.

Systemic Implications of the Bail-In Design

Alissa M. Kleinnijenhuis[1,2], Charles Goodhart[3], and J. Doyne Farmer[1,4]

[1]Institute for New Economic Thinking at Oxford Martin School, University of Oxford, UK
[2]Stanford Institute for Economic Policy Research, Stanford University, US*
[3]London School of Economics, UK
[4]Mathematical and Computational Finance Group, University of Oxford, UK

December 2, 2021

Abstract

The 2007-2008 financial crisis forced governments to choose between the unattractive alternatives of either bailing out a systemically important bank (SIB) or allowing it to fail in a disorderly manner. Bail-in has been put forward as an alternative that potentially addresses the too-big-to-fail problem and contagion risk simultaneously. Though its efficacy has been demonstrated for smaller idiosyncratic SIB failures, its ability to maintain stability in cases of large SIB failures and system-wide crises remains untested. This paper's novelty is to assess the financial-stability implications of bail-in design, explicitly accounting for the multi-layered networked nature of the financial system. We present a model of the European financial system that captures five prevailing contagion channels. We demonstrate that it is essential to understand the interaction of multiple contagion mechanisms and that financial institutions other than banks play an important role. Our results indicate that stability hinges on the bank-specific and structural bail-in design. On the one hand, a well-designed bail-in buttresses financial resilience, but on the other hand, an ill-designed bail-in tends to exacerbate financial distress, especially in system-wide crises and when there are large SIB failures. Our analysis suggests that the current bail-in design may be in the region of instability. While policy makers can fix this, the political economy incentives make this unlikely.

Keywords: *Too big to fail, resolution, bail-in, liquidation, insolvency law, financial crisis, contagion, financial networks, failure, default, bail-out, banks, systemically important banks, loss absorption requirements, bail-in debt, bail-in debt pricing, political economy*

*We thank Tobias Adrian and Thomas Huertes for their valuable feedback on the draft. We thank Viral Acharya, Ben Bernanke, John Cochrane, Rama Cont, Marie Donay, Darrell Duffie, Mark Egan, Grzegorz Halaj, Andreas Lehnart, Robert Merton, Marina Moretti, Jean-Charles Rochet, Emil Verner, Jeremy Stein, John Taylor, Luitgard Veraart, Michael Vincent, and Jianping Zhou for their helpful discussions. We thank the participants of seminars at the International Monetary Fund, Federal Reserve Board of Governors, European Central Bank, European Commission, MIT Sloan School of Management, Mathematical and Computational Finance Group at the University of Oxford, Institute of New Economic Thinking at the Oxford Martin School, Stanford Graduate School of Business and UK Economic & Social Research Council, as well as the staff from the Single Resolution Board and the European Systemic Risk Board. We thank the participants of the 10th Labex ReFi-NYU-SAFE/LawFin Law & Banking/-Finance Conference and the 4th Law & Macroeconomics Conference. We are grateful for Rudy Tanin's excellent research assistance. All errors remain ours.

Contents

1 Introduction

The failure of Lehman Brothers was the defining event of the 2007-2008 financial crisis, bringing the financial system and the real economy to the brink of abyss (Bernanke (2017)). Policy makers were forced to choose between the unattractive alternatives of either bailing-out a systemically important institution or allowing it to fail in a disorderly manner, threatening the stability of the entire financial system (French et al. (2010), Bernanke (2010)). Ordinary bankruptcy procedures at the time were entirely inadequate for dealing with failures of systemically important institutions (Bernanke (2017)). Since then, many have recommended that authorities be given the necessary powers resolve systemically important institutions in an orderly manner (French et al. (2010)). This recommendation has been adopted in many countries. In the European Union, the Bank Recovery and Resolution Directive (BRRD) establishes a common approach to the recovery and resolution of banks and investment firms.[1] In the United States, Title II of the Dodd-Frank Act (DFA) provides a provision for the Orderly Liquidation of large, complex financial institution.

Bail-in has been widely hailed as the primary tool for resolving a failing systemically important bank (SIB) (e.g. FSB (2013), Chennells and Wingfield (2015), BoE (2017)). It potentially ends the "too-big-to-fail" problem by letting investors shoulder the losses – addressing moral hazard – while minimizing the calamitous impact of a bank's failure on the economy and the financial system (FSB (2013)). Bail-in writes down a subset of a bank's debt and converts (part of) it to equity.

The efficacy of bail-in in severe crises scenarios remains an open question. While bail-in has proven relatively successful in dealing with the failure of relatively minor systemically important banks, its ability to handle large bank failures and system-wide crises has not yet been tested. Bail-ins (of sorts) on comparatively small SIBs have been carried out successfully in Cyprus, the Netherlands, Italy, Denmark and the United Kingdom (WBG (2017)).[2] In some cases this required complementary bail-out funds and other measures. For instance, in Cyprus deposits had to be frozen to avoid a bank run (Moulds et al. (2013)). In other cases, such as in Italy (Merler (2016)), policy was distorted due to the widespread ownership of subordinated bank debt by households. Notwithstanding the success in dealing with small SIB failures, leading experts remain skeptical whether bail-in can be effective in a major crisis. Bernanke (2017) notes that: "controversies remain over how effective even a Title II resolution would be in the context of a significant financial crisis". Avgouleas and Goodhart (2015) argue that "the bail-in approach may, indeed, be much superior to bailouts in the case of idiosyncratic failure. In other cases, the bail-in process may entail important risks".

While a COVID-19 financial crisis has been narrowly averted, regulators should be on guard for possible future crises, such as one brought about by disruptions due to climate change. The possibility of enacting improved bail-in policies before a next financial crisis occurs is one of the motivations of this paper. As we show here, whether or not a bail-in is beneficial in a major crisis hinges on the details of the bail-in design. Implementation of the right bail-in policies, and the ability to test them before they occur in real life, could also have the advantage of helping to give regulators the confidence to carry out such policies in times of severe stress rather than falling back on bail-outs.

The novel contribution of this paper is to investigate the systemic implications of the bail-in

[1] The BRRD also establishes the resolution approach in the United Kingdom as it used to be a member state of the European Union.

[2] Only one failing SIB, Banco Popular of Spain, has been resolved under the BRRD so far. The bank failures mentioned here effectively used bail-ins, but these were outside of the approach established by the BRRD.

design, explicitly taking into account the financial system as a whole. For this, we use as a case study the European financial system in which the BRRD applies. Since key design parameters for European bail-ins are typically also key parameters in bail-in designs in other jurisdictions, such as the United States, we expect our results to be relevant across jurisdictions even though details of the bail-in process might differ per jurisdiction.

To evaluate the systemic implications of the applicable bail-in design, we build on the multi-layered network model of the European financial system developed by Farmer et al. (2020). We calibrate this model using *S &P Global Market Intelligence* data and the 2018 European Banking Authority (EBA) disclosures of the stress test results. In line with Greenwood et al. (2015) and Duarte and Eisenbach (2021), our simple and easy-to-calibrate model takes as given: (1) the asset holdings, as well as liabilities and equity, of each financial institution; (2) an adjustment rule applied by institutions when they are hit by adverse shocks; and (3) the price impact of liquidating assets. Our framework extends those developed by these authors by not only capturing fire sale contagion that may take place in the financial network via indirect linkages of common security holdings, but also capturing contagion that may occur via direct linkages of bail-in debt and non-bail-in debt contracts. By incorporating the direct and indirect linkages among institutions jointly, the calibrated system-wide stress testing model can capture the prevailing contagion mechanisms that could endogenously amplify shocks emanating from bail-ins, i.e. exposure loss contagion, overlapping portfolio contagion, funding contagion, bail-in debt revaluations, and bail-in debt runs originating from concerns over expected losses, uncertainty, or similarity to a recently failed bank.

The interactions among these contagion mechanisms arise as follows. A bank bail-in typically results in exposure losses for some of its creditors. *Ceteris paribus*, losses to bailed-in creditors increase their probability of default and decrease their bail-in debt valuation. In turn, their investors may turn unwilling to roll-over maturing bail-in debt to them, since the expected losses on the creditors' bail-in debt have increased now that their probability of default is higher. These creditors then have the obligation to repay the invested funds using their cash buffer or cash raised from selling securities or from pulling back its maturing funding to others. Asset sales can trigger overlapping portfolio contagion, where mark-to-market losses on commonly held securities precipitate further liquidations of securities. Pulling loans can trigger funding contagion, where a halt to refinancing prompts the counterparty to stop rolling over its funding to its counterparties in turn. Alternatively, their investors may decide to cut their exposure to such bail-in debt if they face uncertainty over the losses they may suffer in an impending bail-in or if the bank looks similar to a bank that has just failed. While rolling-over bail-in exposures based on the criteria of expected losses requires investors to be able to rely on accurate pricing of bail-in debt, cutting exposures based on a rule of thumb, such as similarity to a just-failed bank, does not. If the resolution authority holds plenty of discretion over the design of a prospective bail-in – rendering the accurate pricing of bail-in debt difficult – then investors have to resort to a rule of thumb for rolling over exposures. This makes it more likely that the prospect of a bail-in creates a financial panic undermining stability. Taken together, exposure losses from a bail-in can amplify overlapping portfolio and funding contagion which in turn can prompt destabilising de-levering and more bail-ins repeating the cycle.

Our results suggest that *financial stability hinges on the bail-in design* – comprised of what we coin the "bank-specific" and "structural" design parameters. Bank-specific features of the design are those that vary depending on the bank and circumstance in question. We discuss these in the temporal sequence in which they occur. These include the failure threshold, *ad-hoc* debt exclusions from the bail-in process, the recapitalisation target, and debt-to-equity conversion rates. In contrast, the structural aspects of the design are set in advance and are typically common among a given class of banks under the BRRD. These include the requirements on loss-absorbing debt, *a-priori* exclusions of debt from the application of bail-in and uncertainty in the bail-in design. We focus on these key design parameters, since they are typically present in each jurisdiction's bail-in framework, and are most salient to the bail-in design.

Our results suggest that a well-designed bail-in buttresses financial resilience and does not tend to amplify existing financial distress. A well-designed bail-in consists of strong recapitalisation, a "fair" distribution of debt-to-equity conversion rates, high loss absorption requirements, *a-priori* exclusions of short-term debt (i.e. with a time to maturity less than one year) and sufficient certainty about the bank-specific bail-in design – and is done early.

On the other hand, we show that an ill-designed bail-in tends to strongly exacerbate financial distress, especially in system-wide crises and with large SIB failures. It consists of the opposite of a well-designed bail-in, with the exception that only debt with a time to maturity less than seven days to institutions is excluded *a priori*. In line with empirical evidence, we find that the bail-in design matters less for financial stability in the case of smaller SIB failures.

In the case of a severe system-wide crisis, we estimate that the difference in contagious asset losses between a resolution regime with a well-designed and ill-designed bail-in could run up to 4.2 trillion euros, representing 20.3% of asset value of the European banking sector participating in the 2018 EBA stress test (and, for comparison, 22.3% of European GDP as of 2018).[3] In the case of a idiosyncratic failure of the largest European bank by asset size, this difference is estimated to be 2.4 trillion euros (11.6% of these banking sector assets), whereas in the case of a idiosyncratic failure of the smallest European bank by asset size participating in the 2018 EBA stress test, the difference in contagious asset losses is negligible. These quantitative findings suggest that the welfare implications of the choice of bail-in design are significant, especially in times of severe financial crises and large idiosyncratically-failing SIBs. Numerous empirical studies have found that large banking sector losses give rise to contractions in credit supply, thereby depressing macroeconomic output and creating a recession (see e.g. Bernanke (1983) and Reinhart and Rogoff (2009)).[4]

The intuition for why a well-designed bail-in supports stability, even in a financial crisis, whereas an ill-designed bail-in does not, is that the first breaks the de-leveraging cycle, whereas the latter does not. A financial crisis typically unfolds because of a de-leveraging cycle, in which asset losses (e.g. from real-economy exposures or an idiosyncratic bank failure) and a perceived increase in risk, trigger some financial institutions to de-lever (Adrian and Shin (2014)) triggering others to de-lever in turn, thereby repeating the cycle and amplifying system-wide losses. Losses brought about by this dynamic create runs on institutions who have possibly turned insolvent (Brunnermeier (2008)). A well-designed bail-in breaks the de-levering cycle and prevents runs on nearly failing banks. It does so by rapidly bringing down the leverage of nearly failing banks, by recapitalising them strongly. Further, by allowing only long-term debt to count as bail-in debt, and shielding other debt from losses, it also prevents bank runs in anticipation of a bank failure. Moreover, early bail-ins reduce net exposure losses that creditors suffer resulting from bail-ins, reducing de-levering pressures in the first place. Finally, with sufficient certainty in the applicable bail-in design, investors can properly incorporate expected bail-in losses into the price of bail-in debt reducing the risk of panics in anticipation of multiple bail-ins. An ill-designed bail-in at best

[3]Note that this includes the GDP of the UK as it was still a European member state at the time.

[4]Our results also shed light on how the bail-in design affects the no-creditor-worse off (NCWO) principle. This states that resolution authorities should seek to ensure that no creditor or shareholder is expected to incur greater net losses than it would have incurred in winding up the bank under normal insolvency proceedings (see: Article 74(2) of the BRRD). Shareholders or creditors for whom this does not hold are entitled to a difference payment by the resolution financing arrangements (see: Article 75 of the BRRD).

only partially arrests the de-leveraging cycle and at worst may even accelerate it – by doing the opposite of what a well-designed bail-in does.

We further show that the systemic footprint of a bail-in design will not be properly understood without taking account of multiple contagion mechanisms, which can substantially amplify system-wide losses (in a de-leveraging cycle), and the role of non-banks, who hold much of the bail-in debt.[5] We document that multiple poor design choices amplify contagious asset losses. With an overall well-designed bail-in the difference in contagious asset losses in a severe crisis between a resolution regime where bail-in debt is long-term (good design parameter) and one where bail-in debt can be short-term (poor design parameter) is estimated to be 6.7% of banking sector assets, whereas with an overall poorly-designed bail-in this difference is 19.8%. Hence, an additional poor design choice on top of an already poorly designed bail-in amplifies losses by a factor of 3 (for this design parameter) compared to an additional poor design choice on top of a better-designed bail-in. The increase in systemic risk brought about by a poor design choice thus depends on how well the current bail-in mechanism is designed. Connectedly, we document that the degree of amplification among contagion mechanisms is greater with poorer bail-in designs. As a whole, our set of findings is qualitatively robust to an extensive sensitivity analysis.[6] Our results shed light on how a change in parameters of the bail-in design improves or weakens the resilience of the financial system.

Our work suggests that a pivot towards greater stability is in the hands of policymakers. On the positive side we find that a well-designed bail-in can work even in a system-wide crisis, thereby providing evidence that could embolden resolution authorities to opt for well-designed bail-ins rather than bail-outs in a next systemic crisis. On the negative side, our results suggest that the current bail-in design may be in the instability regime, and could make a major financial crisis even worse. Unfortunately, the political economy incentives are such that an improvement in the bail-in design seems unlikely (see Section 6).

Links to the literature. Our paper helps answer the long-standing question of how we can solve the too-big-to-fail problem. Our contribution adds to the literature that discusses, measures, and seeks to address this problem (see e.g. O'hara and Shaw (1990), Acharya and Beck (2014)). Stern and Feldman (2004) warn ahead of the Great Financial Crisis that not enough had been done to reduce creditors' expectations of the too-big-to-fail protection. The too-big-to-fail problem arises because regulatory authorities believe that a bank's failure could impose severe negative externalities upon other actors. Whether a bank is deemed too-big-to-fail does not just depend on its size, but notably also on its interconnectedness, which determines how easily its failure triggers widespread financial contagion with potentially harmful consequences for the real economy (Stein (2013), Bernanke (2016)). If creditors expect too-big-to-fail protection, they reduce their vigilance in monitoring and responding to banks' activities, thereby exerting less market discipline. The lack of market discipline gives incentives for *ex ante* excessive risk taking by banks in the presence of moral hazard, and ultimately risks making the financial system more fragile. "Heads, I win; tails, you lose" always leads to excessive risk taking and systemic risk (Jackson et al. (2015)). *Ex-ante* incentive distortions brought about by reduced risk-sensitivity of capital costs and *ex-post* bail-out

[5]In part because cross-holdings of bail-in debt among banks are discouraged by the eligibility rules for the loss absorption requirements.

[6]For greater confidence in our quantitative estimates, our model should be calibrated to granular data on the multi-layered financial network, including the non-bank holdings of bail-in debt, which is currently typically not known to regulators. We offer the headline numbers as differentials between bail-in designs rather than absolute numbers to show how the bail-in design matters given the adopted calibration to data. We stress that we do not offer an 'optimal' bail-in design, for this requires linking our financial system model with a real economy model to evaluate trade-offs between financial stability and economic growth, as well as more granular data.

costs comprise policy costs of the too-big-to-fail problem (Morrison (2011)).[7]

Calomiris and Herring (2013) discuss how contingent convertibles (CoCos), which upon a *contractually* agreed trigger can be written down or converted to equity, and which have been issued by many banks, can be designed to address the too big to fail problem. Calello and Ervin (2010) are the first to propose bail-in, a *statutory* power of the regulator, to curtail too big to fail. Bail-in is now a core part of the resolution toolkit in many jurisdictions and is the main regulatory tool introduced in the wake of the Great Financial Crisis to resolve too big to fail. Bernanke (2016) argues that the introduction of a resolution authority with bail-in powers raises the probability that creditors will take losses with a SIB failure, thereby reducing the perceived benefits of the too-big-to-fail status. Berndt et al. (2020) report a post-crisis decline of too big to fail for globally SIBs with U.S. headquarters, as suggested by a reduction in market-implied probabilities of government bailout, along with higher cost of capital after controlling for insolvency risk. They argue that this decline is consistent with the post-crisis introduction of a resolution authority. Our paper contributes to the literature on too big to fail by showing quantitatively that the principal post-crisis tool proposed to end to big to fail, bail-in, does not resolve this problem under its current design, since it tends to exacerbate existing distress in a financial crisis. We do find that bailing-in large SIBs is less disruptive than liquidating these and bail-ins of small SIBs work well, regardless of the chosen design parameters, which perhaps dovetails with the findings of Berndt et al. (2020). We further add to the too-big-to-fail literature by, perhaps surprisingly, demonstrating that an improved bail-in design might go a long way in eliminating too big to fail, since a well-designed bail-in (largely) prevents contagion associated with SIB failures.[8]

More narrowly our paper adds to the nascent literature on the systemic effects of bail-in in an interconnected financial system represented as a financial network. This literature aims to quantitatively assess how bail-in affects failure externalities, such as contagion. Klimek et al. (2015) employ a financial network model to evaluate the economic and financial ramifications of bail-in. Hüser et al. (2017) evaluate the systemic implications of bail-in in the European Union, drawing on a calibrated multi-layered network model of bank debt and equity cross holdings. Bernard et al. (2017) investigate how the strategic negotiation with the regulator affects whether banks are willing to participate voluntarily in a bail-in. These papers neither investigate the systemic impact of the bail-in design nor include the prevailing contagion mechanisms and non-banks in their analysis, due to which they underestimate the systemic instability particular bail-in designs can generate.[9]

We assess the systemic implications of the bail-in design by representing the financial system as a network to reflect the high degree of interdependence exhibited by modern financial systems and to capture contagious interactions that may emerge from these interconnections, rendering the use of a financial network model for our economic analysis a natural choice (Allen and Babus (2009)).

[7] Substantially increasing capital requirements alleviates the too-big-to-fail problem (see e.g. Morrison (2011), Admati and Hellwig (2014), Cochrane (2014)), but has remained impracticable due to both heavy opposition by banks enjoying a tax shield on debt and fears by regulators that banks would reduce their real-economy lending.

[8] A recently proposed alternative to bail-in as a mechanism to help end the too big to fail problem is a firm failure mechanism making use of no-fault-default debt (Merton and Thakor (2021)). When a bondholder demands payment associated with no-fault-default debt at maturity, the company can choose to make the payment or surrender equity in the company. This enables the firm to transform its debt claims into equity claims upon default on the debt, thus allowing a reallocation of control rights to bondholders with minimal disruption of the business operations of the firm. Unlike bail-in, which is a core tool for resolving SIBs at present, no-fault-default debt is currently not employed.

[9] Unlike the before-mentioned quantitative papers on the topic which describe the bail-in process largely verbally, our paper enhances the understanding of the bail-in mechanism by specifying the bail-in design in precise formulae, including important legal safeguards such as the NCWO criterion.

Our paper also contributes qualitatively and quantitatively to a conceptual literature that critically evaluates aspects of the bail-in design. This literature includes the papers of Rutledge et al. (2012), Sommer (2014), Persaud (2014), Avgouleas and Goodhart (2015) and Tröger (2018). These papers voice concerns over the timing of the bail-in trigger, whether a recapitalisation can restore confidence in the bank, the possibility of runs on bail-in debt in anticipation of bank failure, and the exorbitant discretion regulators hold in choosing the bail-in design applicable to resolving a failing bank. Our paper adds qualitatively to this literature by expanding the list of concerns, making existing ones more precise (with a theoretical model), and explaining the presence of bail-in design issues given prevailing political economy incentives. Our paper adds quantitatively to this literature by explicitly modelling the key components of the bail-in design giving rise to these concerns, and assessing their joint impact on contagion risk.

Our approach to modelling contagion falls squarely within a widely used approach in the contagion literature. In this literature, the behaviour of institutions is driven by incentives (i.e. increase the return on equity) subject to constraints arising from regulation, contractual obligations, the financial market and internal risk limits (Aymanns et al. (2018)). In times of distress, constraints more often bind (Brunnermeier and Pedersen (2009), Pedersen (2009), Adrian et al. (2013), He and Krishnamurthy (2019)). An institution's behaviour then becomes dominated by survival actions to meet these constraints to avoid failure rather than by normal-time activity (Aymanns et al. (2018)). In line with Cifuentes et al. (2005), Adrian and Shin (2010), Gai et al. (2011), Kok and Montagna (2013) and Greenwood et al. (2015)), who effectively take this observation as a starting point, institutions in our model seek to avoid default by aiming to fulfil contractual obligations and comply with regulatory and market-based constraints. Our model makes use of the system-wide model developed by Farmer et al. (2020), who model the interaction between funding contagion, overlapping portfolio contagion and exposure loss contagion in a multi-layered financial system consisting of banks and non-banks. We go beyond Farmer et al. (2020) by incorporating a model of bail-in design with a novel pricing tool for bail-in debt. Our pricing method is akin to the valuation methods for contingent convertibles proposed by Pennacchi (2010) and Chen et al. (2013), and is in line with the standard approach to asset pricing (Black and Scholes (1973), Merton (1974)) based on jump processes (Merton (1976)).

The literature suggests that uninsured deposits can lead to bank instability and be subject to self-reinforcing runs (Diamond and Dybvig (1983), Goldstein and Pauzner (2005), Egan et al. (2017)). Such feedback mechanisms can even result in multiple equilibria. To cut through the complexity of modelling runs as an equilibrium outcome, we take the simpler approach adopted in stress tests by posing "what-if" scenarios (Schuermann (2014), Engle (2020), Farmer et al. (2021)) under which creditors cease to refinance a bank's bail-in debt. We thereby stay in line with the system-wide stress testing approach taken in the paper. More broadly, our model contributes to the contagion literature (see e.g. Eisenberg and Noe (2001), Cifuentes et al. (2005), Allen and Babus (2009), Gorton and Metrick (2012), Elliott et al. (2014), Acemoglu et al. (2015) and Duarte and Eisenbach (2021)) – which typically assumes all banks are liquidated upon default – by incorporating that the norm for dealing with failing SIBs today is resolution via bail-in, as well as by assessing how the chosen failure procedure affects contagion and the degree to which it is amplified by multiple interacting contagion mechanisms.

Finally, though bail-in has been designed with systemic considerations in mind,[10] this is not enough to attest its suitability on a system-wide scale – as Aymanns et al. (2016) have demonstrated in the context of Basel II. By modeling explicitly its systemic impacts we thus fill an

[10]See: Directive 2014/59/EU of the European Parliament and of the Council.

important gap in the literature.

The remainder of the paper is organised as follows. Section 2 spells out the data we use to calibrate our model. Section 3 sets out the legal framework for bailing-in banks in Europe. Section 4 models the design of the bail-in mechanism as stipulated in this framework. It also delineates our system-wide stress test model capturing the systemic implications of the bail-in design. Section 5 presents our results and sensitivity analysis. Section 6 elucidates the policy implications of our findings in the light of the political economy incentives at play. Section 7 concludes and highlights avenues for future research.

2 Data

The European financial system can be represented as a bi-partite multi-layered financial network. The first set of nodes in this network represent the balance sheets of its banks and non-banks and the second set of nodes are the commonly held securities. To calibrate these nodes to data, we use 2017Q4 data provided by *S&P Global Market Intelligence*, the *ECB Statistical Warehouse*, and the 2018 *European Banking Authority* (EBA) stress test results. We calibrate the consolidated bank balance sheets using *S&P Global Market Intelligence* data. The *S&P Global Market Intelligence* data also allow us to map each bank's liabilities into seniority classes of debt in order to estimate each bank's bail-in debt. We include 40 of the 48 banks that participated in the 2018 EBA stress test.[11] We next use data from the *ECB Statistical Warehouse* to construct a representative balance sheet of leveraged and non-leveraged non-banks in the European financial system.[12]

Interconnections among banks and non-banks in the European financial system can be represented by edges in the network. Each edge represents a financial contract and each layer in the multi-layered network contains different type of edges. A layer contains contractual links of the same type (e.g. interbank loan), seniority, and maturity. Our data specify the total contract value per type of contract on an institution's balance sheet, but do not disaggregate these data into information on the individual contracts, and do not specify counterparty information. Therefore, we have to reconstruct interconnections in the network using the calibrated balance sheets of financial institutions as a data input, in line with Greenwood et al. (2015). Since it is common that these data are missing in systemic risk analysis, by now various methods have developed to reconstruct the network's edges using balance sheet data as an input (see e.g. Anand et al. (2018)). We employ an established probabilistic method for reconstructing the network by Kok and Montagna (2013). In line with the contagion literature (see e.g. Gai and Kapadia (2010) and Caccioli et al. (2014)), we average our results across realisations of the reconstructed networks to remove dependency on the probabilistic realisation. In Online Appendix B, we describe our data to construct the bi-partite multi-layered network of banks, non-banks and their interconnections in more detail. This Online Appendix also depicts the stylised balance sheets of banks and non-banks in our model, as well as the list of EBA banks we model, ordered by asset size. Next, we set out the framework

[11] We excluded 8 banks because of gaps in *S&P Global Market Intelligence* data preventing us from adequately specifying the balance sheets of these banks.

[12] We split proportionally χ percent of the balance sheet of the representative non-bank in a 'leveraged non-bank' and $(1-\chi)$ percent in a 'non-leveraged non-bank' (and remove the representative non-bank from the system), since we do not have a reliable estimate of the relative size of leveraged non-banks. Neither do we have data on the initial leverage of the leveraged non-bank $\lambda^{t0} := \frac{\hat{E}_i^{t0}}{A_i^{t0}}$. As part of our robustness checks, we varied parameters χ and λ^{t0} to investigate how they affect banking sector stability. We found that contagious asset losses in the banking sector increase with χ and λ^{t0}, in line with intuition, and that these parameters do not qualitatively alter our results.

for resolving failing banks that applies in Europe, the Bank Resolution and Recovery Directive (BRRD).

3 The Bail-In Design under the Bank Recovery and Resolution Directive

The motivation for detailing the bail-in design under the BRRD here is that we aim to model the bail-in design (in the next section, Section 4) exactly as it is specified under the BRRD. The BRRD is transposed into law in European jurisdictions and the United Kingdom. Before diving into details of the bail-in design under the BRRD, we explain in a nutshell how the bail-in mechanism works. This makes it easier to understand how the design features of the BRRD connect into a sequential bail-in process, as summarised in Figure 1.

The bail-in design consists of two parts, the bank-specific and structural bail-in design. The bank-specific bail-in design consists of the parameters that resolution authorities must set for each novel bank bail-in (blue boxes). The bank-specific design consists of four key parameters: the failing-likely-to-fail (FLTF) threshold, the *ad-hoc* debt exclusions, the recapitalisation target and debt-to-equity conversion rates. Once these four parameters have been specified, the resolution authority knows how to carry out the bail-in of a failing bank. The structural bail-in design consists of the parameters that tend to apply structurally through time and across a set of banks (green boxes). The regulator must determine these in advance of any bail-in taking place. It consists of four core parameters: *a-priori* debt exclusions, loss absorption requirements, uncertainty in the bank-specific bail-in design and the speed to complete a bail-in.[13]

Let us walk through the steps of the bail-in process at a high-level. Imagine a bank is in distress and possibly insolvent or illiquid. The first step of the resolution authority is to determine whether this bank is deemed failing-or-likely-to-fail (step 1). If the answer is yes and the resolution authority has also determined that resolving the bank is in the public interest and cannot be timely prevented otherwise (e.g. through a private-sector solution),[14] then the bank will be resolved with one of the four resolution tools. The four resolution tools are: (1) sale of business; (2) bridge institution; (3) asset separation; and (4) bail-in.[15] The resolution authority then has to decide which subset of the four resolution tools to use. It will typically pick the bail-in tool, since this is usually the only tool that meets the resolution objectives (Chennells and Wingfield (2015)). The reason is that splitting up large and complex firms, so that critical functions can be preserved while other parts may be wound down, may not be feasible in a timely manner with the other resolution tools.

If the resolution authority has decided that the bank should be resolved with the bail-in tool,

[13] Since we model banks at the level of their consolidated balance sheet, we do not model design options that apply when bailing-in a parent institution with multiple subsidiaries, possibly spread across multiple jurisdictions. For instance, we do not capture the difference in a bail-in design reliant on Single Point of Entry (SPE) or Multiple Point of Entry (MPE). The SPE approach relies on the upstreaming of losses to the parent and the down streaming of capital to an ailing subsidiary to keep the subsidiary in going-concern at out of resolution, whereas the MPE approach relies on bailing-in individual subsidiaries by local resolution authorities (König (2021)). We leave this exercise for future research.

[14] The identification of public interest is a necessary pre-condition for taking resolution action in respect of the failing bank. A resolution action shall be treated as in the public interest if it is necessary for the achievement of and is proportionate to one or more of the resolution objectives and winding up of the institution under normal insolvency proceedings would not meet those resolution objectives to the same extent. The resolution objectives include ensuring the continuity of critical functions and avoiding a significant adverse effect on the financial system, in particular by preventing contagion (see: Article 32 of the BRRD).

[15] See: Article 37(3)(4) of the BRRD.

then AT1 and T2 capital to must immediately be converted to CET1 equity, and if this is already sufficient to recapitalise the bank to target the bail-in stops there.[16][17] If not, then the next step involves determining what set of debt should be excluded on an *ad-hoc* basis from bail-in (step 2), if any. Together with the exclusions of debt from bail-in that have already been made *a-priori* (green box) and the loss absorption requirements that have been determined earlier (green box), this determines the bail-in debt that a bank has to absorb its losses and be recapitalised (grey box). The loss absorption requirements specify how much debt a bank should hold that can absorb losses easily in a prospective bail-in. For instance, short-term debt is considered a less good source of loss absorbing debt than long-term debt, since it is more flight prone and thus might evaporate when it is needed. Hence, short-term debt is excluded from the loss absorption requirements, even though it largely counts as bail-in debt. Since the measure of loss absorbing debt is more restrictive than the measure of bail-in debt, a higher loss absorbing requirement increases indirectly (dashed lines) the amount of bail-in debt a bank holds. *A-priori* exclusions determine directly (solid line) which debt never counts as bail-in debt and the option to exclude debt on an *ad-hoc* basis enables the resolution authority to remove more debt from bail-in on the spot, for instance if including this debt would increase systemic risk.

Once the bail-in debt has been determined, the following step is to recapitalise a bank (step 3). The recapitalisation target determines the haircuts that needs to be imposed on bail-in debt to recapitalise a bank to this target. Haircuts reduce the notional amount of debt of a bank and thereby increase its equity value, since the equity is the residual of assets minus debt. Once haircuts have been imposed (grey box), the next step is to determine the debt-to-equity conversion rates (step 4). The debt-to-equity conversion rate prescribes the number of shares a creditor obtains for each unit of haircut applied to its debt in a given seniority class (grey box). The conversion rate applicable to haircuts on more senior debt is higher than or equal to that on more junior debt. Hence, senior debtors will be at least as well compensated with equity for losses to their debt claim as junior creditors. These rates also determine implicitly to what extent existing shareholders – who are allowed to retain their shares if the bank is still solvent at the start of a bail-in – are diluted. The higher the conversion rates are the more existing shareholders are diluted. Further, the speed of bail-in determines how quickly the debt-to-equity swaps take place following the application of the haircuts in a bail-in (green box). The equity swap can take place immediately after the haircuts take place or months following this. Finally, bail-in design uncertainty determines whether the parameters of the bank-specific bail-in design are clearly established in advance, so that bail-in debt can be properly priced by investors.

Once the debt-to-equity swap has taken place and the resolution financing fund has compensated creditors who are "worse off" in a bail-in than in a hypothetical liquidation of the bank under normal insolvency procedures, the bail-in process is complete. The bail-in process should have helped to restore the short-term viability of the bank by addressing its solvency issues. This will not necessarily address any long-term viability issues. Instead, the business reorganisation plan that the bank is supposed to propose within one month of the bail-in seeks to deliver this.[18]

We will now further delineate the bail-in design under the BRRD.

[16]See: Article 59 of the BRRD.

[17]Under Basel III, the minimum trigger level (in terms of CET1/RWA) required for a CoCo to qualify as AT1 capital is 5.125% (Avdjiev et al. (2013)). Furthermore, these CoCos must be perpetual to qualify as AT1 instruments. CoCos with a shorter maturity date are only eligible to obtain T2 capital status under Basel III. So the way that CoCos enter into a bail-in is that any CoCos qualifying as AT1 and T2 capital, which have not been triggered yet, will be converted into capital at the onset of a bail-in.

[18]See: Article 52 of the BRRD.

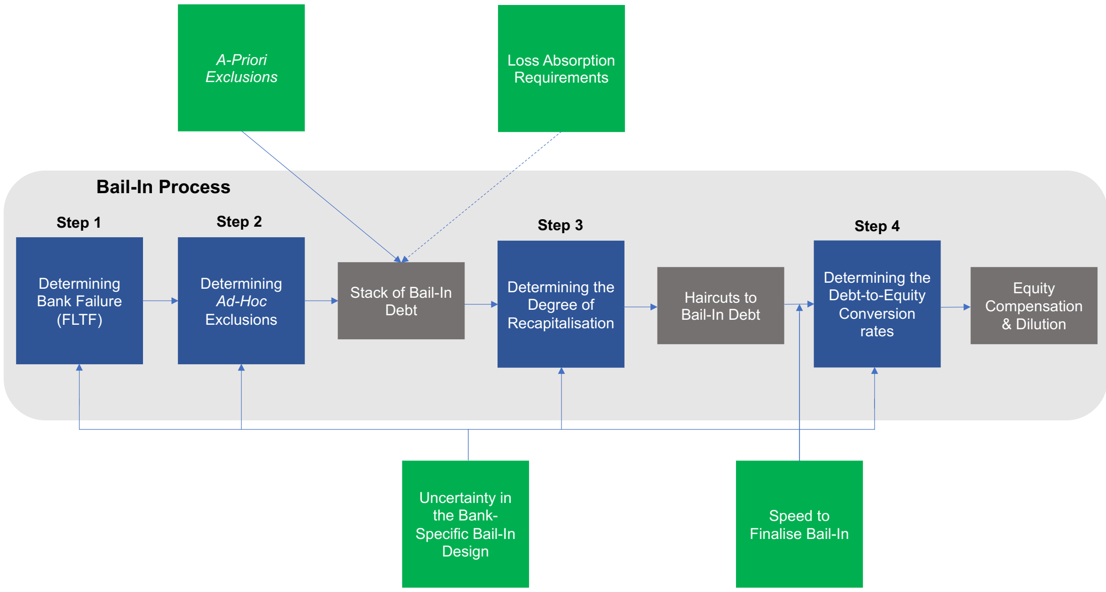

Figure 1: **Schematic overview of the bail-in mechanism.** The blue blocks denote the parameters of the bank-specific bail-in design and the green blocks the parameters of the structural bail-in design.

3.1 Failure Threshold

A bank will be bailed-in if it is deemed to be to *failing or likely to fail* (step 1 in Figure 1), and the other conditions discussed above apply. A bank is considered to be FLTF if at least one of the following circumstances applies:

1. The bank infringes or is likely to infringe upon its requirements for continued authorization (e.g. it slides below its minimum capital requirements);
2. The bank is or is expected to be insolvent in the near future;
3. The institution is or is likely to be illiquid soon.

The determination whether a bank is FLTF is made by an independent valuer in Valuation 1 (EBA (2017b)). Valuation 1 should be consistent with the regular framework of accounting and prudential rules that applies to banks. Nevertheless, the independent valuer is required to apply its independent, skeptical judgement as regards the application of this framework. The appropriate resolution tool and the recapitalisation needs are determined using Valuation 2. Valuation 2 is made on the basis of the bank's economic value to ensure that losses are fully recognised, even if this requires departing from regulator accounting and prudential rules (EBA (2017b)).[19] Valuation 2 is usually more conservative (i.e. lower) than Valuation 1 to avoid under-recapitalising the bank.

3.2 Ad-Hoc Exclusions

Ad-hoc exclusions (step 2 in Figure 1) determine which categories of debt, included in the insolvency hierarchy, should be excluded fully or partially on an *ad-hoc* basis from the application of the bail-

[19] Valuation 2 also informs the choice of debt-to-equity conversion rates by providing an estimate of the post-conversion equity value of new shares transferred or issued.

in tool. The resolution authority may decide to apply *ad-hoc* debt exclusions if at least one of the following conditions is met:[20]

1. The exclusion is strictly necessary and proportionate to avoid contagion;
2. The exclusion is necessary to achieve continuity of critical functions and core business lines;
3. Bailing in the liability is not possible within a short-time frame; or
4. The application of bail-in to these categories of debt would cause a destruction in value such that the losses borne by other creditors would be higher than if those liabilities were excluded from bail-in.

Ad-hoc exclusions may also be made to protect certain creditors from suffering losses, provided some conditions are met.[21] This arguably favors some creditors unfairly. *Ad-hoc* exclusions are permissible in addition to exclusions of debt from bail-in that invariably apply. The set of *a-priori* debt exclusions will be described in Section 3.5.

3.3 Recapitalisation Target

The recapitalisation target (step 3 in Figure 1) is calibrated to aim to restore the bank's short-term viability. This is accomplished by recapitalising the institution so that it complies with the conditions for authorisation and sustains or regains market confidence.[22] The target is met by applying haircuts to the capital and debt instruments of the bank. The bail-in thus focusses on improving solvency rather than liquidity to return the bank to viability in the short-run. Nonetheless, the bail-in may address liquidity issues indirectly, since a better recapitalised bank typically has lower funding costs and less difficulty in retaining access to market funding (Hanson et al. (2011)). A bank's funding costs tends to reflect its credit risk (Flannery and Sorescu (1996), Furfine (2001), Afonso et al. (2011)).

3.4 Debt-to-Equity Conversion Rates

The next step is to apply debt-to-equity conversion rates (step 4 in Figure 1) to debt that has received a haircut. Article 50 of the BRRD tasks the European Banking Authority (EBA) with the duty of providing the conversion rate principles that regulators should use as a guiding principle when setting the debt-to-conversion rate in each priority class (EBA (2017a)). The two principles set out there are:

1. *No Creditor Worse Off (NCWO)*: Resolution authorities should seek to ensure that no creditor or shareholder is expected to incur greater net losses than it would have incurred in

[20] See: Article 44(2) of the BRRD.

[21] Once at least 8% of own funds and liabilities have been bailed-in, the resolution authority is allowed to "bail-out" the bank using the resolution financing fund, in effect excluding sufficiently senior debt from bail-in on an *ad-hoc* basis. (The bail-out amount is capped at 5% of the value of the resolution financing fund, see Article 44(4)(5) and 101(f) of the BRRD.) Under these conditions, the resolution financing fund may be used to prevent equally or higher-ranked creditors from suffering losses or being made worse off, provided some *ad-hoc* exclusions have been made, in effect allowing the application of *ad-hoc* exclusions for any reason outside those listed above.

[22] See: Article 43(2)a of the BRRD.

winding up the bank under normal insolvency proceedings.[23] Claimants are entitled to the difference between the expected loss under resolution and the insolvency proceedings if this is positive, paid from the resolution financing arrangements (RFF).[24]

2. *The Preservation of the Hierarchy of Claims*: Regulators may apply differential conversion rates under certain circumstances, for instance, to ensure that no creditor is worse off or to compensate a haircutted creditor appropriately. Critically, whenever differential conversion rates are applied, the conversion rate to more senior liabilities under the applicable insolvency law must be higher than or equal to that applied to more junior liabilities.

The shares of existing shareholders or other instruments of ownership are cancelled if either the bank is found to be insolvent at the start of bail-in according to Valuation 1 or when all losses are recognised in Valuation 2.[25] Valuation 3 is deployed by the independent valuer to evaluate the NCWO criteria (EBA (2017b)).[26]

We proceed to discuss the parameters of the structural bail-in design, starting with *a-priori* debt-exclusions (green box in Figure 1).

3.5 A-Priori Debt Exclusions

Debt may be excluded from bail-in not only based on *ad-hoc* considerations, but also because of structural concerns. Under the BRRD, relevant *a-priori* exclusions of debt from bail-in are:[27]

1. Covered deposits, including the deposits falling under the deposit guarantee scheme;
2. Secured liabilities up to their collateral value, including covered bonds and liabilities used for hedging purposes;
3. Liabilities to institutions (i.e. credit institutions and investment firms) with a maturity less than seven days; and
4. Liabilities to operator systems with a time to maturity less than seven days.

Although short-term liabilities to institutions and operator systems are excluded, we emphasise that deposits above the deposit guarantee scheme, such as those held by households and retailers, are not – no matter how short term.

3.6 Loss Absorption Requirements

Bail-in is only effective in absorbing losses and recapitalising a bank to a desired capital ratio if the bank has sufficient bail-in debt, or in other words, "loss absorbing capacity". To ensure this, banks

[23] See: Article 74(2) of the BRRD.
[24] See: Article 75 of the BRRD.
[25] See: Article 47(1)a of the BRRD.
[26] Valuation 3 first estimates the actual treatment that shareholders and creditors have received in resolution. Next, it estimates the treatment that shareholders and creditors would have received had the entity under resolution entered insolvency proceedings at the time when the authority decided to resolve the bank (i.e, at 'the resolution date'). The liquidation value is given by the discounted value of the receipts expected to be generated from liquidating the bank over time, taking into consideration applicable rules of the insolvency law (e.g. the allowed time horizon to liquidate the bank) and the market circumstances at the resolution date (EBA (2017b)).
[27] See full list in Article 44(2) of the BRRD.

are required to meet loss absorbing requirements (green box in Figure 1). Each globally systemically important bank is subject to a requirement on its total loss absorbing capacity (TLAC), currently set at 16% relative to its risk-weighted assets and 6% relative to its leverage exposure. Each European bank is subject to a minimum requirement on own funds and eligible liabilities (MREL),[28] where the denominator is the bank's liabilities and own funds. Other than TLAC, which imposes a uniform requirement on all globally systemically important banks, MREL is tailored to the loss absorbing needs of each bank.[29]

Not all debt and bail-in debt instruments are eligible to count towards the loss absorption requirements. While differences exist in the debt eligibility criteria of TLAC and MREL, they have two important rules in common. First, eligible debt must have a time to maturity greater or equal than one year. Second, eligible debt cannot be cross-held by G-SIBs. Further, eligible MREL debt cannot be cross-held by European banks. Cross-holdings are allowed but will be subtracted from the applicable MREL and TLAC measure.

3.7 The Uncertainty and Speed of Bail-In

Bail-in debt can be reliably priced, in the sense of accurately capturing expected discounted losses, only if the bank-specific design parameters (i.e. the four blue boxes in Figure 1) are clearly established in advance by the resolution authority. If the resolution authority has not specified the failure threshold, *ad-hoc* exclusions, recapitalisation target and debt-to-equity rates it intends to apply in a prospective bail-in of a particular bank, then bail-in design uncertainty (green box in Figure 1) will prevail, and bail-in debt and thus risk cannot be properly priced.

Another source of uncertainty is the speed at which a bail-in is completed (green box in Figure 1). The BRRD does not specify how long the resolution may take. As a consequence, European jurisdictions that have transposed the BRRD into law follow distinct approaches. On the one end of the spectrum is a speedy bail-in that is completed in the course of one day, or one 'resolution weekend'. Fast bail-ins are akin to contingent convertible (CoCo) conversions: they fully executed on the day of the trigger. On the other end of the spectrum is a slow bail-in that may take months to conclude.[30]

The resolution authority currently has significant discretion in choosing the appropriate value of the bank-specific bail-in parameters, as the previous sections made clear, and the speed to complete the bail-in. Hence, uncertainty in the bank-specific bail-in design currently prevails. Thus the pricing of bail-in debt in financial markets today cannot capture how the bank-specific bail-in design affects the probability and size of losses to bail-in debt with a given seniority in

[28]See: Article 45 of the BRRD.

[29]Nevertheless, a default MREL exists, which may be deviated from with good reason. The default MREL for a SIB is equal to twice its capital requirements and buffers, because such a bank should be able to both absorb unforeseen losses and be recapitalised to meet its capital standards again. The MREL for a non-SIB is by default set equal to its capital requirements and buffers. The reasoning here is that such a bank will be liquidated under the regular insolvency procedures so has no recapitalisation needs. In effect, small banks thus do not face a MREL requirement, since their MREL coincides with their regular capital requirements. We note the failure of a group of non-systemically important banks can be collectively systemic. Here we think of the "too-many-to-fail" risk elucidated by Acharya and Yorulmazer (2007). Hence, in our opinion even small banks should be subjected to an MREL requirement more stringent than their regular capital requirements.

[30]The Bank of England, for instance, intends to follow the slow approach (BoE (2017)). Under this approach, a bank will be recapitalised over a resolution weekend by applying haircuts to its creditors and by, if need be, cancelling shares of existing equity holders. Only 'some months' (BoE (2017)) later will afflicted creditors be compensated by means of a debt-to-equity swap, if any, and will the NCWO worse off condition be evaluated. While the loss absorption and recapitalisation are thus brought to a quick completion, the entire bail-in is consummated slowly because the shares are returned late.

a prospective bail-in. This is problematic for one of the objectives of the BRRD is to reinstill adequate market discipline by making the pricing of capital reflect actual default risk (Tröger (2018)). This is essential to lower incentives for *ex-ante* excessive risk taking by too-big-to-fail institutions who face reduced risk-sensitivity of capital costs. This is also problematic for increased uncertainty in the bail-in design renders a financial panic when multiple bail-ins are called for in a future financial crisis more likely. According to Tröger (2018) the key *desideratum* is that at least sophisticated investors must be capable to price the risk adequately which requires a reasonably certain *ex-ante* designation of the likelihood and extent of loss-bearing for investors in bail-in debt. Doing so currently requires investors to know the full scope of regulators decision making. While regulatory discretion cannot be altogether avoided, since a resolution of a SIB must to some extent by tailor-made, it should be limited to the indispensable.

4 Model

In this Section, we present our model to evaluate systemic implications of the bail-in design. The first part explains how we model systemic implications of the bail-in design using a multi-layered network representation of the financial system on which we conduct a system-wide stress test. The second part describes our model of the bail-in mechanism and its design. A summary of the notation used to specify the model is given in Table 1 and 2 of Online Appendix A and a summary of our model parameter choices, and conducted sensitivity analysis for each such choice, is given in Table 2 of Appendix A.

4.1 Modelling the Systemic Implications of the Bail-In Design

To evaluate the efficacy of bail-in designs to support financial stability we conduct a system-wide stress test. The time steps t of the stress test are shown in Figure 2.

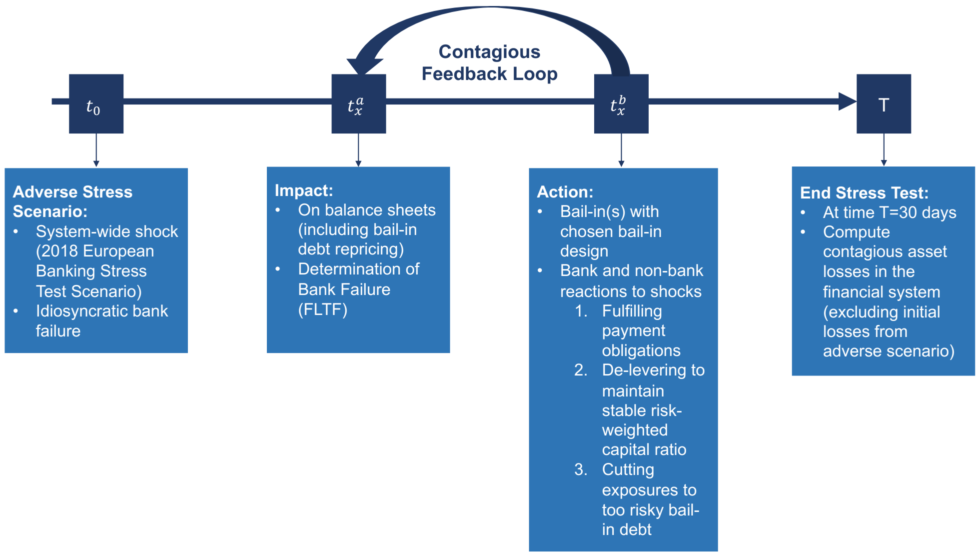

Figure 2: Time line of the system-wide stress test with bank bail-ins. The stress test starts at time t_0 with an exogenous crisis scenario and ends 30 days later at time $T = 30$. In the time steps t_x in between contagion may take place, for $x = 1, .., 29$. At any sub time step t_x^a the impact of shocks is felt and at t_x^b action is taken in response to shocks.

We proceed to walk through the time steps of the stress test using Figure 2 at a high-level. The system-wide stress test starts at time t_0 with an exogenous crisis shock. We consider two types of crisis scenarios, reflecting two distinct cases in which bail-ins can take place and against which we would like to assess the efficacy of different bail-in designs: a system-wide shock and an idiosyncratic shock. The selected exogenous scenario at time t_0 has an impact on balance sheets of (a subset) of banks resulting in asset losses in sub step a of time $t_{x=1}$. If the shocks are severe enough that a bank falls below its failure threshold at time t_1^a, then the bank will be bailed-in at sub step b of time t_1 using the applicable bank-specific and structural bail-in design $\mathcal{D}$. Our model of the bail-in design $\mathcal{D}$ will be spelled out in Section 4.2.

The choice of bail-in design $\mathcal{D}$ determines the exposure losses bailed-in creditors of the bank will suffer and affects the pricing of bail-in debt. The exposure losses generated by the bail-in of the bank at time t_1^b may cause contagious asset losses for its counterparties, triggering a contagious feedback loop (see Figure 2), in which further asset losses on institutions' balance sheets at time t_2^a give rise to further repricing of bail-in debt holdings and bail-ins at time t_2^b. Each time we run through the contagious feedback loop we advance the time from t_x to t_{x+1}.

The selected exogenous scenario at time t_0 may also result in asset losses that prompt an institution to act at time t_1^b. As discussed, institutions in our model are assumed to act when they are hit by adverse shocks according to an adjustment rule in order stay away from binding constraints, in line with Greenwood et al. (2015) and Duarte and Eisenbach (2021). Each of the possible adjustment rules specifying reactions to adverse shocks – action 1 (fulfilling payment obligations), action 2 (de-levering) and action 3 (cutting too risky bail-in exposures) – may contribute to triggering a contagious feedback loop (see Figure 2) involving the interaction among exposure loss contagion, overlapping portfolio contagion and funding contagion. This feedback in turn affects institutions' balance sheets at time t_2^a, including through the repricing of bail-in debt, potentially resulting in more bail-ins and pre-cautionary actions at time t_2^b. All this can bring about another contagious feedback loop. These feedback loops may continue for $x = 1, \ldots, 30$ time steps. The system-wide stress test stops after $T = 30$ days, reflecting our focus on the acute short-term impacts of adverse shocks and contagious feedbacks. Since in the longer-term other forces are likely to dominate market dynamics, we do not focus on this. At the end of the stress test at time T we compute the contagious asset losses associated with bail-in design $\mathcal{D}$. This allows us to evaluate how different bail-in designs affect the financial system. It is important to note that our loss measure captures spillover asset losses; it does not include direct losses due to the exogenously imposed scenario, in line with Greenwood et al. (2015) and Duarte and Eisenbach (2021). This makes our analysis different than but complementary to microprudential stress test analysis that focuses on direct losses for a given adverse scenario (Farmer et al. (2020)).

In Appendix A we motivate how we model each step of the simple stress test depicted in Figure 2 and provide formulas comprising each modelling step. Those less technically inclined can skip reading this.

Interacting Contagion Mechanisms

Prevailing contagion mechanisms may amplify each other resulting in increased contagious losses, as demonstrated by Kok and Montagna (2013), Caccioli et al. (2013), Wiersema et al. (2019) and Farmer et al. (2020). It is therefore key to model prevailing contagion mechanisms jointly, as we do. Figure 3 depicts how interacting contagion mechanisms in our model emerge, thereby summarising our previous explanation of the contagious feedbacks in our model (using Figure 2) through a new perspective. Figure 3 shows that a bail-in (at time t_x^b) can result in exposure losses on bail-in debt (at time t_x^a). This may prompt loss bearing institutions to de-lever (action

2), potentially triggering several rounds of fire sales and funding contagion. Funding contagion results in repayment obligations (action 1); and fire sales generate a price impact resulting in mark-to-market losses for institutions who hold the fire-sold assets (at time t_x^a). We observe that mark-to-market losses on top of exposure losses aggravate downwards revaluations of bail-in debt. These can then trigger the bank's counterparties to stop refinancing its bail-in debt (action 3). In Figure 3, we observe that each contagion mechanisms amplifies the losses of the other mechanisms and renders further bail-ins and contagion more likely. Note that we use the words "can" and "may" often to indicate that whether some event takes place in the stress test depends on the configuration of the experiment. It, for instance, depends on the exogenous scenario imposed, the bail-in design chosen, and the behavioural assumptions on the adjustment rules made.

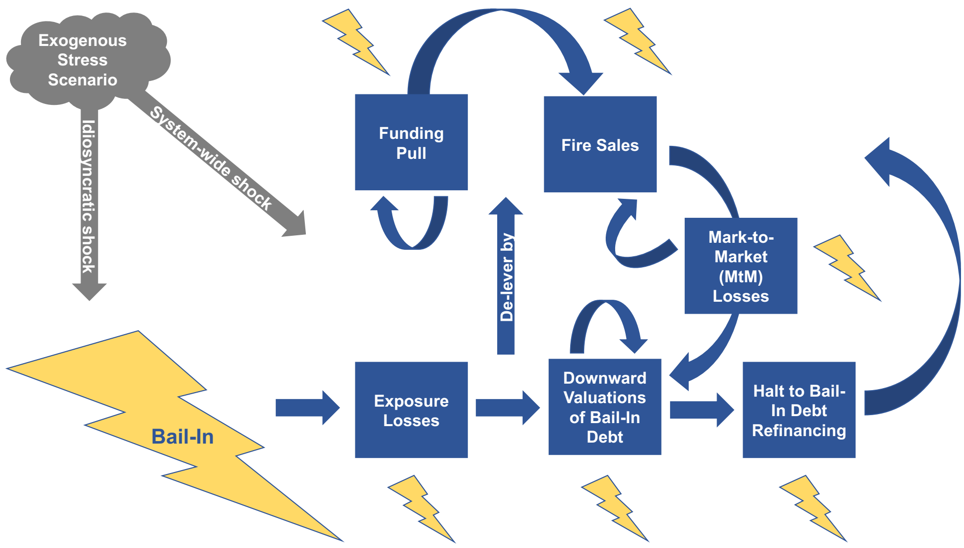

Figure 3: **Interaction among Contagion Mechanisms.** A bail-in (shown with a lightning bolt) may precipitate interactions between the five contagion mechanisms: exposure losses, revaluations of bail-in debt, halts to roll-overs of bail-in debt, funding contagion and firesale contagion. These forms of contagion in turn could cause further insolvency-induced or liquidity-induced bail-ins. An exogenous stress scenario may fuel such interactions among contagion mechanisms directly, rather than via the bail-in mechanism, resulting in bail-ins indirectly.

4.2 Modelling the Design of the Bail-In Mechanism

We now spell out how we model the bail-in design $\mathcal{D}$. Except for a few exceptions, which we note below, we model each of the bank-specific and structural parameters of the bail-in design exactly as these are implemented under the BRRD (as explained in Section 3).

4.2.1 Failure Threshold

A systemically important bank in our model is resolved with the bailed-in tool if it is deemed to be FLTF (step 1 in Figure 1).[31] For simplicity and in line with Hüser et al. (2017), we assume that a bank is subject to only one capital requirement, the minimum CET1 capital ratio. Under

[31] We assume that the public interest test is always satisfied for the 40 SIBs in our model, so do not evaluate this explicitly. Recall that these banks were required to participate in the 2018 European Banking Authority stress test because they are systemically important, globally or domestically. We also assume the absence of timely private-sector solutions.

this assumption, a bank i is bailed-in at time τ_i given by

$$\tau_i = \inf\{t : \rho_i^t < \rho_i^F \text{ or } l_i^t < 0\}. \tag{1}$$

Equation 1 states that a bank i will be bailed-in if its risk-weighted capital ratio ρ_i^t of CET1 equity E_i[32] over risk-weighted assets Ω_i at time t has fallen below the FLTF threshold of ρ_i^F, or if it has become illiquid. The bank is illiquid if the sum of its cash buffer and cash inflows minus outflows, l_i^t, is negative. Because our aim is to assess whether bail-in can resolve the too-big-to-fail problem, we assume the absence of emergency liquidity assistance to delay or avert illiquidity.

The bail-in time τ_i records the state of the bank's balance sheet at the start of bail-in. The evolution of the balance sheet values is endogenously determined in the stress test after the imposition of the exogenous shock. Since Valuation 1-3 are internally consistent in our model and dictated by the stress test losses, we do not model an independent valuer. In our model we thus do not observe asset value shocks between valuations as one might observe in reality.

4.2.2 *Ad-Hoc* Debt Exclusions

Once the bank has been determined to be FLTF at time τ, the next step, taking place at substep τ_e of the bail-in taking place at time τ, is to decide on the *ad-hoc* exclusions $U_i^{k\tau_e}$ in each priority class k (step 2 in Figure 1). The notional amount of bail-in debt in priority class k then reduces to

$$B_i^{k\tau_e} = B_i^{k\tau} - U_i^{k\tau_e}, \tag{2}$$

where $B_i^{k\tau}$ is the bail-in debt in priority class k at time τ (i.e. liabilities minus *a-priori* debt exclusions in priority class k; more on this in Section 4.2.5). Following the *ad-hoc* exclusions, the total bail-in claims of bank i at time $t = \tau_e$ is given by the sum of the bail-in debt contracts B_{ji}^{kmt}

$$B_i^t := \sum_{k\in\mathcal{K}} B_i^{kt} := \sum_{k\in\mathcal{K}} \sum_{j\in\mathcal{F}} \sum_{m\in\mathcal{M}} B_{ji}^{kmt}, \tag{3}$$

where $\mathcal{K} := \{2, ..., 6\}$ is the set of priority classes excluding priority class $k1$, $\mathcal{F}$ is the set of institutions and $\mathcal{M}$ is the set of time to maturities. For modelling purposes, we exclude CET1 equity (k_1) from a bank's bail-in debt count, even though it sits at the lowest rank of its creditor hierarchy, because in our model it is not written down in a bail-in, but instead re-values automatically proportional to revaluations of book equity, as in Hüser et al. (2017). CET1 equity is thus zero if the bank is insolvent.[33]

In this paper, we will assume that the resolution authority does not apply any *ad-hoc* exclusions and leave the study of such impact to future research. The main reason is that it is hard to guess reasonably what debt the resolution authority might exclude on an *ad-hoc* basis in each new bail-in. We do include this step in the bail-in model to enable future analysis of the impact of non-zero *ad-hoc* exclusions without needing to change the model notation.

[32] In line with Farmer et al. (2020), we assume that CET1 equity E_i^t moves one-to-one with book equity $\hat{E}_i^t$ with a data-calibrated offset. See equation 2 and 3 in Online Appendix B.

[33] See modelling details in equations 1-3 in Online Appendix B.

4.2.3 Recapitalisation Target

The next choice concerns to what extent a bank must be recapitalised (step 3 in Figure 1). Given a chosen recapitalisation target ρ_i^{RT} of CET1 equity over risk-weighed assets Ω_i^τ, the total haircuts that must be imposed to recapitalise a bank to target is given by

$$\hat{h}_i^\tau = (\rho_i^{RT} - \rho_i^\tau)\Omega_i^\tau. \tag{4}$$

This shows that the requisite haircuts to recapitalise the bank to target are higher if a bank has a lower capital ratio $\rho_i^\tau = \frac{E_i^\tau}{\Omega_i^\tau}$ at the start of bail-in, or if the recapitalisation target ρ_i^{RT} is set higher . Equation 4 is obtained by rewriting $\rho_i^{RT} = \frac{E_i^\tau + \hat{h}_i^\tau}{\Omega_i^\tau}$, which reflects how much equity $\hat{h}_i^\tau$ should be added to E_i^τ to lift the bank's capital ratio to target. If the bank does not have sufficient bail-in debt to be recapitalised to target, then, assuming the absence of state aid (in line with resolving the too-big-to-fail problem), it will only be recapitalised partially. Hence, the bank's feasible haircuts h_i^τ are given by the minimum of its bail-in debt $B_i^{\tau e}$ and its requisite haircuts $\hat{h}_i^\tau$: $h_i^\tau = \min\{\hat{h}_i^\tau, B_i^{\tau e}\}$. If the bank cannot be recapitalised enough to meet its CET1 minimum capital requirements ρ_i^M, then we assume it will be liquidated via the regular insolvency procedure, as is detailed in Online Appendix C.2.

The total feasible haircuts h_i^τ can be decomposed into the haircuts $h_i^{\tau_a}$ used to lift a bank out of insolvency and the haircuts $h_i^{\tau_b}$ deployed to further recapitalise a bank to its target (see modelling details in Online Appendix C.1.1). The loss absorption phase, a, takes place at substep τ_a and the recapitalisation phase, b, takes place at substep τ_b of the bail-in that happens at time τ. We denote the state of a bank's balance sheet after phase a and b by τ_a and τ_b, respectively.

Two rules exist for how feasible haircuts h_i^τ should be distributed across bail-in investors of the bank. First, haircuts should be sequentially imposed on bail-in debt according to the hierarchy of claims prescribed in insolvency. The haircuts in each priority class $k \in \mathcal{K}$ will then given by $h_i^{k\tau}$ and add up to h_i^τ. Second, haircuts should be imposed equitably across bail-in debt in each priority class.[34] Hence, each contract $B_{ji}^{km\tau}$ in priority class $k \in \mathcal{K}$ will receive a haircut $h_{ji}^{km\tau}$, proportional to its notional value relative to the total notional value in priority class k, summing to a total haircut of $h_i^{k\tau}$ in priority class k. Online Appendix C.1.1 gives the precise formulas describing how these two rules distribute haircuts across priority classes and contracts.

4.2.4 Debt-to-Equity Conversion Rates

Once haircuts have been applied, the next step involves compensating part or all of the bailed-in creditors with an equity stake in the bank (step 4 in Figure 1). The debt-to-equity conversion rate Δ_i^k governs how many shares each creditor in priority class k of the bailed-in bank i receives per unit of haircut applied to its principal amount of bail-in claim B_{ji}^{km} . To enable the regulator to pick conversion rates that adhere to the two conversion rate principles, it is useful to split the conversion rate Δ_i^k in priority class k into that which applies to haircuts made in the loss absorption

[34]Ironically, while bail-in debt must be treated equitably within a priority class, bail-in and non-bail-in debt do not have to be treated equitably under the BRRD, breaking the *pari passu* treatment of debt within a priority class in the insolvency hierarchy. Non-bail-in debt is always loss free, whereas bail-in debt could be fully written down.

phase Δ_{ia}^{k} and that applicable in the recapitalisation phase Δ_{ib}^{k}. Their relationship is given by

$$\Delta_i^k = \frac{\Delta_{ia}^k h_{ji}^{km\tau_a} + \Delta_{ib}^k h_{ji}^{km\tau_b}}{h_{ji}^{km\tau}}. \tag{5}$$

The bank equity share $\epsilon_{ji}^{km\tau_b} \in [0,1]$ that an afflicted creditor receives at time τ_b is then given by

$$\epsilon_{ji}^{km\tau_b} = \frac{\Delta_{ia}^k h_{ji}^{km\tau_a} + \Delta_{ib}^k h_{ji}^{km\tau_b}}{\eta_i^{\tau_b}} \in [0,1], \tag{6}$$

where $\eta_i^{\tau_b}$ denotes the number of outstanding shares of bank i's equity at time τ_b. $\eta_i^{\tau_b}$ is given by the number of existing shares η_i^{τ} plus any newly created shares as part of the bail-in process.[35] We assume that creditors are compensated with equity at time τ_b immediately after haircuts have been imposed.[36]

The time-t value E_{ji}^{kmt} of creditor j's new share $\epsilon_{ji}^{km\tau_b}$ of bank i's equity E_i^t is then given by

$$E_{ji}^{kmt} = \epsilon_{ji}^{km\tau_b} E_i^t. \tag{7}$$

Since the bank's equity value changes over time, its share value also changes over time. This matters because once debt is converted into equity it becomes subject to revaluation, even if no net losses occur because of the conversion itself. The value of j's equity share in bank i will fall if bank i suffers further asset losses due to ongoing contagion after its bail-in has taken place. We will assume that the resolution authority values the bank's equity E_i^t according to its book value.[37]

Conversion Rate Principles

The set of conversion rates must adhere to both the NCWO principle and the preservation-of-hierarchy-of claims principle. Conversion rates must thus be chosen such that the present value of a creditor's claim in bail-in (B) is at least as large as that following a hypothetical liquidation (L) of the bank starting at time τ. Hence, the following inequality must hold: $B_{ji}^{km\tau_e} - h_{ji}^{km\tau_a} - h_{ji}^{km\tau_b} + \epsilon_{ji}^{km\tau_b} E_i^{\tau_b} \geq (1-\zeta_i^{k\tau,L}) B_{ji}^{km\tau_e}$, where $\zeta_i^{k\tau,L}$ is the loss given default in priority class k in a hypothetical liquidation of the bank at time τ (and is defined in Online Appendix C.1.2). This inequality shows that a creditor's claim in a bail-in following the application of haircuts ($h_{ji}^{km\tau_a}$ & $h_{ji}^{km\tau_b}$) and compensation of equity $E_{ji}^{km\tau_b} = \epsilon_{ji}^{km\tau_b} E_i^{\tau_b}$ should be greater than or equal to the creditor's claim following the imposition of the loss given default in a liquidation. Whether the NCWO principle is breached thus depends heavily on the somewhat subjective estimation of the independent valuer regarding the asset recovery value in a hypothetical liquidation of the bank, as well as the valuation method of bank equity. To adhere to the second principle, the conversion

[35] To be precise, the number of outstanding shares $\eta_i^{\tau_b}$ is given by $\eta_i^{\tau} + \sum_{k=2}^{k_a} \Delta_{ia}^k B_i^{k\tau} + \Delta_{ia}^{k_a+1}(h_i^{\tau} - \sum_{k=2}^{k_a} B_i^{k\tau}) + \sum_{k=2}^{k_r} \Delta_{ib}^k B_i^{k\tau_a} + \Delta_{ib}^{k_r+1}(b_i - \sum_{k=2}^{k_r} B_i^{k\tau_a})$, where η_i^{τ} is the number of outstanding shares at the start of bail-in. The definitions of k_a and k_r are given in equation 21 and 30 in Online Appendix C.1.1.

[36] By contrast, in a slow bail-in the debt-to-equity conversion and NCWO evaluation would take place at a later date $\tau_b + \delta$, where δ are the number of days following the initiation of the bail-in and the application of haircuts. So, we assume that Valuation 2, in which haircuts are applied, and Valuation 3, in which the debt-to-equity conversion takes places and the NCWO condition is evaluated, happen in quick succession (both within time step τ), in line with a contingent convertible conversion where the debt-to-equity conversion also happens immediately following the trigger event (Pennacchi and Tchistyi (2019)). We discuss the presumed stability implications of the bail-in speed in Section 6, but leave its quantitative study for future research.

[37] As explained in Section 3.4, the resolution authority is allowed to value the equity claim based on its book value, market value, franchise value or something else. This choice will have implications for whether a creditor is deemed to be worse off. The resolution authority can use this to its advantage. It may, for instance, decide to value the equity according to the franchise value of the firm if this is higher than its book (or market) value to avoid breaching the NCWO condition.

rate in a priority class k must be greater than or equal to that in priority class $k-1$, i.e. we must have that $\Delta_i^k \geq \Delta_i^{k-1}$, $\forall k \in \mathcal{K}$.

Fair and unfair conversion rates

Because the two conversion rate principles must hold, the total set of possible conversion rates is restricted to a set of eligible conversion rates that comply with these principles. Equation 47 in Online Appendix C.1.2 shows that the two types of conversion rates that typically comply are (what we refer to as) "fair" and "unfair" conversion rates. Fair conversion rates are set such that creditors in the loss absorption phase of bail-in suffer pure write downs, amounting to a debt-to-equity conversion rate of zero (i.e. $\Delta_{ia}^k = 0$, $\forall k \in \mathcal{K}$). Creditors in the recapitalisation phase obtain a debt-to-equity conversion rate such that the book value $E_{ji}^{km\tau_b}$ of their new equity share equals the value of their haircuts $h_{ji}^{km\tau_b}$. With fair conversion rates, haircuts to lift a bank out of insolvency are thus not compensated, while those imposed to further raise a bank's capital ratio to its recapitalisation target are compensated fully resulting in no net loss in book value terms at the time of conversion. We call these conversion rates fair because they most equitably administer losses and offer compensation, while adhering to the conversion rate principles. While (what we refer to as) "proportional" conversion rates, which would hand out equity proportional to haircuts imposed, seem superficially fairer than "fair" conversion rates, they typically do not adhere to the conversion rate principles. The reason is that by compensating junior creditors more senior creditors usually can no longer be compensated enough (since the equity value that can be allocated is fixed), resulting in their breaching of the NCWO principle. "Zero" conversion rates or "pure write-downs", in which none of the bailed-in creditors receives compensation for its haircuts, do not satisfy the conversion rate principles either.

While there are many ways to make conversion rates unfair, we define "unfair" conversion rates as the most extreme set of unfair rates that still complies with the conversion rate principles (assuming no *ad-hoc* exclusions).[38] Under our unfair scheme, creditors in junior priority classes will be treated as badly as in a hypothetical liquidation of the bank. With higher estimated liquidation costs of hypothetically liquidating a bank, the number of creditors who receive zero conversion rates thus increases under unfair rates. Since more write-downs are applied in a bail-in with unfair rates than the haircuts necessary to lift the bank out of insolvency, an excess of equity is left non-allocated. The excess equity is then distributed proportionally relative to haircuts imposed on creditors who have remained in the recapitalisation phase of the bail-in. Relative to fair conversion rates, this particular example of unfair conversion rates thus imposes excessive losses on the most junior creditors and distributes excessive gains to the most senior creditors. Senior creditors could even make a net profit. This could be argued to be unfair. The formulaic definition of fair and unfair conversion rates is provided in Online Appendix C.1.2. Figure 1 in this Appendix illustrates how fair conversion rates work and Figure 2 shows how unfair rates compare with fair ones.

It is possible that neither fair nor unfair conversion rates satisfy the conversion rate principles. This could, for instance, occur if the *ad-hoc* exclusions are substantial and the asset recovery rates in liquidation are estimated to be high, in addition to a bail-in taking place late such that the bank is insolvent at the start of bail-in. It occurs because the no-creditor-worse-off principle will then be breached, since the loss given default in bail-in will be higher than that in the hypothetical liquidation of the bank. See Online Appendix C.1.2 for details.

Cancellation or Dilution of Existing Shareholders

[38]Conversion rates are frequently made to deviate from fair conversation rates in some unfair way, whenever the resolution authority applies *ad-hoc* exclusions.

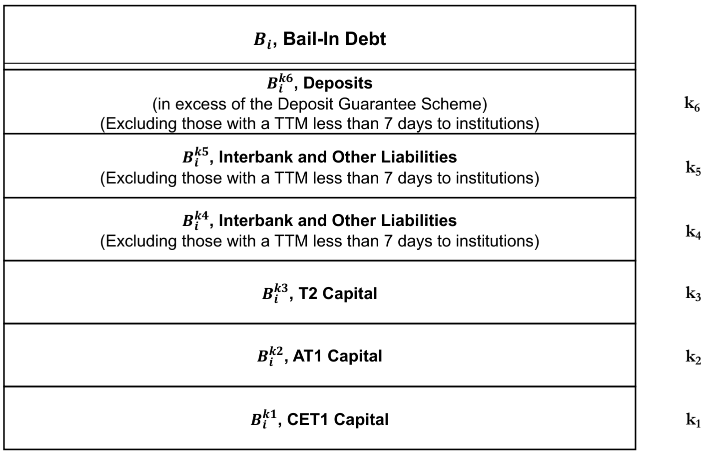

Figure 4: Break-down of bail-in debt B_i^k per priority class $k = k_1, ..., k_6$. Bail-in debt in k_1 is most junior and therefore first in line to absorb losses. Bail-in debt in priority class k_1, k_2 and k_3 is given by CET1, AT1 and T2 capital, respectively. Bail-in debt in priority class k_4 and k_5 consists of interbank and other liabilities excluding contracts with a time to maturity (TTM) less than 7 days with an institution. Bail-in debt in k_6 consists of deposits in excess of the deposit guarantee scheme, excluding contracts with a TTM less than 7 days with financial institutions.

Existing equity holders are not subject to a debt-to-equity swap and are treated differently than debt holders in a bail-in. Shares of existing shareholders are cancelled if the bank is found to be insolvent at the start of bail-in τ,[39] while existing shareholders' ownership will be diluted if the bank is found to still be solvent at the start of bail-in.[40] Their dilution increases if more haircuts are applied and debt-to-equity conversion rates are more generous, as shown in Online Appendix C.1.2.

4.2.5 *A-Priori* Debt Exclusions

We now describe how we model the structural bail-in design. As noted earlier, some debt contracts are excluded *a priori* from bail-in. Given these exclusions, the composition of bail-in debt B_i of each bank in our model is shown in Figure 4.[41][42]

Formally described, the bail-in debt in priority class k available at the start of bail-in at time τ is given by the bank's liabilities $L_i^{k\tau}$ in that priority class minus the *a-priori* exclusions in that priority class $P_i^{k\tau}$:

$$B_i^{k\tau} := L_i^{k\tau} - P_i^{k\tau}, \tag{8}$$

Given the set of *a-priori* exclusions, the bank's total bail-in debt is found by summing over its bail-in debt in each priority class, as equation 2 has shown.

[39] See: Article 47(1)a of the BRRD.

[40] See: Article 47(1)b of the BRRD.

[41] In our model, we make the following simplifications with regards to *a-priori* exclusions. First, we exclude collateralised debt, in line with Hüser et al. (2017). Second, we assume that liabilities to institutions are excluded when the time to maturity is less than 7 days, rather than when the maturity is less than 7 days. This means that longer-term debt that soon matures is excluded in our model, whereas it would be included in reality. We make this simplication to render the type of debt exclusions consistent with the debt that is ineligible to count towards the loss absorption requirements: i.e. debt with a *time to maturity* less than a year.

[42] A minimum requirement on a new debt class, senior non-preferred debt, is being phased-in in Europe. This debt class would reside between rank k_4 and k_5. Since the phase-in is incomplete at the time of writing this paper, we do not include this debt class in this paper.

4.2.6 Loss Absorption Requirements

As noted before, each globally systemically important bank is subject to TLAC and each European bank is further subject to MREL. These stipulate the minimum amount of loss absorbing capacity L_i^M each bank should hold. Online Appendix C.1.3 provides the formulas for the loss absorption requirements used in our model.

4.2.7 Uncertainty in the Bail-In Design

The literature suggests that uninsured deposits can lead to bank instability and be subject to self-reinforcing runs (Diamond and Dybvig (1983), Goldstein and Pauzner (2005), Egan et al. (2017)). The reason is that uninsured deposits are often impaired in a bank default and thus prone to runs. To cut through the complexity of modelling runs as an equilibrium outcome, we take the simpler approach adopted in stress tests by posing three "what-if" scenarios (Schuermann (2014), Engle (2020), Farmer et al. (2021)) under which creditors cease to refinance a bank's bail-in debt. We thereby stay in line with the system-wide stress test approach in our paper. Nonetheless, our runs scenarios stay in the spirit of Diamond and Dybvig (1983) by having investors run as soon as they think their bail-in holdings are at high risk of being impaired.

When there is uncertainty about the bail-in design, creditors cannot incorporate an accurate estimation of expected losses in the pricing of bail-in debt. They thus cannot rely on making roll-over decisions of bail-in debt based on expected losses, as is our "what-if" scenario for creditor roll-over behaviour when there is sufficient design certainty. We model the effect of uncertainty using two possible "what-if" heuristics. One poses that investors roll-over bail-in debt based on the bank's distance-to-default, and the other poses that they do it based on the similarity to a bank that has recently failed (Temzelides et al. (1997)), as has been a frequent heuristic in the past when uncertainty over asset impairment prevailed (see e.g. He and Manela (2016) on investors' acquisition of noisy signals due to uncertainty over asset losses). The formulaic definition of these three criteria (which we refer to as expected, uncertainty and similarity runs) is given in Appendix A. We flag experiments with design uncertainty with $\nu = 1$ and those with sufficient design certainty with $\nu = 0$.[43]

In sum, the bail-in design of a bank in our model is given by $\mathcal{D}_i = (\rho_i^F, \mathbf{U_i^{\tau_e}}, \rho_i^{RT}, \mathbf{\Delta_i}, \mathbf{P_i}, L_i^M, \nu)$. See table 1.

[43] In our model, in contrast to bail-in debt, non-bail-in debt is not subject to rollover risk driven by impairment expectations/guesses, as it is practically immune from losses if issued by SIBs for which resolution (via bail-in) is the norm. As noted, all modelled SIBs in our model will be resolved upon failure. Our assumption on the presence or absence of rollover risk is motivated by the empirical literature showing that uninsured deposits (e.g. deposits falling under the deposit guarantee scheme; and, *de facto*, non-bail-in debt) tend to experience outflows if the bank is under distress, whereas insured deposits do not (e.g. Egan et al. (2017)). In practise, if there is uncertainty whether a failing bank would be resolved or not then non-bail-in debt may also be subject to rollover risk.

Table 1: Bail-in design $\mathcal{D}_i$.

	Bail-in design parameter	**Description**
Bank-specific bail-in design	ρ_i^F	Failure threshold
	$\mathbf{U_i^{\tau_e}}$	Set of *ad-hoc* exclusions
	ρ_i^{RT}	Recapitalisation target
	$\mathbf{\Delta_i}$	Set of debt-to-equity conversion rates
Structural bail-in design	$\mathbf{P_i}$	Set of *a-priori* exclusions
	L_i^M	Minimum loss absorption requirement
	ν	Uncertainty in the bank-specific bail-in design

5 Results

In this section we investigate how the bail-in design affects systemic risk. We do this by studying how the financial system responds to shocks. We impose two different types of shocks corresponding to the two cases in which bail-ins occur. In the first case we assume idiosyncratic bank failures. That is, we assume that the balance sheets of the largest n banks by asset size suffer an asset loss resulting in their failure and subsequent bail-in.[44] In the second case we simulate a system-wide crisis. We do this by imposing a scaled version of the system-wide shock used in the 2018 European Banking Authority (EBA) stress test. We vary the magnitude of the scaling factor for the shock over the interval $s \in (0, 2)$, where a scale factor of one corresponds to a shock equal in magnitude to the original EBA shock.[45] We measure systemic risk by tallying contagious asset losses in European banking sector participating in the EBA stress test, excluding direct losses emanating directly from the imposed shocks.

We then study different bail-in designs by varying parameters and observing how they affect the response of the financial system to the shocks. The parameter space is too large to investigate exhaustively. We simplify the analysis by identifying two extreme values for each parameter, one corresponding to a "good" bail-in design, and the other corresponding to a "poor" bail-in design. These are roughly bounded by the parameter values that the resolution authority indicates they could plausibly pick (see Section 3). We chose these values through a combination of intuition and experimentation. In our results, we always choose the "poor" structural parameters by default, unless otherwise indicated. The reason is that these parameters form the *status quo*. We set the bank-specific parameters equal to their "good" values by default, since regulators have discretion in setting these. Online Appendix D.1 and D.2 also show the results using the poor design choice for both the structural and bank-specific parameters. This exercise confirms that our results qualitatively hold no matter what baseline is used. Figure 5 shows the baseline values for the good and poor bail-in design. The other baseline settings used to generate the results are listed in Table 2 in Appendix A. The key parameters for the bail-in design are given in Table 1. As noted in Section 2, all results shown are averages over 50 different realizations, corresponding to variations in the randomly reconstructed interbank and common asset holdings networks based on calibrated balance sheets.

[44] In practise, we think it is unlikely that more than 3 idiosyncratic failures happen at the same time. We show the results for up to $n = 5$ idiosyncratic failures, however, to explore how contagious losses increase in the number of large bank failures.

[45] Appendix A describes our idiosyncratic and systemic shock scenarios in detail.

Because of insufficiently granular data on bail-in debt cross holdings and because our model does not capture the bail-in design's real-economy effects, we cannot perform an optimal calibration of the bail-in design, nor can we claim full quantitative accuracy of our results. They are, however, qualitatively robust to extensive sensitivity analysis, as will be discussed in Section 5.4. We believe they therefore offer good insight about bail-in designs that make the financial system resilient.

We decompose the study of the systemic implications of the bail-in design into three parts. First, we study the impact of the bank-specific bail-in design. Second, we study the impact of the structural bail-in design. Third, we assess the role that the various contagion channels play in amplifying the systemic implications of bail-ins.

	Bail-In Design Parameters	Good Bail-In Design	Poor Bail-In Design
Bank-Specific Bail-In Design	Failing-Likely-To-Fail Threshold	**Early Bail-In** ρ^F equal to MCR	Late Bail-In $\rho^F = 0\%$ (equal to insolvency point)
	Recapitalisation Target	**Strong Recapitalisation** ρ^{RT} equal to MCR plus twice CBR	Weak Recapitalisation ρ^{RT} equal to MCR plus half CBR
	Debt-to-Equity Conversion Rates	Fair Conversion Rates See equation 31 & 42 in Online Appendix C.1.2	Unfair Conversion Rates See equation 38 in Appendix C.1.2
Structural Bail-In Design	*A-Priori* Debt Exclusions from Bail-In	Short-Term Debt Exclusion Debt excluded with TTM<1Y	Short-Term Debt Inclusion Debt included with 7d<TTM<1Y
	Loss Absorption Requirements	Higher than *Status Quo* to Counteract Debt Exclusions i.e. debt exclusions with 7d<TTM<1Y	Current MREL and TLAC minimum loss absorption requirements See Online Appendix C.1.3
	Uncertainty in the Bank-Specific Bail-In Design	No/Limited Discretion Roll-over based on expected bail-in loss; See equation 25 in Appendix A	Yes Roll-over based on rule of thumb, since bail-in debt cannot be properly priced; See equation 26 in Appendix A

Figure 5: Shows the "good" and "poor" baseline configurations of the bail-in design. By default we show results for the good bank-specific design (highlighted in green) and the poor structural design (highlighted in red). MCR stands for minimum capital requirements ρ_i^M, CBR stands for combined capital buffer ρ_i^{CB} and TTM refers to the time to maturity.

5.1 Stability Impact of the Bank-Specific Bail-In Design

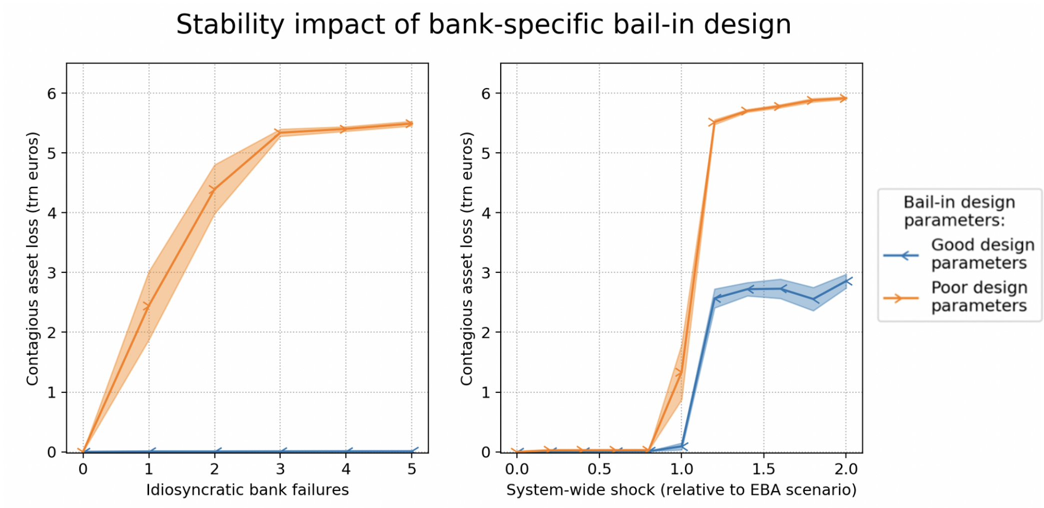

Figure 6: Stability impact of the bank-specific bail-in design, for an idiosyncratic and EBA shock scenario. The blue line uses good and the orange line uses the poor bank-specific design parameters. Both lines use the poor structural design parameters. See Table 5 for the good and poor parameters.

Figure 6 shows how the financial system responds to shocks differently for a bank-specific bail-in design with the good and poor bank-specific parameters, keeping the structural parameters equal to their default (poor) setting. The left panel shows the systemic response to idiosyncratic shocks and the right panel shows the systemic response to the 2018 EBA stress test. The magnitude of the contagious asset loss to the banking system is measured in trillions of euro. We observe that when both the poor bank-specific and poor structural parameters are applied the system's loss can reach up to 6 trillion euros, whereas when the bank-specific parameters are turned to their good values system losses stay below 3 trillion dollars in extreme system-wide crisis and stay close to zero when multiple banks idiosyncratically fail. To put these results in perspective, the size of European banking system that we model is 18.8 trillion euro, so 6 trillion euros represents 31.9% of the system value. By comparison, the asset losses emanating directly from the endogenously imposed EBA shock ($s = 1$) are only about 0.3 trillion euro, which is 1.6% of system value. The contagious amplification of initial losses is thus substantial.

Figure 6 reveals how the choice of the bank-specific design parameters greatly alters systemic risk by a, perhaps surprisingly, large magnitude of multiple trillions of euros (compare orange and blue lines). On the one hand, we observe that "good" bank-specific parameters consisting of an early bail-in, strong recapitalization and fair conversation rates almost completely eliminate contagion in the case of idiosyncratic failures of large SIBs (blue line, left plot). Yet, the good parameters curb but do not crush contagion in the case of increasingly severe system-wide distress (blue line, right plot).[46] On the other hand, we observe that "poor" bank-specific parameters consisting of a late bail-in, weak recapitalization and unfair conversion rates precipitate contagion in the case where large SIBs idiosyncratically fail (orange line, left plot) and exacerbate financial turmoil in case of a systemic crisis (orange line, right plot). In contrast, the bank-specific bail-in design is inconsequential if smaller SIBs idiosyncratically fail, perhaps owing to their smaller systemic impact, as Figure 4 in Online Appendix D.1 shows. This is in line with our experience of

[46]We will find in Section 5.2 that the structural bail-in design also must be well-calibrated to further reduce contagion in systemic crises to more acceptable levels. A well-calibrated bank-specific bail-in design is typically not enough by itself to ensure financial resilience.

bailing-in small SIBs.

In the case of a severe system-wide crisis 1.2 as severe as the 2018 EBA shock (see point $s = 1.2$ on the x-axis), we estimate that the difference in contagious asset losses between a resolution regime with a well-designed and ill-designed bank-specific bail-in could run up to 2.9 trillion euros, representing 15.4% of asset value of the European banking sector participating in the 2018 EBA stress test.[47][48] In the case of a idiosyncratic failure of the largest European bank by asset size, this difference is estimated to be 2.4 trillion euros (11.6% of these banking sector assets), whereas in the case of a idiosyncratic failure of the smallest European bank by asset size participating in the 2018 EBA stress test, the difference in contagious asset losses is negligible.

All in all, the observations in Figure 6 suggest that *financial stability hinges on the bank-specific bail-in design.* This finding should provide regulators with insight into the "good" region of the bank-specific parameter space in which stable financial systems are to be expected, as well as the "poor" region which is best avoided. This Figure does not yet inform us as to the underlying mechanism behind this result. In order to gain intuition as to why we observe a wedge in stability between the good and poor bank-specific design, we have to disentangle how each bank-specific policy parameter shifts the system towards stability or peril, starting with the failure threshold.

5.1.1 Stability Impact of the Failure Threshold

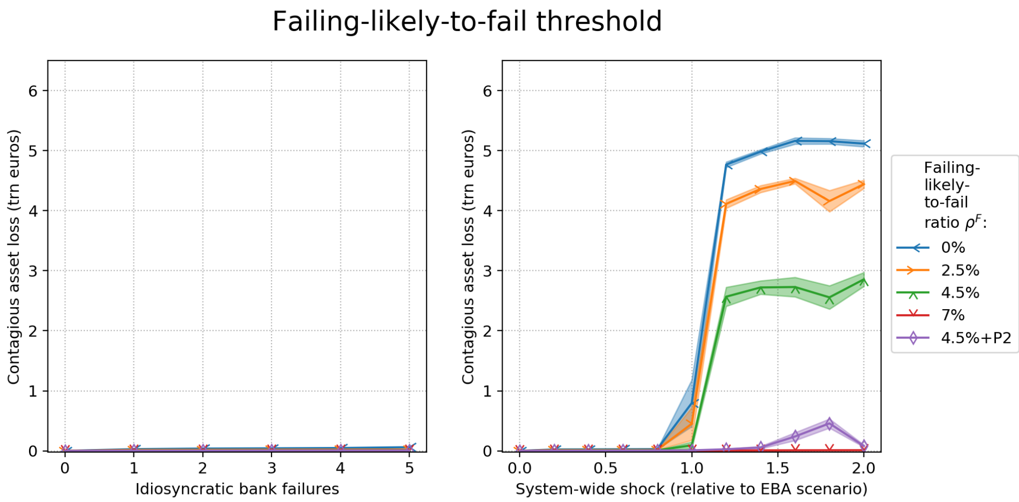

Figure 7: Financial stability impact of varying the failing-likely-to-fail threshold, in response to an idiosyncratic and system-wide shock. All other bank-specific parameters are set to their good values. As before, structural parameters are set to poor values. See Table 5 where the baseline parameters can be found.

Figure 7 studies the stability impact of the failing-likely-to-fail ratio (ρ_i^F). We observe that contagion is higher if the bank is bailed-in only once it has become insolvent (blue line) than when it is bailed-in earlier, for instance at its Pillar I minimum capital requirements of 4.5% of CET1 equity relative to RWAs (green line). Contagion is even lower if that bank is bailed-in even earlier, for example at its Pillar I and II minimum capital requirements of $\rho_i^F = 4.5\% + P2_i$ (purple line). Interestingly, we observe that contagion kicks in non-linearly in the right plot only once

[47]Interestingly, contagious asset losses shoot up just around the 2018 EBA shock size of $x = 1$ and are absent before that, making it seem plausible that banks calibrated their stress test results so as to survive the stress test by a small margin.

[48]For a system-wide shock size below $x = 0.75$ (i.e. if the financial crisis is not severe enough) no bail-ins take place, rendering the choice of bail-in design irrelevant.

the system-wide shock becomes sufficiently severe. Before that our analysis shows that the shock is not severe enough to trigger bail-ins, precautionary delevering, or exposure reductions to risky bail-in debt (plots not shown). It is well known in the literature that contagious losses can increase non-linearly as a function of the adverse shock because of shock amplifications in the financial system (e.g. Cont and Schaanning (2017)). Contagious losses increase more slowly and seemingly plateau after the non-linear spike corresponding to a phase transition.[49] After the spike the extra number of institutions that fail or act in systemically destabilising ways for a given increase in the exogenous shock size is smaller than during the spike.

We observe in Figure 7 that contagion remains small, in the order of multiple millions, in the case of idiosyncratic failures (left plot) irrespective of the level of the FLTF threshold. The reason is that as long as the other bank-specific parameters are well-designed and all surviving banks are well capitalized, the exposure losses resulting from a bail-in of an idiosyncratically failing SIB do not trigger further bail-ins or precautionary action (analysis plots not shown). Contagion in this case does not extend beyond exposure losses resulting from the idiosyncratic bail-ins. (In the case that the other bank-specific bail-in parameters are ill-designed – i.e. weak recapitalisation and unfair conversion rates – then the level of FLTF threshold starts to also matter greatly in the case of idiosyncratic bank failure as Figure 5 in Online Appendix D.1 shows.) The just-explained reasons for why we observe barely any contagious losses with good bank-specific parameters in case of idiosyncratic failures, and why we observe a non-linear increase in losses followed by a slowly increasing plateau do not only hold for this plot, but also for the previous plot and the plots that follow.

These observations point to our finding that *the later a bank is bailed-in (i.e. the lower the FLTF threshold is set), the higher the system's contagious losses are.* The reasons why an earlier bail-in promotes stability are twofold. The less important but obvious reason is that a bank that is bailed-in earlier requires fewer haircuts to be recapitalised to a given recapitalisation target, thereby limiting the potential scope of exposure losses to the bank's creditors which could spur further contagion. The more important but non-obvious reason is that banks that are bailed-in early enough such that they remain solvent at the onset of bail-in can have their creditors compensated with an equal amount of equity (in book value terms) per incurred haircut without breaching the NCWO principle, thereby avoiding *net* exposure losses altogether at the point of conversion. Hence, an early bail-in can limit contagion.

Let's explain this point in more detail by means of a stylised example. Imagine that a bail-in is started late, such that the bank is insolvent at the onset of bail-in and has liabilities worth 15 euros and assets worth 10 euros. The regulator would like to recapitalise the bank to a target capital ratio of equity over assets equal to $\frac{1}{2}$. It must then apply 10 euros of haircuts so that the liabilities are reduced to 5 euros giving a capital ratio of $\frac{5}{10} = \frac{1}{2}$. A total of 5 euros of newly created equity can be distributed to bailed-in creditors. The remaining 5 euros worth of haircuts cannot be compensated resulting in exposure losses to part of the bailed-in creditors. The question is who should receive the compensation and who should be suffering the losses. The conversion rate principles dictate that no creditor should be worse off in a bail-in than in a hypothetical liquidation of the bank and senior creditors should receive at least as much compensation as junior creditors. With these two rules, it makes sense to apply pure write downs to junior creditors and compensate senior creditors with equity. This does not only preserve the hierarchy of claims, it also helps to ensure that senior creditors are not worse off: in a liquidation junior creditors would have also

[49] We confirmed that asset losses keep slowly increasing, rather than plateau by plotting the results for a more severe exogenous shock on the x-axis than shown in this paper (i.e. going beyond 5 idiosyncratic bank failures and beyond a system-wide shock twice as severe as the 2018 EBA shock).

suffered 100% loss given default since the bank is insolvent, whereas senior creditors would have only suffered a positive loss given default if the liquidation costs of the assets are high. Hence, senior creditors must be compensated with equity to ensure they are not worse off. This is what fair conversion rates do. They impose uncompensated haircuts on those creditors who help to lift a bank out of insolvency, while giving equity compensation to creditors who further recapitalise the bank to target.[50]

Imagine now that the bail-in is started early such that the bank is still solvent, but possibly in breach of its minimum capital requirements, with liabilities worth 8 euros and assets worth 10 euros. All the haircuts worth 3 euros can now be compensated with equity worth 3 euros, resulting in no net losses to creditors at the point of conversion. In sum, an early bail-in can avoid net exposure losses (i.e. haircuts imposed greater than equity received) altogether, whereas a late bail-in where the bank in insolvent at the start of bail-in cannot. Therefore, stability is enhanced by an earlier bail-in.

An unrecognised point is that to reduce the likelihood that the bank is found to be insolvent at the start of bail-in, the FLTF threshold ρ_i of CET1 equity of RWA should be set far above the insolvency point of 0%. One reason is that the bank's capital ratio can suddenly drop due to (contagious) asset losses; this happens in our model as well as in reality. So even if the bank is above the threshold in one time period its capital ratio can fall significantly below it in the next. A second reason is that losses are typically recognised only at the onset of bail-in (not modelled). Hence, even if the bank still solvent by a small margin when it gets bailed-in it may be found insolvent in Valuation 2 when haircuts are determined.

Given these reasons, if the FLTF threshold is set at $\rho_i^F = 0\%$, then the bank will for sure be insolvent at the onset of bail-in, resulting in the largest exposure losses and consequent contagion. If it is instead set at $\rho_i^F = 2\%$, then it will still be quite likely that the bank is insolvent at the start of bail-in due to a downwards jump in asset values in our model resulting from contagion. With $\rho_i^F = 4.5\%$, insolvency at the beginning of bail-in will already become less likely and with an even higher FLTF, set equal to the minimum requirements plus Pillar II buffers, it will become even less likely. This explains why a higher FLTF threshold ρ_i^F reduces contagion.

Our finding chimes with the recommendation of Rutledge et al. (2012) that the trigger for bail-in power should be set at the point when an institution would have breached the regulatory minima but before it became balance-sheet insolvent, to allow for a prompt response to a banks' financial distress. In contrast to our paper, Rutledge et al. (2012) do not provide qualitative reasons or quantitative evidence why this is important.

[50] Equation 47 in Online Appendix C.1.2, referred to earlier, makes this discussion precise. It shows which conversion rates do and which do not adhere to the no-creditor-worse-off principle under the applicable bail-in circumstances.

5.1.2 Stability Impact of the Recapitalisation Target

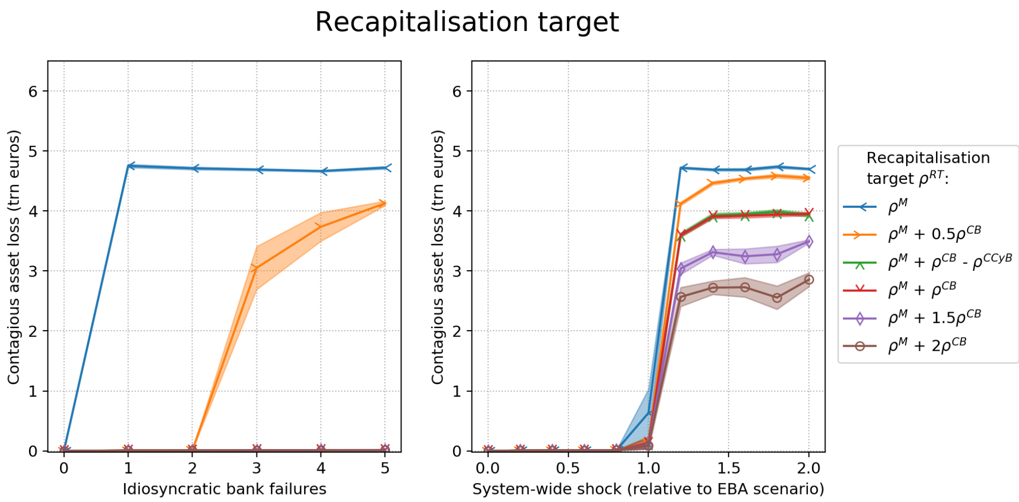

Figure 8: Systemic losses due to the recapitalisation target ρ^{RT}. ρ_i^{CB} gives a bank's combined regulatory buffer and ρ_i^{CCyB} its countercyclical capital buffer. ρ^M denotes the minimum capital requirement. Note that the case $\rho^{RT} = \rho^M$ is worse than the poor baseline case, i.e. $\rho^{RT} = \rho^M + \frac{1}{2}\rho^{CB}$. See Table 5 for the baseline parameters.

Figure 8 studies the stability impact of the recapitalisation target. We observe that contagion remains more limited if the bank is recapitalised strongly, for instance to its minimum capital ratio ρ_i^M plus twice its combined capital buffers ρ_i^{CB} (brown line), than if it is recapitalised weakly, for example to its minimum capital ratio ρ_i^M (blue line). Note that we included the blue line, which is worse than the orange line corresponding to the poor bank-specific baseline, since regulators could plausibly recapitalise a bank to just its minimum capital requirement. Unlike in the previous plot, contagion now also amplifies non-linearly in the case of idiosyncratic bank failures with a poor recapitalisation target choice (orange and blue line). The reason is that a weakly recapitalised bank will be prone to de-lever, continuing funding outflows, and repeated failure in our model, thereby imposing losses on others who are then be dragged into preventative action or a bail-in, resulting in further contagion.

Taking together, the observations in Figure 8 show that *contagion remains more limited, if bailed-in banks are more strongly recapitalised.* Superficially, one might have guessed that a stronger recapitalisation undermines stability, since a higher recapitalisation results in more haircuts increasing the scope for exposure losses that could ignite further contagion. So why does a stronger recapitalisation reduce systemic risk?

The first part of the answer is that a stronger (re)capitalisation means that a bank finds it easier to retain market access to funding in our model and keeps its funding cost low (Hanson et al. (2011)). Indeed, if the target is set too low, the reputational damage that the failure has of itself caused could lead, as it has in the past, to continuing outflows of funds and liquidity problems (Carlson and Rose (2016)), which could threaten the longer-term viability of the bank as well as the stability of the financial system it is embedded in. This part is relatively well-understood in the literature. We model outflows by creditors who stop rolling over maturing funding to bail-in debt that think will experience losses (recall Section 4.2.7).

The second part of the answer is that a stronger recapitalisation means that a bank in our model has fewer incentives to de-lever to attain a more resilient capital ratio, one that is further removed from its minimum capital ratio where failure is likely (recall Section 4.1).

A point that is not appreciated so far by resolution authorities, however, is that recapitalising a bank quite strongly such that it meets both its minimum capital requirements plus its combined capital buffers might still compromise stability on the order of trillions of dollars (compare red line with purple or brown line). The reason is that a bank may not be willing to use its regulatory capital buffers (Goodhart et al. (2008), Goodhart (2013), Farmer et al. (2020), FED (2020), Kleinnijenhuis et al. (2020)), since buffer use comes with penalties in its ability to make discretionary payments such as dividends and bonuses. A bank that is thus recapitalised to the edge of its buffer zone will, with any further (contagious) asset loss, be inclined to de-lever to get out of the penalty zone. Especially in a crisis, it is highly likely that a bank will suffer further losses triggering it to de-lever. In our model, a bank starts de-levering once it has used 50% of its regulatory buffers (see Appendix A). De-levering might be individually rational but tends to be destabilising for the system as a whole. In sum, the novel point is that banks should be recapitalised sufficiently in excess of their minimum capital requirements and regulatory buffers (see e.g. purple and brown line) so that they are unlikely to de-lever or suffer funding outflows following the bail-in.

Another point that is not sufficiently appreciated is that a greater amount of haircuts to meet a higher target does not necessarily imply greater net losses for creditors. With fair conversation rates, creditors in the recapitalisation phase of the bail-in can be compensated fully, resulting in zero net losses in book value terms at the point of conversion. Therefore, a higher recapitalisation target will not exert a strong destabilising pressure via the exposure loss channel as long as conversion rates are fair and bail-in happens early.

Even though a higher recapitalisation can be realised without imposing net losses, debt holders whose claims are converted to equity are unlikely to be happy with their new status lower down in the hierarchy. The equity value fluctuates constantly, exposing the new equity holder to potential losses. It is therefore useful to gauge how much bang for buck a regulator gets in terms of lowering contagious losses by increasing the recapitalisation target and thus haircuts. A back-of-the-envelope calculation reveals that the return on investment (i.e. the drop in contagious losses/increase in haircuts) is significantly greater than one in all cases. As a matter of fact, the haircuts for both the orange and brown line are around 1 trillion euros for a EBA shock of $s = 1.5$, while the drop in contagious losses from the orange to brown line is around 2 trillion at $s = 1.5$. Hence, the return in investment is nearly infinite. The reason is that a stronger the recapitalisation makes it less likely that a bailed-in bank will fail again or act in a destabilising manner.

5.1.3 Stability Impact of the Debt-to-Equity Conversion Rates

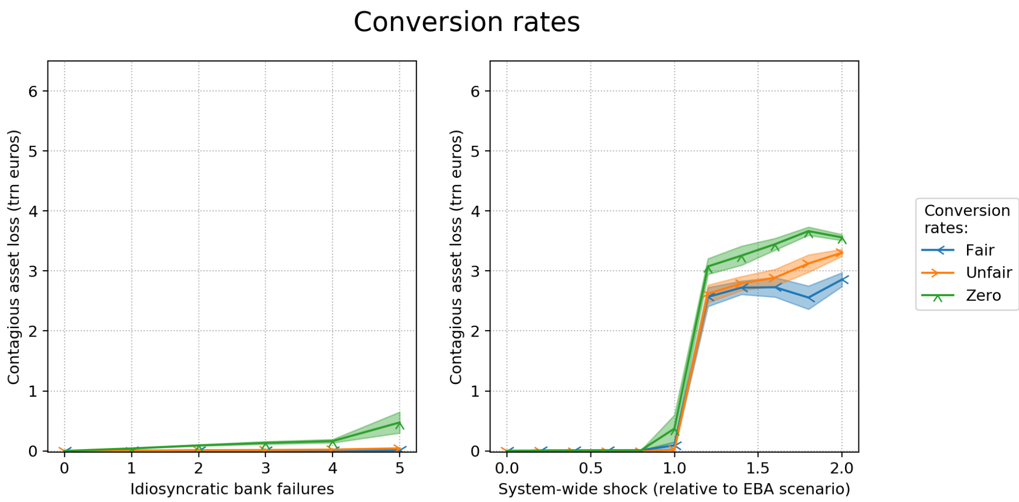

Figure 9: Shows the stability impact of the third bank-specific bail-in parameter: the debt-to-equity conversion rate.

Figure 9 studies the systemic impact of the debt-to-equity conversion rates. We observe that conversion rates matter less for financial stability than the previous two bank-specific parameters do. The loss difference between zero conversion rates (green line) and fair rates (blue line) remains a little less half a trillion euros in the idiosyncratic scenario and 1 trillion euros in the EBA scenario. The difference between fair (blue line) and unfair rates (orange line) is just below half a trillion in the EBA scenario and nearly negligible in the idiosyncratic scenario.

The recapitalisation target choice is more important than the conversion rate choice, because a weak recapitalisation generates substantial contagious amplifications (as a result of de-levering and funding outflows), whereas an unfair conversion rate choice does not. The conversion rate choice just affects how exposure losses are distributed, it does not necessarily affect the size of spill-overs as long as counterparties have sufficient buffers to absorb exposure losses. The failure threshold choice is also more important than the conversion rate choice, since the failure threshold affects the overall exposure loss amount, whereas the conversion rates only affects the distribution thereof.

Nevertheless, the choice between fair and unfair conversion rates still gives a non-negligible difference in losses on the order of half a trillion for severe systemic crises. Fair conversion rates fare slightly better because they distribute losses more evenly than unfair conversion rates, resulting in fewer contagious spill-overs occurring via institutions that would have born the brunt of the losses with unfair rates. Whether fair or unfair conversion rates fare better will in practise depend on the distance-to-default of the bailed-in creditors. The debt-to-equity conversion rates could in principle be tailored such that institutions with the least (most) loss-absorbing ability are compensated best (worst), as long as the conversion rate principles are adhered to. This is, however, not very fair. Moreover, choosing unfair conversion rates based on the distance-to-default in an *ad-hoc* fashion creates uncertainty as to the losses that would be suffered in a prospective bail-in and undermines the pricing of bail-in debt, thereby generating instability as the results in Section 5.3 will show.

Finally, we observe that our zero conversion rate benchmark, which is not compatible with the conversion rate principles, undermines stability most. This is expected, since zero conversion rates imply pure write-downs resulting in a maximum exposure loss for each bailed-in creditor. What is less expected is that the financial stability outcome may correspond to the zero-conversion regime modelled here, even though the resolution authority will eventually apply positive conversion rates (since not doing so will breach the NCWO condition). If the bail-in is slow in the sense that creditors receive their equity compensation late (e.g. BoE (BoE (2017))) creditors will suffer a pure exposure loss in the interim period. Even though this loss is eventually erased, the distress caused by the temporary loss can inflame further contagion. Hence, resolution authorities exacerbate distress in a slow bail-in.

5.2 Stability Impact of the Structural Bail-In Design

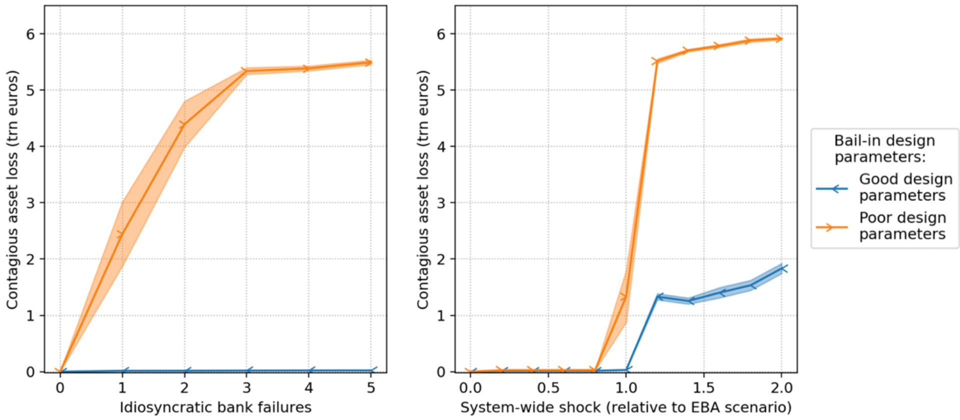

Figure 10: The financial stability impact of the bail-in design on financial stability. This is the key result of the paper and shows the dramatic difference a well-designed bail-in makes in terms of supporting financial stability compared to a poorly-designed bail-in. Here we alter the structural bail-in design relative to its *status quo* (i.e. poor baseline value). The blue (orange) line uses both the good (poor) bank-specific and structural design parameters.

Figure 10 evaluates how the structural bail-in design (consisting of *a-priori* debt exclusions, the loss absorption requirements and uncertainty in the bank-specific bail-in design) matters for financial stability. Compared to Figure 6, in which we only studied the systemic impact of the bank-specific design leaving the poor structural parameters unchanged, we now set the structural bank-specific parameters to their good value in the blue line, and leave it to their poor value in the orange line as before. So Figure 10 shows the overall impact of the choice of bail-in design on financial stability, and is the key result of our paper.

We observe that the orange line in Figures 6 and 10 are identical, while the blue line in this figure is substantially lower in systemic crises – widening the wedge between the good and poor design. This shows that the bail-in design matters even more for financial stability if the structural design is also taken into account in the calibration. Choosing both the bank-specific and structural bail-in design well results in relatively acceptable levels of contagion even in extreme system-wide crises, up to twice as severe as the crisis scenario envisioned by the 2018 EBA stress test (see blue line). Hence, a well-designed bail-in makes bailing in banks a credible option even in severe system-wide crises. This optimistic finding offers tentative evidence to dampen concerns of some leading thinkers that bail-ins could never work in systemic crises. They remain right that ill-designed bail-ins are a recipe for disaster.

In the case of a severe system-wide crisis 1.2 as severe as the 2018 EBA shock (see point $s = 1.2$ on the x-axis), we estimate that the difference in contagious asset losses between a resolution regime with a well-designed and ill-designed bail-in could run up to 4.2 trillion euros, representing 20.3% of asset value of the European banking sector participating in the 2018 EBA stress test (and, for comparison, 22.3% of European GDP as of 2018).[51] In the case of an idiosyncratic failure of the largest European bank by asset size, this difference is estimated to be 2.4 trillion euros (11.6% of these banking sector assets). These quantitative findings suggest that the welfare implications of the choice of bail-in design are significant, especially in times of severe financial crises and large idiosyncratically-failing SIBs.

In a similar way as before, we next vary each of the structural parameters one at a time, keeping the other parameters equal to their default settings (i.e. bank-specific parameters to their good values and structural parameters to their poor values) to study how the individual structural parameters affect stability.

5.2.1 The Stability Impact of the Loss Absorbing Requirements & Debt Exclusions from Bail-In

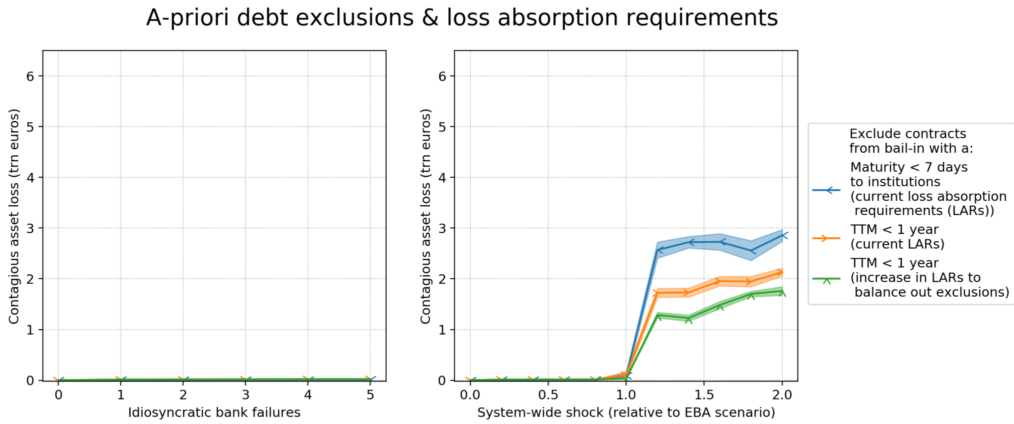

Figure 11: The stability impact of two of the structural bail-in parameters: debt exclusions and loss absorbing requirements, using the baseline set-up of a good bank-specific and poor structural design.

Figure 11 studies the stability impact of two structural parameters: *a-priori* debt exclusions and loss absorption requirements. In the blue line, debt with a time to maturity less than 7 days is excluded from bail-in (as in the poor baseline), whereas in the orange line debt with a time to maturity less than 1 year is excluded (as in the good baseline). The green and orange lines thus focus on the effect of changing the debt exclusions. (In both the green and orange line the loss absorption requirements are set to their poor structural baseline, as per usual.) In the green line, the loss absorption requirements are increased relative to the poor baseline by an amount such that reduction in the loss absorption ratio due to the debt exclusions with a time to maturity between 7 days and 1 year are cancelled out. Increasing the loss absorption requirements namely nudges banks to hold more long-term bail-in debt with a time to maturity greater than one year, since only that debt is eligible to count towards the requirements. So in our model, going from the orange to the green line simply means swapping short-term debt with long-term debt.

We observe that the financial system is less stable when short-term debt with a time to maturity between 7 days and 1 year is included in bail-in (blue line) than when this debt is excluded (orange line) for the case of severe EBA shocks. The stability difference is slight for idiosyncratic shocks though, if the good bank-specific baseline is used, as it is here. In general, however, short-term debt exclusions matter too with idiosyncratic bank failures, as we will show in Figure 12.

[51]This difference shrinks to 2.95 trillion euros if only the bank-specific bail-in design is changed from its poor to good design, but the structural design is kept constant and equal to its *status quo* poor settings.

While we observe that excluding short-term debt increases stability overall, two competing effects are at play. On the one hand, short-term debt exclusions enhance resilience, since fewer bail-in runs can take place, reducing the chance of a "bail-in debt collapse" and resulting systemically costly asset liquidations.[52] On the other hand, these debt exclusions weaken resilience since less bail-in debt renders an unsuccessful bail-in in which the bank must be liquidated (in absence of state aid) more likely. We discuss these in turn.

On the one hand, short-term debt is more susceptible than long-term debt to refinancing risk. If short-term debt is excluded from bail-in, then investors have no incentive to run on the short-term debt of a nearly failing bank, since this debt has become practically immune from losses should the bank fail as it is not subject to haircuts in a bail-in. Runs on included short-term bail-in debt are harmful because the bank is then forced to liquidate assets to meet outflows, thereby triggering contagion.[53] They are also harmful because they reduce a bank's bail-in debt, rendering its successful recapitalisation less likely as only bail-in debt can be used to recapitalise a bank. If a bank cannot be recapitalised sufficiently to meet its minimum capital requirements it must be bailed out or (disorderly) liquidated under the regular insolvency law. Hence, in our model such a bank will be liquidated.[54], because we assume the absence of state aid. This assumption allows us to examine whether a well-designed bail-in can solve the too-big-too-fail problem. Even if a bank can be sufficiently recapitalised that it meets its minimum capital requirements, it may not be possible to recapitalise it fully to target. This will undermine stability, since a weak recapitalisation does, as Section 5.1.2 showed.[55]

On the other hand, the pile of short-term bail-in debt is not only at risk of falling due to runs if it remains included, it also per definition falls if short-term debt is excluded *a priori*. Excluding short-term bail-in debt exerts a negative pressure on stability for one of the same reasons that a bail-in collapse is bad: it renders its successful recapitalisation less likely. Differently, it does not involve risk of asset liquidations to meet outflows. So, moving from a regime with short-term debt inclusions (blue line) to one with short-term debt exclusions (orange line) still improves stability.

To alleviate the downside of excluding short-term debt, we study whether increasing the loss absorption requirements helps. We observe that if short-term exclusions of bail-in debt are *paired* with an equivalent increase in long-term debt, made possible by raising the minimum loss absorption requirements (green line), then systemic risk falls more sharply than in the case where only short-term exclusions are made (orange line). The reason is that augmenting the loss absorption requirements nudges banks to lengthen the maturity of their liabilities, for only debt with a time to maturity greater than one year counts towards the loss absorption measure. Pairing short-term debt exclusions with correspondingly increased loss absorption requirements thus fortifies financial stability, since it swaps short-term with long-term bail-in debt which is more stable and less prone to destabilising runs.

A downside of higher loss absorption requirements is, however, that banks' financing cost might go up, since long-term debt is typically more expensive. This decreases banks' profits and

[52]Recall that in our model (see Section 4.2.7) refinancing risk manifests itself by investor decisions to stop rolling over short-term bail-in debt if a bank's expected losses increase substantially.

[53]Recall that in our model asset liquidations involve fire sales of securities and the retraction of maturing funding to counterparties.

[54]Online Appendix C.2 explains how we model a liquidation.

[55]Another important reason, outside the realm of financial stability considerations, for why short-term debt exclusions have negative implications is that it increases the risk that the NCWO condition is breached. Intuitively, this could be understood as follows. Any loss that needs to be absorbed will rise up through the priority ranks more quickly if more debt is excluded from bail-in compared to the speed at which it would have risen up in a regulator insolvency procedure where all debt is included. Hence, with more exclusions it becomes more likely that a creditor suffers greater losses in bail-in than in liquidation.

therefore retained earnings which is a source of fresh capital (not modelled). Hence, the benefits of higher loss absorption requirements might be offset partially by lower profitability. We leave the inclusion of this channel for future research.

Figure 12 shows the same result as Figure 11, but uses the poor rather than good bank-specific baseline. Interestingly, a comparison between these figures reveals that multiple poor design choices amplify contagious asset losses. With an overall well-designed bail-in (except for retaining design uncertainty) the difference in contagious asset losses – for a systemic crisis shock $s = 1.2$ times the 2018 EBA shock – between a resolution regime where bail-in debt is long-term (good parameter, see green line) and one where an equal amount of bail-in debt can be short-term (poor parameter, blue line) is estimated to be 6.7% of banking sector assets (or 1.3 trillion euros; see Figure 11), whereas with an overall poorly-designed bail-in this difference is 19.8% (or 3.73 trillion euros; see Figure 12). Hence, an additional poor design choice on top of an already poorly designed bail-in amplifies losses by a factor of 3 (for this parameter) compared to an additional poor design choice on top of a better-designed bail-in. Next, we study how the contagion mechanisms at play amplify the systemic impact of bail-in designs.

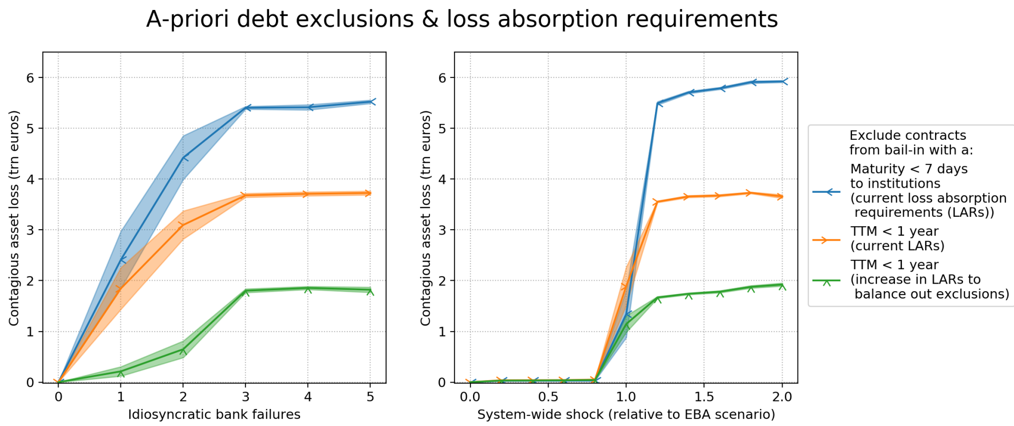

Figure 12: The stability impact of two of the structural bail-in parameters, debt exclusions and loss absorbing requirements, using a poor bank-specific and poor structural design.

5.3 Contagious Amplifications of the Bail-in Design

Figure 13 studies how multiple contagion mechanisms amplify the systemic implications of the bail-in design. We also study the role of non-banks, who hold bail-in debt, by turning them on in the right plot. This enables us to study how contagious asset losses to the banking sector (shown on the y-axis) are amplified by the behaviour of non-banks (i.e. delevering or stopping to roll-over risky bail-in debt) in response to losses on their bail-in debt holdings of banks. If non-banks are turned off we assume they do not react to losses, so no feedback takes place. To study how each contagion mechanism contributes to amplifying losses, we turn different subsets of contagion mechanisms on. In the blue line only exposure loss contagion is turned on. In the orange line, exposure loss contagion, overlapping portfolio contagion and funding contagion are turned on. In the green line, bail-in debt re-valuations are additionally turned on, as well as the "what if" scenario posing that investors stop rolling over bail-in debt whenever expected losses become too high. This scenario presumes investors can rely on bail-in debt prices to inform expected losses. The red and purple line show two different "what-if" scenarios for non-roll overs of bail-in debt whenever bail-in debt cannot be accurately priced, because of uncertainty in the bail-in design. These are that creditors stop rolling over maturing bail-in debt if a bank gets too close to its minimum capital ratio (purple line) or if it looks similar to a previously failed bank (red line). (Revisit Appendix A for a formulaic definition of these three scenarios.)

The results in the paper correspond to the baseline settings where all contagion mechanisms, as shown in the purple line, are turned on (revisit Table 2 in Appendix A for baseline set-up). "Turned on" means that we allow a contagion channel to generate asset losses or liquidity shocks. For instance, when overlapping portfolio contagion is turned off, we set the price impact of sales to zero rendering the mark-to-market losses from fire sales zero, whereas if it is turned on then the price impact is positive. Figure 13 corresponds to poor bank-specific and structural parameters in order to highlight how much poor bank-specific parameters amplify the interaction among contagion mechanisms, rather than to good bank specific parameters as in the rest of the paper. We discuss how a different bail-in design baseline affects contagious amplifications below.

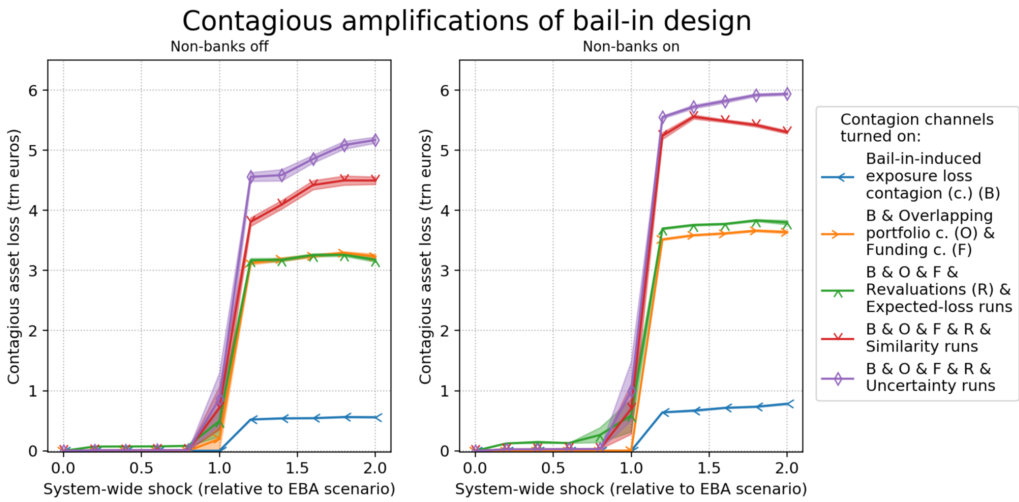

Figure 13: Contagious amplifications of the poor bank-specific and structural bail-in design – as a function of system-wide shocks. The green, purple and red lines correspond to the three different run scenarios we specified in Section 4.2.7 (i.e. expected, uncertainty and similarity runs).

The first thing we observe from Figure 13 is that considering merely the exposure loss contagion that could ensue from bail-ins (blue line), as Hüser et al. (2017) have done, would suggest incorrectly that the European financial system remains resilient in the face of severe system-wide shocks.[56] We observe that if instead, rightly, four more prevailing contagion mechanisms (i.e. overlapping portfolio -, funding -, revaluation - and various forms of run contagion) are taken into account as is shown with the green, red and purple line, then contagion may shoot up for severe enough shocks. The reason that the systemic consequences of the bail-in design tends to be underestimated if multiple contagion channels are not captured jointly is that contagion mechanisms tend to amplify each other. We observe that overlapping portfolio and funding contagion (orange line) amplify exposure loss contagion (blue line) massively, by roughly 3 trillion euros. Revaluation of bail-in debt (green line) amplifies contagious shocks to a lesser extent.

Second, we observe from Figure 13 that a financial system with significant uncertainty about how bank-specific bail-in design parameters will be set in a prospective bail-in (purple line) amplifies contagious losses more than one in which regulators have less discretion (green line) and bail-in debt can thus be properly priced. The intuition is that if the bail-in design parameters are clearly specified in advance, then creditors in our model will only cut exposure to such bail-in debt as is expected to bear large losses (green line). Then, both bail-in debt that is senior enough to escape

[56]The finding of Hüser et al. (2017) is consistent with that of Glasserman and Young (2015), who have shown that systemically destructive exposure loss contagion is unlikely.

haircuts in a bail-in and that which will suffer haircuts yet is amply compensated with equity (resulting in small or no net losses) will be at lower refinancing risk, since their losses are expected to be small. Whereas with uncertainty, creditors are more likely to reduce exposure across priority classes since they are not sure which contracts will be on the hook for losses (red and purple line).

While our finding that *less discretion and more clarity in the bail-in design promotes stability* makes intuitive sense, we have not seen this point made before.

Third, we observe that the degree of amplification among contagion mechanism set off by bail-ins critically depends on the bail-in design. Contagious amplifications fall dramatically if not only the bank-specific parameters (see Figure 23 in Online Appendix D), but also the parameters of the structural bail-in design (see Figure 24 in Online Appendix D) are *well* designed. Moreover, the magnitude of amplification differs starkly between an overall well-designed bail-in (see Figure 24 in Online Appendix D) and an ill-designed one (as in Figure 13 above). From these figures it can be inferred that in a severe financial crisis induced by $s = 1.2$ times the EBA shock, the amplification of losses by the exposure, fire sale and funding contagion mechanisms taken jointly (orange line) relative to the losses generated by the exposure contagion mechanism taken on its own (blue line) is estimated to be 1.7 for a good bail-in design, 3.0 for a good bank-specific but poor structural design, and 5.4 for a poor bail-in design. Hence, an ill-designed bail-in not only increases contagious losses it also increases the degree by which individual contagion mechanisms amplify each other's losses.

Fourth, we observe that the system is more unstable when creditors face uncertainty about the bank-specific bail-in design (purple and red line) than when they have certainty (green line). The reason is that when certainty is provided bail-in debt can be priced properly so creditors will only run on bail-in debt whose expected losses soar (green line). This means that even if a bank is in trouble, only a subset of its creditors would run and senior creditors that are expected to remain loss-free in the bail-in will keep rolling over bail-in debt. Whereas when there is uncertainty, even if the bail-in will leave some bail-in creditors loss free, all creditors will pull back their maturing bail-in exposures to a troubled bank (measured by a sharp fall in its capital ratio, as in the purple line; or measured in terms of similarity to a previously failed bank, as in the red line), since they simply cannot know whether they will be facing losses since bail-in debt prices do not accurately reflect expected losses.

Finally, we observe from Figure 13 that ignoring non-bank holdings of bail-in debt results in a non-negligible underestimation of the systemic impact of bail-ins in the banking system (here on the order of approximately a trillion euros).[57] The reason is that non-banks hold a significant amount of bail-in debt since banks are discouraged from cross-holding it through the eligibility rules of the loss absorption requirements. This raises a question about the efficacy from a financial stability perspective of pushing the risk of bail-in losses outside of the banking sector onto the non-banking sector, as losses on non-banks can feedback via contagion onto the banking sector.

5.4 Sensitivity Analysis

Our sensitivity analysis is split into two parts. The first part deals with the bail-in design. The second part deals with the simulation of the evolving financial network and the contagious shocks that may emerge.

[57] With better data on non-bank holdings of bail-in debt, we could more accurately estimate the degree by which banking stability could be overestimated if non-banks are not taken into account.

Sensitivity Analysis on the Systemic Impact of Bail-In Design

1. We showed that our findings on the systemic effect of the choice of bail-in design parameters hold regardless of whether a bank fails for idiosyncratic reasons or due to a system-wide crisis.

2. We showed our findings hold regardless of the number of idiosyncratic bank failures or severity of a system-wide crisis.

3. We investigated how moving the value of each bail-in design parameter away from the baseline set-ups of the "good" and "poor" bank-specific and structural bail-in design affects systemic risk. This shows robustness in at least four dimensions:

   (a) It provides insight into the directional effect that each bail-in parameter has on financial stability.

   (b) It shows that our results are robust to the choice of default baseline of the bail-in design. Our results are also robust to various other baselines (not shown). These show that the qualitative effect on financial stability of the choice of bail-in parameter is the same regardless of whether the default setting is the good or poor bail-in design, or some other design setting. This indicates strongly that our findings generally hold and do not result from local conditions in the bail-in design parameter space.

   (c) It gives insight into the relative importance of each bail-in parameter in shifting system-wide outcomes.

   Results showing deviations from the good and bank-specific bail-in design are found in Section 5 and Online Appendix D, respectively. Section 5 also showed the systemic impact of changing from a poor bank-specific to a good bank-specific bail-in design, as well as from a poor structural to a good structural bail-in design.

Sensitivity Analysis on the Contagious Amplifications of the Bail-In Design

1. We showed that our findings on the systemic implications of the bail-in design hold for distinct decompositions of prevailing contagion mechanisms that jointly determine systemic outcomes.

2. We studied how the systemic impact of the bail-in design is altered by allowing contagious feedbacks between banks and non-banks and by disallowing this (i.e. turning non-banks off). Non-banks tend to amplify bail-in losses and feed these back onto the banking sector, but the qualitative impact of the choice of bail-in design remains unaltered.[58]

3. We showed that our findings on the systemic impact of the bail-in design are robust to alternative specifications of the market liquidities and institutions' adjustment rules following a shock, which we take as a given in line with Greenwood et al. (2015) and Duarte and Eisenbach (2021). We have asserted that price formation parameters (such as the price

[58] As part of this study, we also examined whether the characteristics of the structure of the non-banking system affect qualitatively the outcome of the efficacy of the bail-in design. We find that increasing non-bank leverage and the relative size of the leveraged non-banking system results in larger feedbacks, but, again, the qualitative results remain the same (not shown). Hence, deviating from data-driven balance sheets compositions, which we take as given, does not modify our fundamental findings.

impact) and behavioural parameters (such as the buffer value at which banks start to de-lever, the speed at which they de-lever, the assets they liquidate in order to de-lever, and the target value that they seek to reach once they are de-levering) do not qualitatively affect outcomes. We also showed that our findings on the systemic implications of the bail-in design hold for different "what-if" scenarios (i.e. adjustment rules) regarding when creditors stop rolling over bail-in debt.

4. We have conducted a host of other sensitivity analyses on our model of the contagion dynamics in the financial system, which are summarised in Table 2 in Appendix A. Among others, we have studied how the jump parameters of the asset price process governing the bail-in debt pricing model, as well as the parameters of the bail-in design affect the pricing of bail-in debt and systemic outcomes (see Online Appendix C.2.2). The outcomes of these test all conform to intuition. They also reinforce the validity our findings. For instance, we find that bail-in debt does not sharply revalue as long as the debt-to-equity swap is at least as valuable as the haircut imposed in a bail-in. This is consistent with our finding that contagious losses do not increase sharply as long as compensated equity stake equals the haircut imposed in a bail-in.

Taken together, our extensive sensitivity analysis points out that our findings on the systemic implications of the bail-in design are qualitatively robust to alternative specifications of baseline settings.

6 Policy Implications: The Bail-In Design and its Political Economy

The previous section showed, perhaps surprisingly, that a possible shift towards stability remains in the hands of policymakers – even in systemic crises. Our evidence presented here also suggests, however, that the current policy parameters might be in the regime of instability. The present bail-in design will allow idiosyncratic shocks to be handled effectively for smaller SIBs, but its application to more systemic crises and larger SIBs remains, as of now, problematic.

Our evidence further suggests that the political economy around bail-in design renders an improvement of the current unstable design towards a more stable one unlikely. We will explain some of the incentives and concerns that the main parties to the running of the bail-in system, notably the regulated banks and the regulatory authorities, but also other stakeholders, will have had. It was difficult enough to get the main principles of bail-in resolution accepted and endorsed by all concerned. For obvious reasons quite a lot of leeway was then left, especially to the regulatory authorities, to interpret and vary the parameters as they, the regulators, thought best. But the devil is in the detail, and as we shall show, that devil has been active.

In this exercise we focus on the bank-specific and structural bail-in design parameters in the same order as before. For each parameter we shall revisit the range that best promotes stability and then discuss briefly how the protagonists currently seek to interpret them.

6.1 Failure Threshold

We showed that an early bail-in promotes stability. Yet regulators in the European Union and United States are prone to bail-in a bank late. Regulators in the United States explicitly plan to bail-in a bank late, only when it is or nearly is insolvent (McAndrews et al. (2014)). European regulators, though legally required to bail-in a bank early when it breaches its minimum capital requirements, may, in practice, bail-in a bank so late that it is close to insolvency. A first reason is that suffered losses are often recognised late in the book equity value of a bank. A second reason is that asset price uncertainty is always an issue, rendering it easy to argue that the bank is in fact compliant with its minimum capital requirements. A third reason is that in contrast to the market equity value of the bank, the book equity of the bank is not (sufficiently) forward-looking masking impending difficulties (Hanson et al. (2011)). Because of these three reasons, the market value of equity oftentimes drops sharply prior to a financial crisis (Baron et al. (2021)), whereas the book equity stays relatively flat (Hanson et al. (2011)). A fourth reason is that a bank's asset value may fall sharply due to sudden asset losses in a crisis. This suggests setting the failure threshold sufficiently high such that, even if unrecognised losses are discovered, more asset losses materialise, and asset value uncertainty prevails, a bank will typically be found to be solvent at the onset of bail-in.

Complementing the regulatory failure trigger with a market-based measure of capital adequacy (see e.g. Sarin and Summers (2016) and Brownlees and Engle (2016)) may combat regulatory forbearance manifest in the tendency to bail-in too late. The potential for regulatory forbearance could be further limited if hitting the trigger automatically leads to a resolution (via bail-in), in similar fashion as hitting the trigger of a contingent convertible (CoCo) automatically leads to a debt write-down or conversion (Calomiris and Herring (2013)). An alternative solution could be a metric based on multiple requirements (Buckmann et al. (2021)). A tricky issue that must be resolved for any failure trigger to be effective is how to avoid equity death spirals in anticipation of the trigger event.

We will discuss the self-interest of each of these parties:

1. Regulatory Authorities

There is no such thing as a foolproof, or perfectly accurate, valuation of a bank, or any other holders', assets. The current value of an asset depends on the future returns that it will bring, discounted back to the present, but the future path of such returns, (and of future interest rates), is unknown, and unknowable, though it is perfectly possible to estimate probabilities, at least up to a point. Indeed, the market value of assets, interacting with their volatility, can be used to make estimates of the distance to, or probability of, default; and, indeed, of the likelihood of any such default having a contagious effect on other similar institutions (Segoviano and Goodhart (2009)). Indeed, such estimates of DD, or PD, have, in the past at least, been a considerably more accurate predictor of actual defaults than accounting valuations.

Nevertheless, they are not, and cannot, be entirely accurate, and, if they were used as policy triggers, they would surely be subject to even greater manipulation than at present (Goodhart's Law, Goodhart (1975)). So, if a regulator was to trigger a failure threshold on the basis of a somewhat mechanical algorithm based on market valuations, and volatilities, it would almost certainly be subject to lawsuit, claiming that the algorithms and valuations were either inherently faulty or inappropriate. The last thing that regulators want is to have each failure leading to a long-drawn-out legal battle. And if they were to lose such a battle, as might occur, their credibility would be shot, whereas the failing bank has no such constraint.

Indeed, for such reasons, the failure of a bank is hardly ever triggered early by an assessment of its solvency, whether based on market or accounting values. Instead, what normally happens is that various signals, including market valuations, lead informed investors to believe that a bank may be subject to possible failure; so, to avoid loss and also having their claim frozen, they withdraw funding. The normal sequence of failure is that a bank thus becomes illiquid, and appeals to the central bank for liquidity support, almost always claiming that it is fully solvent. The scale of support needed in the context then causes the regulatory authorities to send in an independent valuer, (valuation one in the process), who must be independent from any public authority, to inform the FLTF criteria.[59] And it is this valuation that typically leads the regulatory authority to claim, at a late stage, that the failure threshold is triggered. In an act of regulatory forbearance, the central bank may also provide liquidity support even though the bank might already be insolvent, further resulting in any eventual bail-in to be started late, when the bank is already severely insolvent.

Another concern about triggering a failure threshold early is that this could well lead the market in turn to reduce the exposure to whichever bank is the next weakest "antelope in the herd" (Blinder (2013)), especially if there is pronounced uncertainty over the losses creditors could suffer in a prospective bail-in. So, after Bear Stearns, the market focused on Lehman Brothers; after Lehman Brothers, Merrill Lynch; and after Merrill Lynch, everyone else; with a similar story occurring after Northern Rock in the UK. With regulators, like bank managers, having relatively short periods in office, there is a natural inclination to defer taking a step which could have serious and widespread consequences, in the hopes that (market) conditions might improve, and the need for resolution may go away; or, at least, be deferred until someone else takes up position.

2. The bank management and other equity shareholders

Bank managers will naturally not want their period in office likely brought to an early close by an earlier and higher trigger and equity shareholders will not want the reduction in equity prices that an early trigger would induce. While the regulatory authorities could, in theory, ignore the wishes of bank management and their shareholders, there are several reasons why they are likely to listen intently to the concerns of the regulated, in this case the banks and their shareholders. First, regulation becomes much easier if the regulated are reasonably content with the working of the system. Second, if the regulated are not contented, the banks concerned can wield considerable lobbying power and funding in support of efforts to relax the regulatory constraints. Third, if the regulated dislike the form of the regulation, they can deploy legal and other professionals to find ways around such regulations; and the pay that they can offer such professionals is much greater than the regulatory authorities can offer. Indeed, there is cynicism about the revolving door, whereby those trained and having become expert in regulation among the authorities then leave for much higher paying jobs among the regulated, where they can tell their new bosses how best to avoid, and even evade, the regulations.

The political, and financial, muscle that the banks can wield is probably greater in the USA than in most other developed countries, but without the USA no international system of financial regulation will work.

3. The accountants

Basing the failure threshold trigger on an accounting valuation, even though it is normally induced by illiquidity problems, provides accountants with a more prominent role and greater business. At the same time, the involvement of accountants gives the regulatory authorities someone to blame

[59] See: Article 36(1) and 36(4)a of the BRRD and EBA (2017b).

if things go wrong, and any legal suit may then be aimed at the accountants rather than the authorities. Moreover, the process of valuations, and setting debt to equity conversion rates will be normally handled by the accountants rather than by the regulatory authorities themselves. So, having the accountants and accounting values taking a prime role helps the regulatory authorities both to avoid blame and some tricky technical issues.

4. Other creditors
Other creditors are likely to ambivalent. An earlier trigger will protect them further against loss, should the downwards spiral continue and subsequent resolution become necessary. On the other hand, an earlier trigger is almost certain to lead to some immediate downgrade in their valuations, especially if the equity valuation in a prospective debt-to-equity swap is valued below the haircut. Further, some creditors may not like, or even be allowed, to hold equity in lieu of debt. So, they are not likely to form a strong body of political opinion to fight for early threshold triggers.

6.2 Recapitalisation Target

We demonstrated that a *strong recapitalisation* fortifies resilience. Yet none of the three recapitalisation options advanced by the BRRD recapitalises a bank sufficiently; these being recapitalisation to make a bank meet its minimum capital requirements; recapitalisation to bring the bank's capital ratio into line with its peers; and recapitalisation to ensure a bank also meets its regulatory buffers. All three are too weak. And so is the recapitalisation target proposed under BRRD II, which is equaled to a bank's minimum capital requirement plus its combined regulatory buffer minus its applicable CCyB buffer. Instead, a bank should be recapitalised in excess of both its minimum capital requirements and its regulatory buffers. This permits the bank to suffer inevitable further losses in a systemic crisis without being pushed into the regulatory buffer zone, where it will be inclined to act in destabilising ways to delever in order to avoid restrictions on discretionary payments and stay away from default, as shown in the results. Recapitalisation to a level where a bank's capital ratio outstrips comfortably both its minimum requirements and regulatory buffer standards is shown to be a good design choice boosting stability. But this is unlikely to happen. The recapitalisation target requires that it complies with the conditions for continued authorization, and sustains or regains market confidence. This, obviously, is a somewhat subjective matter, and leaves a lot of leeway for the concerns and incentives of those involved. Again we turn to these:

1. The regulatory authorities
The higher the recapitalisation target that is required and achieved, the more likely the long-term resilience of the bank and the system at large, is likely to be obtained. Indeed, if the target is set too low, the reputational damage that the failure has of itself caused, could lead, as it has in the past, to continuing outflows of funds and liquidity problems (Carlson and Rose (2016)). Having agreed to the recapitalisation target in the first place, the central bank would then be virtually bound to provide, almost unlimited, liquidity support to the recapitalised bank. That, in turn, could lead to financial and reputational loss should the bank being resolved not manage to succeed after all.

But, while the self-interest of the authorities is for a high target, the self-interest and incentives of all those involved in the bank itself are for the lowest possible target. It would be feasible again for the commercial bank, and/or its various creditors, to go to law to claim that the target was too high. In court they could argue that their property rights were infringed upon unduly or that the no-creditor-worse-off principle was violated. Again, as with the trigger, the regulators cannot easily impose a system on the regulated, which the regulated dislike and will try intensely to avoid or evade.

2. Bank managers, shareholders and other creditors
The higher the recapitalisation target ratio is, the more the equity shareholders will be diluted, and the less their expected returns. Similarly, the other bail-in creditors will suffer greater haircuts, if a high target is imposed. This will generate a net loss if their haircut is not compensated fully with a debt-to-equity swap. Even if creditors do not encounter net losses as a result of the swap, the swap exposes the new equity holders to any future losses by being thrust to the bottom of the claims hierarchy. Some creditors for this reason will not like holding equity. Others may not be allowed institutionally to hold equity. So for all these reasons, those involved in the commercial banking sector will pressurise the regulators, and the politicians, to impose as low a recapitalisation target as seems consistent with a reasonable chance of successful recovery.

6.3 Debt-to-Equity Conversion Rates

We showed that *fair debt-to-equity conversion rates* tend to buttress stability more than unfair or zero conversion rates. Yet current debt-to-equity conversion rates may not be in the regime of stability. The EBA explicitly refrained from meeting a request from market participants to provide detailed examples indicating how conversion rates might be set (EBA (2017a)). They responded that they "were not able to identify a model for application of conversion rates", implying that their choice does not clearly favour well-designed and clearly-specified in-advance fair rates. Surely, if regulators were planning to apply typically fair conversion rates, they would have pronounced so in the interest of stability. Specifying pricing parameters in advance, such as the conversion rate, allows bail-in debt to be more confidently priced. This reduces uncertainty and promotes stability. Instead, the EBA (2017a) reveals that regulators prefer to retain discretion to elect unfair conversion rates, so as to accommodate *ad-hoc* exclusions and adhere to the NCWO principle. We will discuss the self-interest of the stakeholders now individually:

Refusal to specify in advance how debt-to-equity conversion rates will be applied gives regulators an extra degree of freedom to respond to (political) pressures for additional *ad-hoc* exclusions and the (unknown) context of the bail-in. Regulators will prize such additional flexibility.

While banks will suffer in so far as the extra uncertainty will lead to greater risk premia on bail-in debt, such uncertainty is probably dwarfed by the time-varying uncertainty over PD. Moreover such bank managers will not, in general, be actual holders in person of such bail-in debt, and will have less direct interest in this issue.

In contrast such debt is most likely held largely in non-bank financial intermediaries. While the model of bail-in remains untried in a systemic crisis, such intermediaries may be reluctant to challenge the regulatory authorities, to limit their discretionary flexibility, perhaps relying on their own powers of (legal) protest *ex post facto*, if this should lead to large, and contagious, losses in the course of a systemic crisis.

6.4 Debt Exclusions and Loss Absorption Requirements

We showed that systemic risk tends to fall if short-term debt with a time to maturity less than one year is excluded from the application of the bail-in tool,[60] especially if this exclusion is paired with a proportional increase in the loss absorption requirements. Such exclusion would make bail-in debt consistent maturity-wise with loss absorbing debt in MREL and TLAC.[61] Yet,

- Deposits not falling under the Deposit Guarantee Scheme;
- Liabilities to institutions with a maturity between seven days and one year;
- Liabilities to operator systems with a time to maturity between seven days and one year.

are bail-in debt. Such short-term bail-in debt is highly vulnerable to runs. Runs could precipitate a bail-in debt collapse placing stability in grave danger.

Next, we found that pairing short-term debt exclusions with an *increase in the loss absorption requirements* relative to the *status quo* reduces contagion in a system-wide crisis. An increase in the TLAC or MREL loss absorption requirements ensures that sufficient bail-in debt remains after short-term debt exclusions from bail-in have been made. Furthermore, it nudges banks towards lengthening the maturity profile of their debt, making these more stable. Short-term debt exclusions paired with commensurately increased loss absorbing requirements have not been implemented by authorities so far. But the exclusion of short-term liabilities from bail-in would have the effect of either making the hit to creditors who remain in the scope of bail-in more severe in the case of failure, or prompting the authorities to raise the ratio of TLAC/MREL requirements to total liabilities in order to reduce the risk of rendering remaining creditors worse off and/or lacking sufficient loss-absorbing capacity. Since only debt with a time to maturity greater than one year counts towards TLAC/MREL, the latter has the effect of lengthening the maturity profile of bank liabilities. Either way, the proposal to extend the range of *a-priori* exclusions would increase the funding costs of banks; so the bankers are bound to resist.

Moreover, if failure did lead to greater losses to TLAC/MREL creditors, it could make the banking system as a whole more prone to contagion, since it could lead to the collapse of the market for TLAC/MREL liabilities and/or would raise the required interest rate quite sharply. And, again, the more claims on banks are excluded from bail-in, the more likely that those subject to bail-in might claim that the no-creditor-worse-off principle was breached.

Regulators might be hesitant to safeguard a larger chunk of short-term debt from losses upon bank failure, because of worries that this would de-incentive risk monitoring by bank creditors thereby increasing moral hazard. This concern would, however, largely be alleviated if short-term debt is swapped with long-term debt with increased loss absorption requirements. But this latter

[60] In jurisdictions where deposits form the largest share of the banks' liability base, it may not be straightforward to achieve this. It would require a wholesale revamp of the banks' business models. In these jurisdictions more thinking is needed to find a solution.

[61] The main reason for this difference in approach is that the a-priori exclusions are mainly there for two reasons, in order to allow greater operational security and efficiency and, second, to protect retail investors. On the other hand, the definition of TLAC and MREL claims was intended to ensure that banks hold sufficient, easily loss absorbing debt, so that the bail-in tool can be effective in absorbing losses and recapitalising a bank to a desired capital ratio. It was understood that liabilities in longer duration format, and of a kind that could be subject to bail-in with more limited adverse macroeconomic consequences, were easier to bail-in without putting macro stability at risk. So, the considerations based on *a-priori* exclusions, and the definition of TLAC/MREL came from different corners. Nevertheless, it would seem logically considerably tidier and neater to make any liability, subject to bail-in and non-excluded, co-terminous with TLAC and MREL. This could most easily be done by extending a priori exclusions to all bank liabilities with a time to maturity of under one year.

point is typically not thought out.

With the commercial banks probably strongly against extending the range of *a-priori* exclusions, and no constituency strongly in favour, little is likely to be done to change the present situation. But the present situation not only remains somewhat illogical, but it also leads to the likelihood of political pressures for further *ad-hoc* exclusions in the course of any major systemic banking crisis which is being resolved through such a bail-in process.

6.5 The Uncertainty and Speed of Bail-In

We found that *uncertainty over the bank-specific bail-in design* renders markets prone to a panic in a systemic crisis with failing SIBs. Yet, resolution authorities have left considerable uncertainty over the bank-specific bail-in parameters governing a prospective bank bail-in. The effective FLTF threshold and the debt-to-equity conversion rates can take a wide range of values. *Ad-hoc* exclusions, though not modelled in this paper, further compound uncertainty over the bank-specific bail-in design. So did the recapitalisation target. Fortunately, in moving from BRRD I to BRRD II, the resolution authorities seem to have narrowed down the target range. To be sure, some discretion by the resolution authority is warranted, but default values for the bank-specific parameters should exist and not deviated from without good reason. In this way, bail-in debt and thus risk could be priced more securely.

A final important source of uncertainty stems from the speed with which the bail-in process is concluded, which is often slow. A faster Valuation 3 in which equity shares are returned by means of debt-to-equity swap has various (stability) benefits: when a bail-in of a prominent bank is triggered in the course of a systemic downturn, it will likely be seen by the market as a negative shock, and asset values will tumble. That would cause the difference between that bank's prior book value and its current valuation to become even greater, and very possibly to overshoot. If (and only if) Valuation 3 becomes the valuation that becomes public knowledge, while Valuations 1 and 2 remain confidential to the authorities, there could be a case for deferring Valuation 3 until asset markets have stabilised. Even then, however, a faster Valuation 3 has several advantages.[62] For all the reasons set out in footnote 62, our tentative position is that, as a generality, it would be better to complete the debt-to-equity swap (based on Valuation 3) quickly rather than slowly, as the proposers of modern bail-ins originally envisioned (Calello and Ervin (2010)); but it is a fairly

[62] First, in the case where the bank is not insolvent, so equity shareholders become diluted but not wiped out, there would be no advantage whatsoever in delaying Valuation 3, and the completion of the bail-in process. A late valuation 3 poses financial stability risks, as will be noted below. Furthermore, under fair conversion rates no creditor or equity holder will be better off with a late return of equity regardless of whether the bank's equity value further recovers or not. An upswing in the equity value following a bail-in benefits existing and new equity holders even if equity is returned early. Second, deferring Valuation 3 leaves creditors and markets uncertain of the value of their assets in the meantime. Unless bailed-in stakeholders get priority-class-specific formal notes which, even then, might have a relatively low value, they might not be able to realise and access any value in the interim, if such creditors themselves needed to raise extra funds. A temporary exposure loss stemming from a late return of equity renders contagion more likely. Uncertainty as to the ultimate allocation of losses could generate panic in financial markets. Third, the regulatory authorities cannot be sure that deferring Valuation 3 would lead to a higher valuation, and they would look stupid if it did not. Fourth, while there might seem to be a greater confidence effect, if there was a later higher valuation and less hits and write-downs, Valuation 2, which has to take account of the effects of the shock of the failure itself, is known by creditors undergoing haircuts to recapitalise the bank sufficiently, and may become public knowledge, which would mean that the deferment of Valuation 3 would have no beneficial effect on confidence. Fifth, while a later valuation with higher prices would probably benefit existing equity holders and junior creditors, it might harm more senior creditors, so timing in the case where the bank was insolvent at the start of bail-in will affect distribution, rather than benefit all creditors. The reason is that existing equity holders and junior creditors no longer have to be wiped out if the bank's asset value recovers following the bail-in and face less funding pressure. They can now be compensated with equity. Senior creditors will then be in lesser control of the bank.

fine call.

Consider again the incentives of the stakeholders: The accountants, who will be doing the Valuation 3, would like more time, since that should make them less prone to error (and hence law suit), and enable them to charge more. The regulators will also be rushed and their staff over-stretched by pressures to do everything in a great hurry. So both would prefer a later Valuation 3. Those bank managers who are fired will not care one way, or the other; but some managers may be retained to help in completing a lengthy, detailed valuation. An argument will be that rushed valuation would be an incorrect valuation.

7 Conclusion

The contribution of this paper has been to explore the systemic implications of the bail-in design in our networked financial system. We developed a multi-layered network model of the European financial system in which the BRRD framework for resolving banks applies. The model captures the systemic footprint of the bank-specific and structural bail-in design parameters by jointly including the chief endogenous amplification mechanisms. Our quantitative study had the following main results.

- First, we showed that financial stability *hinges* on the bail-in design.
- Second, our results suggest, perhaps surprisingly, that bail-ins could be a credible tool for resolving failing SIBs in system-wide crises, if bail-ins are well-designed. This finding should alleviate concerns that bail-ins are unfit for dealing with failures in (severe) financial crises. And thereby instill regulators with the confidence to opt for well-designed bail-ins rather than bail-outs in the next financial crisis. In contrast, we found that ill-designed bail-ins could render an already bad systemic crisis worse. While our results suggests that the bail-in design matters when resolving a large European SIB, we found that smaller European SIB failures are insufficiently systemic to trigger sweeping contagion irrespective of the bail-in design, consistent with experience (WBG (2017)).
- Third, our results suggest that a crisis-proof bail-in design involves the following stability-enhancing (i.e. "good") bank-specific parameters: an early bail-in, a strong recapitalisation and fair conversion rates. Stability is further reinforced by a trio of "good" structural policy settings. These are: the exclusion of short-term debt with a time to maturity < 1 year from the application of the bail-in tool, higher loss absorption requirements than the *status quo*, and sufficient certainty about the applicable parameters of the bank-specific bail-in design. On the other hand, we found that a crisis-prone bail-in design consists of the opposite parameters. Our findings on the directional effect of bail-in design parameters on financial stability are qualitatively robust to extensive sensitivity analysis.
- Fourth, we showed that the systemic footprint associated with a particular bail-in design risks being underestimated unless multiple interaction contagion mechanisms and non-bank holdings of bail-in debt are captured. We further showed that an ill-designed bail-in amplifies contagious shocks substantially more than a well-designed one. We relatedly showed that a poor design choice added to an otherwise poorly designed bailed-in results in a greater amplification of losses (i.e. is more harmful) than a poor design choice added to an otherwise well-designed bail-in.

Our evidence suggests that a possible pivot towards stability remains in the hands of policymakers, but it also suggests, however, that the current policy parameters might be in the regime of instability. Given political economy considerations it is unlikely that regulators will shift the ill-designed bail-in parameters towards well-designed ones.

If the main concern is to ensure that the banking system remains resilient even in the face of a systemic crisis we would argue, as set out above, that both the failure threshold trigger and the recapitalisation target should be set considerably higher than at present. The reason why this does not happen is because it runs contrary to the incentive structures both of the regulators and, at least as important, of the regulated banks themselves. While, in theory, it might be possible to change the capitalist system fundamentally in order to cause a revision of that incentive structure, that would be such a fundamental change that it lies well outside the remit of this paper. Rather what we might suggest is that the government, should it so wish, could require the regulatory authorities to undertake a special evaluation of any bank which was found to have a PD above a certain value. This would raise the question of manipulation, and there are various ways of dealing with that, e.g. basing the PD on an average of market values, rather than on any particular day, and perhaps not disclosing the algorithmic basis on which the PD was based, or even changing that according to context. Similarly, a government could, if it so wished, require a higher and stronger recapitalisation target ratio, in other words telling the authorities what minimum target ratio they should aim to achieve.

We also suggest four further ways of changing the parameters of bail-in to meet the objective of moving towards the regime of stability, but with somewhat less emphasis and confidence in our approach. These are that exclusions from bail-in should be extended to include short-term claims with a time to maturity less than a year that are currently subject to bail-in but excluded from TLAC/MREL. Short-term bail-in debt is prone to runs in anticipation of bank default risking a bail-in debt collapse, with systemically destabilising consequences, right when the bail-in debt is most needed to recapitalise banks. To compensate for the loss in "loss absorption capacity" stemming from the short-term debt exclusions, the loss absorption requirements should be lifted. This also has the effect of lengthening the maturity profile of bank liabilities as only debt with a time to maturity greater than a year is allowed to count towards the loss absorption requirements, further enhancing stability.

On balance, we would propose, again as a generality, that the speed of Valuation 3 should be as reasonably fast as possible, rather than deferred in the expectation of a widespread market recovery. Quick haircuts to recapitalise the bank but a slow conversion of debt to equity gives affected creditors a loss in the interim and uncertainty as to their recovery, both of which undermine stability. Finally, uncertainty in the bail-in design should be reduced by specifying in advance the bank-specific bail-in parameters that regulators will typically apply in a bail-in. This would enhance the pricing of risk and reduce episodes of financial panic stemming from an inability to estimate loss exposures.

Future Research

Though we have studied carefully the systemic implications of the bail-in design, we have by no means exhausted the research agenda on this topic, on many matters deserving fresh attention or further scrutiny. Perhaps the most pressing need would be to design suitable FLTF triggers, such as one that is partially market-based, that are neither prone to regulatory capture nor to accounting manipulation. What is, unfortunately, obvious is that the main ways of reinforcing the stability of the bail-in process, as earlier described, run counter to the self-interest of bank equity shareholders, amongst them, notably, senior bank managers. A desideratum, therefore, would be a quantification of the probable effect on bank profits of each of these possible stability reinforcing measures. Besides ignoring the effect on bank profitability, our exercise has not been imbedded in a wider model of the real economy. Unless bail-in is triggered very early, it will generally lead to a subsequent shortfall in that bank's TLAC and MREL requirements. There has been very little analysis of the optimal speed of rebuilding, and how that could be enforced.

Other somewhat more technical exercises include: (1) better calibrating the multi-layered network model to data of bail-in debt holdings in the banking and non-banking system, thereby enabling more accurate quantitative insights; (2) assessing the systemic implications of novel bail-in regulations (specifically BRRD II and MREL II) and debt classes (including senior non-preferred debt); (3) comparing the merits and flaws of different bail-in designs in distinct jurisdictions (especially the European Union and the United States); (4) studying the financial stability implications of cross-border and cross-subsidiary resolution, rather than working from a consolidated balance sheet; (5) investigating the systemic implications of two unexplored bail-in design parameters, i.e. the *ad-hoc* exclusions of debt from bail-in and the speed with which to complete a bail-in; and (6) studying deeper the relative efficacy and interaction among bail-in design parameters.

Concluding Note

We end this paper by returning to Ben Bernanke's words: "Have we ended bail-outs? [...] We cannot guarantee that a future administration, fearful of the economic consequences of a building financial crisis, will not authorise a financial bail-out. But the best way to reduce the odds of that happening is to have in place a set of procedures to deal with failing financial firms that those responsible for preserving financial stability expect to be effective" (Bernanke (2017)). The too-big-to-fail problem will be alleviated if bail-in is a credible alternative to bail-out. Our paper shows that *the credibility of bail-in critically depends on the bail-in design.*

References

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

Acharya, V. V. and Beck, T. (2014), The social value of the financial sector: too big to fail or just too big?, Vol. 29, World Scientific.

Acharya, V. V. and Yorulmazer, T. (2007), 'Too many to fail—an analysis of time-inconsistency in bank closure policies', Journal of financial intermediation **16**(1), 1–31.

Admati, A. and Hellwig, M. (2014), The bankers' new clothes: What's wrong with banking and what to do about it, Princeton University Press.

Adrian, T., Brunnermeier, M. K. and Nguyen, H.-L. Q. (2013), 4. hedge fund tail risk, in 'Quantifying Systemic Risk', University of Chicago Press, pp. 155–174.

Adrian, T. and Shin, H. S. (2010), 'Liquidity and leverage', Journal of financial intermediation **19**(3), 418–437.

Adrian, T. and Shin, H. S. (2014), 'Procyclical leverage and value-at-risk', Review of Financial Studies **27**(2), 373–403.

Afonso, G., Kovner, A. and Schoar, A. (2011), 'Stressed, not frozen: The federal funds market in the financial crisis', The Journal of Finance **66**(4), 1109–1139.

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

Almgren, R. and Chriss, N. (2001), 'Optimal execution of portfolio transactions', Journal of Risk **3**, 5–40.

Anand, K., van Lelyveld, I., Banai, Á., Friedrich, S., Garratt, R., Hałaj, G., Fique, J., Hansen, I., Jaramillo, S. M., Lee, H. et al. (2018), 'The missing links: A global study on uncovering financial network structures from partial data', Journal of Financial Stability **35**, 107–119.

Avdjiev, S., Kartasheva, A. and Bogdanova, B. (2013), Cocos: a primer, Technical report, Bank of International Settlements Quarterly Review.

Avgouleas, E. and Goodhart, C. (2015), 'Critical reflections on bank bail-ins', Journal of Financial Regulation **1**(1), 3–29.

Aymanns, C., Caccioli, F., Farmer, J. D. and Tan, V. W. (2016), 'Taming the Basel leverage cycle', Journal of Financial Stability **27**, 263–227.

Aymanns, C., Farmer, J. D., Kleinnijenhuis, A. M. and Wetzer, T. (2018), 'Models of financial stability and their application in stress tests', Handbook of Computational Economics **4**, 329–391.

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

Bernanke, B. (2016), 'Ending "too big to fail": What's the right approach', Brookings Institution's Blog .

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

Bernanke, B. S. (2010), Remarks on the Squam Lake Report – fixing the financial system, Technical report, Squam Lake Conference.

Bernanke, B. S. (2017), Why Dodd-Frank's Orderly Liquidation Authority should be preserved, Technical report, Brookings Institute.

Bernard, B., Capponi, A. and Stiglitz, J. E. (2017), 'Bail-ins and bail-outs: Incentives, connectivity, and systemic stability'.

Berndt, A., Duffie, D. and Zhu, Y. (2020), 'The decline of too big to fail', Available at SSRN 3497897 .

Bertsimas, D. and Lo, A. W. (1998), 'Optimal control of execution costs', Journal of Financial Markets **1**(1), 1–50.

BIS (2015), Revisions to the standardised approach for credit risk, Technical report, Bank for International Settlements, Basel Committee on Banking Supervision.

Black, F. and Scholes, M. (1973), 'The pricing of options and corporate liabilities', Journal of Political Economy **81**(3), 637–654.

Blinder, A. S. (2013), 'After the music stopped', The Financial Crisis, the Response, and the Work Ahead. .

BoE (2017), The Bank of England's approach to resolution, Technical report, Bank of England.

Brownlees, C. and Engle, R. F. (2016), 'Srisk: A conditional capital shortfall measure of systemic risk', Review of Financial Studies pp. hhw060

Brunnermeier, M. K. (2008), Deciphering the liquidity and credit crunch 2007-08, Technical report, National Bureau of Economic Research.

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

Buckmann, M., Gallego Marquez, P., Gimpelewicz, M., Kapadia, S. and Rismanchi, K. (2021), 'The more the merrier? Evidence from the global financial crisis on the value of multiple requirements in bank regulation'.

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., 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.

Calello, P. and Ervin, W. (2010), 'From bail-out to bail-in', https://www.economist.com/finance-and-economics/2010/01/28/from-bail-out-to-bail-in.

Calomiris, C. W. and Herring, R. J. (2013), 'How to design a contingent convertible debt requirement that helps solve our too-big-to-fail problem', Journal of Applied Corporate Finance **25**(2), 39–62.

Carlson, M. A. and Rose, J. (2016), 'Can a bank run be stopped? Government guarantees and the run on Continental Illinois'.

Chen, N., Glasserman, P., Nouri, B. and Pelger, M. (2013), 'Cocos, bail-in, and tail risk', Office of Financial Research Working Paper (0004).

Chennells, L. and Wingfield, V. (2015), 'Bank failure and bail-in: an introduction'.

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

Cochrane, J. H. (2014), 'Toward a run-free financial system', Across the great divide: New perspectives on the financial crisis **197**, 214–15.

Coen, J., Lepore, C. and Schaanning, E. (2019), 'Taking regulation seriously: Fire sales under solvency and liquidity constraints', Bank of England working paper .

Cont, R. and Schaanning, E. F. (2017), 'Fire sales, indirect contagion and systemic stress testing', Norges Bank Working Paper, 2017/2 .

Deloitte (2018), Valuation of difference in treatment: Banco Popular Español, Technical report, Single Resolution Board.

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

Duarte, F. and Eisenbach, T. M. (2021), 'Fire-sale spillovers and systemic risk', The Journal of Finance **76**(3), 1251–1294.

EBA (2017a), Final guidelines on the rate of conversion of debt to equity in bail-in, Technical report, European Banking Authority, EBA/GL/2017/03.

EBA (2017b), Regulatory technical standards on valuation for the purposes of resolution and on valuation to determine difference in treatment following resolution under directive 2014/59/eu on recovery and resolution of credit institutions and investment firms, Technical report, European Banking Authority.

Egan, M., Hortaçsu, A. and Matvos, G. (2017), 'Deposit competition and financial fragility: Evidence from the us banking sector', American Economic Review **107**(1), 169–216.

Eisenberg, L. and Noe, T. H. (2001), 'Systemic risk in financial systems', Management Science **47**(2), 236–249.

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

Engle, R. F. (2020), 'Stress testing with market data', Handbook of Stress Testing, 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., Kleinnijenhuis, A. M., Nahai-Williamson, P. and Wetzer, T. (2020), 'Foundations of system-wide stress testing with heterogeneous institutions', Bank of England Staff Working Paper No. 861 .

Farmer, J. D., Kleinnijenhuis, A. M. and Wetzer, T. (2021), 'Stress testing the financial macrocosm', Forthcoming in Handbook of Financial Stress Testing .

FED (2020), 'SR 20-5: Questions and answers (Q&A) on statement regarding the use of capital and liquidity buffers', https://www.federalreserve.gov/supervisionreg/srletters/sr2005.htm.

Flannery, M. J. and Sorescu, S. M. (1996), 'Evidence of bank market discipline in subordinated debenture yields: 1983–1991', the Journal of Finance **51**(4), 1347–1377.

French, K., Baily, M., Campbell, J., Cochrane, J., Diamond, D., Duffie, D., Kashyap, A., Mishkin, F., Rajan, R., Scharfstein, D. et al. (2010), 'The Squam Lake report: fixing the financial system', Journal of Applied Corporate Finance **22**(3), 8–21.

FSB (2013), Progress and next steps towards ending too-big-to-fail (TBTF), Technical report, Financial Stability Board.

FSB (2016), Guiding principles on the temporary funding needed to support the orderly resolution of a global systemically important bank (G-SIB), Technical report, Financial Stability Board.

Furfine, C. H. (2001), 'Banks as monitors of other banks: Evidence from the overnight federal funds market', The Journal of Business **74**(1), 33–57.

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.

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

Goldstein, I. and Pauzner, A. (2005), 'Demand–deposit contracts and the probability of bank runs', the Journal of Finance **60**(3), 1293–1327.

Goodhart, C. (1975), Problems of monetary management: The UK experience, in A. C. (ed.), ed., 'Inflation, Depression and Economic Policy in the West.'.

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

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

Halaj, G. (2018), 'Agent-based model of system-wide implications of funding risk'.

Hanson, S. G., Kashyap, A. K. and Stein, J. C. (2011), 'A macroprudential approach to financial regulation', Journal of economic Perspectives **25**(1), 3–28.

He, Z. and Krishnamurthy, A. (2019), 'A macroeconomic framework for quantifying systemic risk', American Economic Journal: Macroeconomics **11**(4), 1–37.

He, Z. and Manela, A. (2016), 'Information acquisition in rumor-based bank runs', The Journal of Finance **71**(3), 1113–1158.

Hüser, A.-C., Hałaj, 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 **38**, 81–97.

Jackson, T. H., Scott, K. E. and Taylor, J. B. (2015), Making Failure Feasible: How Bankruptcy Reform Can End Too Big to Fail, Hoover Press.

Kleinnijenhuis, A. M., Kodres, L. E. and Wetzer, T. (2020), 'Usable bank capital', https://voxeu.org/article/usable-bank-capital.

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.

König, E. (2021), 'Single Point of Entry - a resolution strategy addressing the home – host issue in Europe's Banking union'.

Kyle, A. S. (1985), 'Continuous auctions and insider trading', Econometrica: Journal of the Econometric Society **53**(6), 1315–1335.

McAndrews, J., Morgan, D. P., Santos, J. A. and Yorulmazer, T. (2014), 'What makes large bank failures so messy and what to do about it?', pp. 229–244.

Merler, S. (2016), 'Italy's bail-in headache', https://www.bruegel.org/2016/07/italys-bail-in-headache/.

Merton, R. C. (1974), 'On the pricing of corporate debt: The risk structure of interest rates', The Journal of Finance **29**(2), 449–470.

Merton, R. C. (1976), 'Option pricing when underlying stock returns are discontinuous', Journal of Financial Economics **3**(1-2), 125–144.

Merton, R. C. and Thakor, R. T. (2021), 'No-fault default, chapter 11 bankruptcy, and financial institutions', Journal of Banking & Finance p. 106066.

Morrison, A. D. (2011), 'Systemic risks and the 'too-big-to-fail'problem", Oxford Review of Economic Policy **27**(3), 498–516.

Moulds, J., Smith, H., Traynor, I., Elder, M. and Treanor, J. (2013), 'Cyprus banks remain closed to prevent run on deposits', https://www.theguardian.com/world/2013/mar/26/cyprus-banks-closed-prevent-run-deposits.

O'hara, M. and Shaw, W. (1990), 'Deposit insurance and wealth effects: the value of being "too big to fail"', The Journal of Finance **45**(5), 1587–1600.

Pedersen, L. H. (2009), When everyone runs for the exit, Technical report, National Bureau of Economic Research.

Pennacchi, G. (2010), 'A structural model of contingent bank capital', FRB of Cleveland Working Paper 10-04 .

Pennacchi, G. and Tchistyi, A. (2019), 'Contingent convertibles with stock price triggers: The case of perpetuities', The Review of Financial Studies **32**(6), 2302–2340.

Persaud, A. (2014), 'Why bail-in securities are fool's gold', Available at SSRN 2643702 .

Reinhart, C. M. and Rogoff, K. (2009), This time is different: eight centuries of financial folly, princeton university press.

Rutledge, V., Moore, M., Dobler, M., Bossu, W., Jassaud, N. and Zhou, J. (2012), 'From bail-out to bail-in: mandatory debt restructuring of systemic financial institutions'.

Sarin, N. and Summers, L. H. (2016), 'Understanding bank risk through market measures', Brookings Papers on Economic Activity **2016**(2), 57–127.

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

Schuermann, T. (2014), 'Stress testing banks', International Journal of Forecasting **30**(3), 717–728.

Segoviano, M. B. and Goodhart, C. (2009), 'Banking stability measures', IMF working papers pp. 1–54.

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

Sommer, J. H. (2014), 'Why bail-in? and how!', Economic Policy Review .

Stein, J. C. (2013), 'Regulating large financial institutions', https://www.federalreserve.gov/newsevents/speech/stein20130417a.htm.

Stern, G. H. and Feldman, R. J. (2004), Too big to fail: The hazards of bank bailouts, Brookings Institution Press.

Temzelides, T. et al. (1997), 'Are bank runs contagious?', Business review (Nov), 3–14.

Tröger, T. H. (2018), 'Too complex to work: a critical assessment of the bail-in tool under the european bank recovery and resolution regime', Journal of Financial Regulation **4**(1), 35–72.

WBG (2017), Understanding bank recovery and resolution in the EU - A guidebook to the BRRD, Technical report, World Bank Group.

Wiersema, G., Kleinnijenhuis, A. M., Wetzer, T. and Farmer, J. D. (2019), 'Scenario-free analysis of financial stability with interacting contagion channels'.

A Detailed Description of Modelling Framework

Here we motivate how we model each step of the simple stress test depicted in Figure 2 and provide formulas of comprising each modelling step. Our model follows the same philosophy – and frequently uses the same (type of) modelling assumptions – as Greenwood et al. (2015) and

Duarte and Eisenbach (2021). It is useful to have read Section 4.2 describing the model of the bail-in design beforehand. The baseline parameters of the model are summarised in Table 2 in this Appendix. Recall that our framework takes as given institutions' balance sheets, adjustment rules following adverse shocks, and market liquidities, in line with Greenwood et al. (2015) and Duarte and Eisenbach (2021). Our sensitivity analysis in Section 5.4 points out that changing these baseline assumptions does not qualitatively change our findings on the systemic implications of the bail-in design.

t_0 - Adverse stress scenario

We would like to assess the efficacy of different bail-in designs in the case of both idiosyncratic bank failures and financial crises.

Scenario 1: Idiosyncratic bank failure(s)
For the scenario of idiosyncratic failures, We assume a solvency shock to external assets such that the risk-weighted capital ratios of the largest n banks by asset size drops $q = 4\%$ below their failing-likely-to-fail (F) threshold ρ_i^F.[63] If the threshold ρ_i^F is set at 0%, for instance, the bank will be insolvent at the onset of bail-in under this assumption, whereas if the failure threshold is set at $\rho_i^F = 4.5\%$ the bank will still be solvent but in breach of its minimum capital requirements. Instead of a solvency shock, we could have alternatively imposed a liquidity shock resulting in an idiosyncratic bank failure through illiquidity.[64] We will show results for up to $n = 5$ idiosyncratic bank failures to investigate how systemic losses grow as the number of SIB failures goes up. In practise, we think it is unlikely that more than 3 idiosyncratic bank failures happen at a time.

Scenario 2: Financial crisis
To simulate a financial crisis scenario, we impose a scaled version of the system-wide shock used in the 2018 European Banking Authority (EBA) stress test. We vary the magnitude of the scaling factor for the shock over the interval $s \in (0, 2)$, where a scale factor of one corresponds to a shock equal in magnitude to the original shock. For simplicity, contrary to the EBA scenario, we assume asset losses hit banks' external assets only (in line with Cont and Schaanning (2017)), resulting in a diminished capital ratio, whereas in the EBA scenario all of the banks' assets suffer losses as a consequence of adversely set values of macroeconomic and financial variables. The 2018 EBA results specify both the pre-shock ρ_i^{data} and the post-shock capital ratio $\rho_i^{EBA,adverse}$ of each bank. Given these, the system-wide shock s in our model gives each bank i a risk-weighted capital ratio equal to $\rho_i^{t_0} = \rho_i^{data} - s(\rho_i^{data} - \rho_i^{EBA,adverse})$ as a direct consequence of the systemic shock scenario. $\rho_i^{t_0}$ should be thought of as the capital ratio of a bank i following the microprudential stress test results but before any spill-overs from contagion have been modelled. If $s = 0$ (no shock) then $\rho_i^{t_0} = \rho_i^{data}$, so the capital ratio is equal to that in normal times. If $s = 1$ then $\rho_i^{t_0} = \rho_i^{EBA,adverse}$, so the capital ratio is equal to capital ratio following the EBA shock. If $s = 2$, then the scenario is twice as severe as the EBA shock.

The banks' average risk-weighted capital ratio before the system-wide shock takes place is 15% (with a standard deviation of 3.6%). After the system-wide shock takes place of severity $s = 1$, the average is 11.5% (with a standard deviation of 3.5%). Thus the original EBA shock

[63] The results are qualitatively robust to this choice. A larger q simply means that more losses have to be absorbed with haircuts to recapitalise a bank to its recapitalisation target, resulting in a bigger shock to the system. It does not qualitatively affect our results.

[64] Compared to the solvency shock we impose such shock gives more stress to the financial system at the outset because the bank has to liquidate assets in a disorderly manner if it does not have a sufficient cash buffer to meet the repayment obligations. Changing the exogenous idiosyncratic shock to a liquidity shock does not qualitatively change our results.

corresponds to a loss of about 23% in the risk-weighted capital ratio relative to the pre-shock ratio. The reductions in bank capital ratios are a consequence of adverse shocks the 2018 EBA scenario imposes on variables, such as GDP, inflation, unemployment, asset prices, and interest rates (ESRB (2018)). To give an idea of the severity of the 2018 EBA shocks, equity prices are projected to decline by 41% and 30% in the United States and European Union, respectively. EU GDP falls by 2.5% to 7.4% relative to the baseline level, residential and commercial property prices fall by approximately 27%, and yields on long-term rates are assumed to rise sharply, by 235 basis points in the United States and 83 in Europe.

$\mathbf{t_x^a}$ - Impact

To next step involves determining the impact of exogenous shocks and endogenous shocks (if losses are contagious spillovers from step t_x^b) on the balance sheets of institutions, as well as determining whether a bank failure has occurred as a consequence of these shocks. These exogenous shocks also result in valuation losses to external assets. Consequently, these shocks result in a revaluation of bail-in debt. The endogenous shocks emanating from contagious spillovers of institutions' reactions to shocks undertaken in step t_x^b include:

1. Liquidity shocks from bail-in debt or non-bail-in debt that is not rolled over;
2. Net exposure losses on bail-in debt;
3. Mark-to-market losses on common asset holdings;
4. Revaluation of bail-in debt.

We now discuss the impacts of exogenous shocks in detail and will detail the impacts of endogenous shocks in our discussion of step t_x^b.

(Re)pricing of bail-in debt
The exogenous or endogenous shocks give losses to a bank's assets at time t_x^a resulting in a downwards re-valuation of its bail-in debt. Our novel method for pricing bail-in debt uses a standard approach of discounting the expected value of future payoffs, in line with Black and Scholes (1973) and Merton (1974). Under our valuation method, bail-in debt becomes less valuable if its future pay-off is likely to decrease. The value of a bank's bail-in debt thus drops if its failure becomes more likely and/or if the pay-off to its creditors in a prospective bail-in will be reduced. Under some bail-in designs $\mathcal{D}$ the pay-off in a priority class will be less than in others. Both the bail-in design and the bank's capital ratio following the exogenous and endogenous shocks will thus influence the value of bail-in debt in our model.

With the intuition of what drives bail-in debt valuation made clear, we now proceed to describe our pricing method. We price bail-in debt with notional amount B_{ji}^{km} as the discounted expected value of its future pay-offs. The value V_{ji}^{kmt} at time t of j's claim on bank i's bail-in debt in priority class k which matures at time $T = t + m$ is given by

$$V_{ji}^{kmt} = \exp^{-r(T-t)} \mathbb{E}^{\mathcal{Q}}[P_{ji}^{kT}(A_i^{\tau}, A_i^{T})]. \tag{9}$$

Equation 9 shows that the value V_{ji}^{kmt} at time t of a bail-in debt B_{ji}^{kmt} claim is given by the expected pay-off P_{ji}^{kT} under the risk-neutral measure $\mathcal{Q}$ discounted by the risk-free rate r.[65] The value V_{ji}^{kmt}

[65] For simplicity we take $r = 0$ in line with the current low interest rate environment. We could have just as easily taken the 30Y treasury yield as a proxy for the risk-free rate without modifying our results qualitatively.

of a bail-in debt claim B_{ji}^{km} is approximated by the average pay-off P_{ji}^{kTn} over $n = 1, ..., N$ Monte Carlo runs. That is,

$$V_{ji}^{kmt} \approx \exp^{-r(T-t)} \frac{1}{N} \sum_{n=1}^{N} P_{ji}^{kTn}(A_i^{\tau}, A_i^{Tn}). \tag{10}$$

The pay-off P_{ji}^{kTn} at time T in Monte Carlo run n is given by the notional of the bail-in debt claim B_{ji}^{kmt} at time t in the stress test if no bail-in takes place prior to the maturity T of the contract.[66] This is the case if the risk-weighted capital ratio ρ_i^{sn} in Monte Carlo run n stays above the failure threshold ρ_i^F in the time interval $s \in [t, T]$.[67] If not, then the pay-off P_{ji}^{kTn} is given by bail-in pay-off, which consists of the sum of the haircutted notional $(B_{ji}^{kmt} - h_{ji}^{km\tau_a n} - h_{ji}^{km\tau_b n})$ and the returned share $\epsilon_{ji}^{km\tau_b n}$ of the bank's post-bail-in capital $E_i^{\tau_b n}$. We have that

$$P_{ji}^{kTn} = \begin{cases} B_{ji}^{kmt}, & \text{if } \rho_i^{sn} \geq \rho_i^F, \ \forall s \in [t, T] \\ B_{ji}^{km\tau} - h_{ji}^{km\tau_a n} - h_{ji}^{km\tau_b n} + \epsilon_{ji}^{km\tau_b n} E_i^{\tau_b n}, & \text{if } \rho_i^{sn} < \rho_i^F, \text{ for a } s \in [t, T] \end{cases} \tag{11}$$

where $B_{ji}^{km\tau} = B_{ji}^{kmt}$ in the lower line, since the notional claim does not change in the Monte Carlo simulation unless a bail-in takes place.

We highlight two implicit assumptions we make when valuing the received equity claim $E_{ji}^{km\tau_b n} = \epsilon_{ji}^{km\tau_b n} E_i^{\tau_b n}$. First, the equity claim $E_{ji}^{km\tau_b n}$ can be converted to cash right after the completion of bail-in, at time τ_b, at no liquidation cost (i.e. no price impact), so that the equity claim can be counted as a payoff. Second, the new equity holder wishes to immediately convert its equity claim $E_{ji}^{km\tau_b n}$ to cash after bail-in rather than at a later time, for instance at the original maturity date of the contract T.[68][69] Another assumption we make is that equity will be valued at its book value (see Section 4.2.4), in line with Greenwood et al. (2015).

To evaluate the pay-off P_{ji}^{kT} of the bail-in contract contract B_{ji}^{km}, we need a process that

[66] We make the admittedly strong assumption that the yield on bail-in debt is zero. This means our model cannot capture how issuing bail-in debt in times of distress might become more costly. We leave incorporating bail-in yields for future research.

[67] Our pricing method does not incorporate (near) illiquidity as a reason for bailing-in a bank, though in the stress test simulation this does count as a reason to bail-in a bank. We leave the incorporation of illiquidity (as reason for bailing-in a bank) in bail-in debt pricing for future research.

[68] We note that the right time to count the cash flow arising from the debt-to-equity conversion for the purpose of pricing bail-in debt is not obvious. The CoCo literature, for instance – which has provided us with the analogue for pricing bail-in debt – does not discuss what the right time is to count any cash flows associated with the new equity claim, as far as we are aware.

[69] An alternative to our current assumption is to assume that the new equity holder liquidates its equity claim to generate a cash flow only at the maturity of the original debt contract. Here we work out the pay-off in the case where equity holders liquidate their new equity claim only at the maturity date T of the original contract. In this case, the received equity claim E_{ji}^{kmTn} should be valued at time T instead. Further, the receivers of bank equity will now be exposed to asset losses resulting from any reduction in the value of the bank's CET1 equity E_i^n following the bail-in. In such case, even if the bail-in haircuts and equity compensation do not result in net losses in book-value terms as a direct consequence of the bail-in, claimants may nonetheless suffer losses due to a drop in the bank's equity after a bail-in. The pay-off P_{ji}^{kTn} of a bail-in claim in bail-in (i.e. if $\rho_i^{sn} < \rho_i^F$, for a $s \in [t, T]$) is then given by

$$\begin{aligned} P_{ji}^{kTn} &= B_{ji}^{km\tau_b n} + \epsilon_{ji}^{km\tau_b n}(E_i^{\tau_b n} + \sum_{s=\tau_i^{mn}}^{T} (E_i^{s+1,n} - E_i^{sn})) \\ &= B_{ji}^{km\tau_b} + \epsilon_{ji}^{km\tau_b n}(E_i^{\tau_b n} + E_i^{Tn} - E_i^{\tau n}), \end{aligned} \tag{12}$$

where $B_{ji}^{km\tau_b} = B_{ji}^{km\tau} - h_{ji}^{km\tau_a n} - h_{ji}^{km\tau_b n}$ and $E_i^{\tau_b n}$ denotes value of bank i's equity following the hypothetical bail-in in simulation n. Further, E_i^{Tn} and $E_i^{\tau n}$ denote the value of bank i's equity in simulation n at the maturity T of the bail-in contract and the start of the bail-in τ_i^{mn}. Both values are taken to be not affected by bail-in and follow the original Monte Carlo path. Yet another alternative is to assume that an equity holder never liquidates its equity claim and instead reaps the dividends of the stock over time.

governs the evolution of the asset value A_i. In line with Pennacchi (2010), who introduces a way to price the contractual analogue of bail-in debt, CoCos, we propose to model the asset value A_i evolution according to a jump process. A jump process captures, unlike a simple geometric Brownian motion, that asset values are prone to sizable asset losses, which is especially common in financial crises on which we focus. Specifically, we apply Merton's jump-diffusion process (Merton (1976)) with log-normal jumps. That is, the risk-neutral jump process for the asset value A_i follows

$$\frac{dA_i^t}{A_i^t} = (r - \lambda_i \bar{j}_i)dt + \sigma_i dW_i^{\mathcal{Q}t} + j_i dq_i^t,\ A_i^t = A_i^{t_s}. \tag{13}$$

The solution to equation 13 is given by

$$A_i^{t+1} = A_i^t \exp^{(r-\lambda_i \bar{j}_i - \frac{\sigma_i^2}{2})\Delta t + \sigma_i \sqrt{\Delta t} Z_i)} (1 + j_i)^{\phi_i},\ A_i^t = A_i^{t_s}. \tag{14}$$

The initial condition in equation 13 stipulates that asset value A_i at time t equals the asset value in the stress test simulation at time t, which is denoted by t_s. The jump events of bank i are governed by a compound Poisson process q_i^t with jump intensity λ_i, which gives the mean number of arrivals per unit time t. The magnitude of bank i's random jump is given by j_i, where $1+j_i$ is log-normally distributed with mean μ_i^J and standard deviation σ_i^J. The magnitude should be interpreted as the fractional increase or decrease in the bank's total assets A_i. The instantaneous variance of bank i's asset returns conditional on the Poisson event not occurring is given by σ_i^2. $dW_i^{\mathcal{Q}t}$ is a standard Gauss-Wiener process under the risk-neutral dynamics. Further, dq_i^t and $dW_i^{\mathcal{Q}t}$ are assumed to be independent. Z_i is distributed according to the standard normal distribution. ϕ_i equals to one with probability λ_i and equals to zero with probability $1 - \lambda_i$. Hence, a jump occurs if $\phi_i = 1$ and no jump occurs if $\phi_i = 0$.

Using equation 14, the path of bank i's asset value A_i can be generated over the remaining lifetime of the bail-in contract B_{ji}^{km} ($s \in [t, T]$). When we impose a invariant-liability assumption similar to that used in Merton's structural credit risk model (Merton (1974)), we can run n Monte Carlo paths of the equity value E_i and risk-weighted capital ratio ρ_i. Under this assumption, bail-in debt and liabilities remain equal to their time-t value up to a bail-in, and stay equal to their post-bail-in value τ_b up to maturity T; that is, $B_i^s = B_i^t$ and $L_i^s = L_i^t$, $\forall s \in [t, \min\{\tau_i^{mn}, T\}]$; and $B_i^s = B_i^{\tau_b n}$ and $L_i^s = L_i^{\tau_b n}$, $\forall s \in [\tau_b^{mn}, T]$. To generate the equity E_i^n path in run n, we also use the invariant-liability assumption, as well as equation 2, 3 and 4 in Online Appendix B. To generate the risk-weighted capital ratio ρ_i^n path in run n, we again use the invariant-liability assumption, as well as the definition of ρ_i (defined in Section 4.2.1) and the approximation $\Omega_i^{t+1} \approx \frac{A_i^{t+1}}{A_i^t}\Omega_i^t$. These generated sample paths allow us to compute the payoff P_{ji}^{kTn} (see equation 11) in each Monte Carlo run n.

Figure 8, 9, and 10 in Online Appendix D.3 show and discuss how the value V_{ji}^{kmt} of a bail-in debt contract depends on: (i) the jump-process parameters (λ_i, σ, μ_i^J and σ_i^J); (ii) the seniority class k of bail-in debt; and (iii) a selection of bank-specific bail-in parameters, as well as the size of adverse shock.

Determination of bank failure (FLTF).
Exogenous and endogenous shocks can also lead to bank failure. A bank i will be bailed-in at stopping time τ_i whenever it is failing or likely to fail (see equation 1 in Section 4.2.1).

$\mathbf{t_x^b}$ - Action

The impacts of adverse shocks on balance sheets prompt institutions to act to avoid (getting close to) default, and prompt the regulator to bail-in SIBs that have been determined to be failing-likely-to-fail.

Bail-in(s) with applicable design $\mathcal{D}$
Given that structural design choices have already been made in advance of any bail-in taking place, bank-specific design choices left to decide on at the onset of the bail-in of bank i at time τ_i are to what recapitalisation target ρ_i^{RT} the bank should be recapitalised, and how creditors who received a haircut should be compensated with a debt-to-equity swap Δ_i^k applicable in each priority class k. As noted, we assume *ad-hoc* debt exclusions are zero. We detail our model of the bail-in mechanism and its design in Section 4.2 and Online Appendix C.

The net exposure loss to creditor j – holding a bail-in contract in priority class k with maturity m of notional $B_{ji}^{km\tau}$ – resulting from bailing in bank i at time τ is given by the positive difference of the value of its equity compensation resulting from the debt-to-equity swap $E_{ji}^{km\tau_b} = (\epsilon_{ji}^{km\tau_b} E_i^{\tau_b})$ and its haircuts $h_{ji}^{km\tau} (= h_{ji}^{km\tau_a} + h_{ji}^{km\tau_a})$.[70] The loss suffered by an existing equity holder j whose ownership in bank i is diluted or cancelled as a consequence of the bail-in is given by any positive difference between $E_{ji}^{k_1\tau}$ and $E_{ji}^{k_1\tau_b}$, where k_1 is the priority class corresponding to CET1 equity. Any net exposure losses (or gains) from bail-in will result in valuation shocks to balance sheets of bailed-in investors at time t_{x+1}^a.

Bank and non-bank reactions to shocks
As noted, our model takes the adjustment rules following adverse shocks as a given, in line with Greenwood et al. (2015) and Duarte and Eisenbach (2021). Institutions first act to meet payment obligations, then act to avoid obtaining a dangerously risky capital ratio, and finally act to reduce exposures to excessively risky bail-in debt in our model. This order of action reflects the priority that needs to be given to liquidity issues over solvency issues to avoid failure. While asset value uncertainty and regulatory forbearance may help a bank who breaches its minimum capital requirements or is insolvent once its losses are fully recognised keep its licence to operate as a bank and stay out of resolution, illiquidity kills an institution fast (recall, for instance, the failure of Lehman Brothers (Brunnermeier (2008))). A payment obligation that is not met constitutes a default.

(1) Fulfilling payment obligations
If a counterparty j stops rolling over a bail-in debt contract to bank i with notional B_{ji}^{mkt} at the maturity of the contract (i.e. if the time to maturity is zero, $m = 0$), then bank i has to repay B_{ji}^{mkt} amount of cash the next time period, otherwise it will fail and be bailed-in because of illiquidity. In our model repayment obligations arise for two reasons: (1) if an investors stops rolling over bail-in funds (see reaction 3); and (2) if a bank or non-bank raises cash by reducing maturing funding (see action 2) in order to de-lever.

The order in which assets are liquidated to raise cash is dictated by a liquidity "pecking order" to minimise liquidation costs. In line with Cifuentes et al. (2005), Halaj (2018), we assume that assets that are least costly to reduce will be reduced first to raise cash to meet payment obligations. Practically, this means that an institution i will draw down its cash buffer first C_i^t, pull back maturing bail-in funding B_i^t (i.e. contracts for which $B_{ij}^{k,m=0,t}$) and non-bail-in funding next, and resort to liquidating securities T_i^t last, as the latter may have a price impact but the

[70] If the difference is negative creditor j enjoys a profit as a consequence of the bail-in. It may still suffer a loss later if bank i's equity value falls.

former do not.[71] Within an asset class, we assume that an institution i will raise a given amount of cash by reducing individual contracts proportionally (in line with Greenwood et al. (2015), Duarte and Eisenbach (2021); these authors assume proportional liquidation across asset classes rather than within an asset class as we do). Cifuentes et al. (2005), Greenwood et al. (2015) and Farmer et al. (2020) show that altering the liquidation pecking order does not qualitatively change their results.

Tradable assets must be marked-to-market on institutions' balance sheet. In our model, the price impact of asset sales is approximated with a price impact function that is linear in the net asset sales, in line with Kyle (1985), Bertsimas and Lo (1998), Almgren and Chriss (2001), Greenwood et al. (2015), Duarte and Eisenbach (2021) and Cont and Schaanning (2017). Given this approach, the price p^t_{am} at time t of an individual security m of type $a \in \mathcal{A}$, (where set $\mathcal{A}$ consists of the following asset types: government bonds, corporate bonds, equities, other securities) is given by

$$p^t_{am} = p^{t0}_{am}(1 - \beta_{am} f^t_{am}).^{72} \tag{15}$$

Here $f^t_{am}\%$ denotes the cumulative percentage of net asset sales of asset m of type a up to time t, relative to the asset's market capitalisation, and β^t_{am} is the asset's price impact parameter. In line with Schnabel and Shin (2004), Cifuentes et al. (2005) and Gai and Kapadia (2010), our model takes as a given that the price of a security falls by 5% if 5% of the market capitalisation has been sold on a net basis. This corresponds to picking $\beta_{am} = 1$. In line with Greenwood et al. (2015), our price impact parameter is uniform across assets $\{a, m\}$. Any mark-to-market losses resulting from reductions in the price of assets will be felt as valuation shocks on the balance sheets of institutions at time t^a_{x+1}.

Recall that balance sheets of banks and non-banks are described in Section 2 and detailed in Online Appendix B.

(2) De-levering to maintain a stable (risk-weighted) leverage ratio

We take as a given that banks will de-lever in response to asset losses in order to maintain a stable risk-weighted capital ratio ρ_i, stay away from their failure threshold, and avoid regulatory penalties applicable in the regulatory buffer zone. Non-banks, likewise, will de-lever to maintain a stable leverage ratio λ_i in order to avoid margin calls (margin calls are not modelled). In our model, banks de-lever to return to their internal target ρ^T_i whenever their losses eat into over $u = 50\%$ of their Basel III "combined capital buffer" ρ^{CB}_i, in line with Farmer et al. (2020).[73] Non-banks in our model de-lever to their target λ^T_i whenever their leverage ratio falls below 95% of their initial leverage ratio $\lambda^{t_0}_i$, in line with Cont and Schaanning (2017).[74] The assumption of

[71] We assume that the discount rate and funding cost of bail-in debt are zero. The former is a reasonable assumption in today's low interest rate environment. The latter is less realistic as interest rates on bail-in debt can steeply rise with financial distress, incorporating this is left for future research.

[72] An alternative specification of a linear price impact function used by e.g. Greenwood et al. (2015) and Cont and Schaanning (2017) is that the relative price change of asset $\{a, m\}$ following liquidation size q is given by $\frac{\Delta p_{am}}{p_{am}} = -\psi_{am}(q)$, where $\psi_{am}(q) = \frac{q}{D_{am}}$. The market depth is given by $D_{am} = c\frac{ADV_{am}}{\sigma_{am}}$, where ADV_{am} is the average daily trading volume and σ_{am} is the daily volatility of asset $\{a, m\}$. Greenwood et al. (2015) assume $D_{am} = 10^{13}$ meaning that 10bn euros of trading imbalances lead to a price change by 10bp. Duarte and Eisenbach (2021) formulate a linear price impact function capturing how the wealth of potential buyers of a security alleviates the price impact generated by (forced) sellers of a security, in line with Shleifer and Vishny (2011).

[73] As explained in Farmer et al. (2020), a bank's combined risk-weighted capital buffer is given by: $\rho^{CB}_i := \rho^{CCoB}_i + \rho^{CCyB}_i + \max\{\rho^{GSIB}_i, \rho^{DSIB}_i, \rho^{SR}_i\}$, where ρ^{CCoB}_i is the capital conservation buffer, ρ^{CCyB}_i is the applicable countercyclical capital buffer (if any); and ρ^{GSIB}_i, ρ^{DSIB}_i, and ρ^{SR}_i are the G-SIB surcharge, D-SIB surcharge, and systemic risk buffer respectively. For each bank in our stress test we use the values for ρ^{CB}_i obtained by Farmer et al. (2020).

[74] Our findings are qualitatively robust to the choice of buffer threshold at which an institution starts to de-lever

leverage targeting is widely used in the contagion literature (see e.g. Greenwood et al. (2015) and Duarte and Eisenbach (2021)) and grounded in the empirical research of Adrian and Shin (2010) who provide evidence showing that banks maintain a relatively stable capital ratio over time and manage leverage to offset shocks to asset values.[75] In line with Duarte and Eisenbach (2021), we assume that a bank returns to its internal target gradually: by no more than $\Lambda_i\%$ per time step, thereby limiting self-inflicted harm induced by high liquidation costs and contagious feedbacks (Caccioli et al. (2014) and Aymanns et al. (2016)).

In sum, once a bank's capital ratio ρ_i^t falls below its internal buffer ρ_i^B capital ratio given by its minimum capital requirement ρ_i^M plus one minus its perceived usability of regulatory buffers u_i times its combined regulatory buffer ρ_i^{CB}, i.e.

$$\rho_i^B = \rho_i^M + (1 - u_i)\rho_i^{CB}, \tag{16}$$

it will seek to gradually return to a target capital ratio ρ_i^T (where we must have that $\rho_i^T \geq \rho_i^B$) given by its minimal capital requirement plus $v_i \in \mathcal{R}_0^+$ times its combined regulatory buffer ρ_i^{CB}, i.e.

$$\rho_i^T = \rho_i^M + v_i \rho_i^{CB}, \tag{17}$$

with its intermediate target capital ratio $\tilde{\rho}_i^T$ given by

$$\tilde{\rho}_i^T = \min\{\rho_i^t + \Lambda_i, \rho_i^T\}.^{76} \tag{18}$$

The intermediary target is applicable until the bank hits its target ρ_i^T, after which it will act only once it falls below its internal buffer ρ_i^B again.

As explained in Section 4.2.1, for simplicity we assume that a bank is subject to only one capital requirement, the Pillar I minimum CET1 capital ratio ρ_i^M. In practise, under Basel III a bank is also subject to a minimum capital ratio on CET1, AT1 and T2 capital combined, as well as Pillar II additional capital requirements. Our model could easily incorporate additional requirements.

The leverage target of a non-bank is given by λ_i^T to which it seeks to return once its leverage λ_i^t falls below its leverage buffer $\lambda_i^B = 0.9\lambda_i^T$, in line with Cont and Schaanning (2017). Were we to choose $\lambda_i^B = \lambda_i^T$ instead, then we revert to Greenwood et al. (2015).

From observing the set-up of our model it can be understood that our (risk-weighted) leverage targeting model offers three methodological contributions relative to Greenwood et al. (2015): (1) institutions act to revert to target only once they move too far below their target; (2) institutions return to target gradually, in line with Duarte and Eisenbach (2021); and (3) banks target a risk-weighted capital ratio. The logic of the first has already been explained, so we proceed to explain the logic of the second. The one-sided nature of constraints means that institutions react asymmetrically to large losses and large gains (Cont and Schaanning (2017)). In the light of this and because adjusting a portfolio involves trading costs, we assume that institutions de-lever only

to return to its target capital ratio.

[75] Their evidence suggests that banks do not issue new equity following a negative asset shock. A contributing factor is that if debt overhang is severe, raising equity dilutes existing shareholders as the gains from the reduction in risk accrue disproportionately to debt holders. Another contribution factor, which holds especially in times of crisis, is that issuing equity may be infeasible or financially unappealing.

[76] An alternative way to build the partial adjustment model to a risk-weighted leverage target is as Duarte and Eisenbach (2021) have done. Under their approach, the new risk-weighted capital ratio will be set equal to $\rho_i^{t+1} = b\rho_i^T + (1-b)\rho_i^{p,t+1}$ where ρ_i^T is the target, ρ_i^{t+1p} the passive risk-weighted capital ratio, and b the adjustment speed. $b = 1$ means the bank moves immediately to target, whereas $b = 0$ means the bank follows a passive strategy for its risk-weighted capital ratio.

once they get too close to their binding constraints.

Banks face capital constraints prior to hitting their minimum regulatory capital ratio (Goodhart (2013) and Farmer et al. (2020)), including market-based capital constraints and regulatory constraints. Basel III introduced regulatory capital buffers, which sit on top of capital requirements, which are explicitly meant to be "usable" in the sense that banks can use these buffers to absorb losses. However, since using regulatory buffers comes with a penalty in the form of restrictions in discretionary payments (such as dividend and bonus payments) – which get increasingly more severe as more quadrants of the buffer are used up – banks have been shown hesitancy to use their regulatory buffers (Kleinnijenhuis et al. (2020)). We therefore assume that banks start to de-lever whenever their losses eat up $u\%$ of their regulatory buffers, where $u = 0\%$ means banks are not willing to use their regulatory buffers at all.[77] So, our de-leveraging model reflects that regulatory buffers can be seen as semi-binding. $u\%$ can alternatively be interpreted as determining the point where the market-based capital constraint becomes binding.

Another reason why we adopt a leverage target model in which institutions only revert back to target once they fall too far below their target is that Cont and Schaanning (2017) have shown that leverage targeting models (where $\lambda_i^B = \lambda_i^T$) overestimate the magnitude of liquidations, especially at smaller shock levels, but underestimate the acceleration (convexity) of liquidations for larger shock sizes, present in the threshold model.

The logic of the third innovation stems from the observation that in 2018 the risk-weighted capital ratio rather than the leverage constraint determined whether banks passed the EBA 2018 stress test.[78] The risk-weighted capital ratio was thus the effective binding constraint for EBA banks in 2018. For (hedge) funds we stick to leverage since they do not face a risk-weighted capital requirement. Using the risk-weighted capital ratio rather than a leverage ratio to determine whether a bank is failing-to-likely-to fail, from a solvency perspective, is moreover in line with Hüser et al. (2017).

De-levering amount. A non-bank returns to its leverage target λ_i^T by delevering d_i^t amount given by

$$d_i^t = [A_i^t \frac{1}{\lambda_i^T} - \hat{E}_i^t] \mathbb{1}_{\{\ \lambda_i^t \leq \lambda_i^B\}}, \tag{19}$$

which follows from

$$\lambda_i^T = \frac{(A_i^t - d_i^t) - (L_i^t - d_i^t)}{A_i^t - d_i^t} \mathbb{1}_{\{\ \lambda_i^t \leq \lambda_i^B\}} = \frac{\hat{E}_i^t}{A_i^t - d_i^t} \mathbb{1}_{\{\ \lambda_i^t \leq \lambda_i^B\}}, \tag{20}$$

where A_i^t, L_i^t and $\hat{E}_i^t$ represent the assets, liabilities, and book equity of non-bank i (for comparison, recall that E_i^t denotes the CET1 equity), respectively. It delevers d_i^t amount by liquidating d_i^t amount of assets according to the liquidation pecking order described in action (1).

Risk-weighted de-levering amount. A bank's risk-weighted capital ratio is given by its CET1 equity E_i^t over its risk-weighted assets Ω_i^t, where its risk-weighted assets are given by the weighted sum of its assets carrying different risk-weights under Basel III, i.e. $\Omega_i^t = \sum_{p \in \mathcal{P}} \omega_p A_{ip}$. The set of asset types $\mathcal{P}$ is defined as $\mathcal{P} := \{1, ..., 9\}$, where $\{A_{i1}, ..., A_{i9}\} =$
$\{C_i, Y_i, T_{ia}^1, T_{ia}^2, T_{ia}^3, T_{ia}^4 I_i, R_i, O_i\}$. Asset type $p = 1, ..., 9$ correspond respectively to cash, external assets, government bonds, corporate bonds, equities, other securities, interbank assets, reverse repos, and other assets. Each asset type A_{ip} bears a different risk weight ω_p under the standard

[77]Farmer et al. (2020) shows that financial stability increases if banks are more willing to use their regulatory buffers.

[78]In fact, multiple banks saw their leverage ratio drop below the Basel III 3% minimum leverage ratio and still passed the stress test.

approach in Basel III. In line with the Basel III standardised approach (BIS (2015)), we set the risk weights ω_p for $p = 1, ..., 8$ (i.e. except $p = 9$) equal to $\{0, 0.35, 0, 1, 0.75, 1, 0.4, 0.1\}$. For the "other asset" class $p = 9$ we choose the risk-weight such that it acts as a balancing item to ensure that total RWAs Ω_i^t match the 2018 EBA data. (See Online Appendix B for a description of the asset side of bank balance sheets and its correspondence to bail-in debt holdings.)

A bank returns its risk-weighted capital ratio target ρ_i^T whenever its risk-weighted capital ratio ρ_i falls below its buffer ρ_i^B and it is not in resolution (i.e. $\rho_i \geq \rho_i^F$ and $l_i^t \geq 0$, see Section 4.2.1). We assume that a bank returns to its target ratio ρ_i^T by reducing assets A_{ip} with non-zero risk-weights starting with the most high risk-weight assets, as this may present an effective way to quickly get back to its capital ratio target ρ_i^T.[79] We assume that external assets, even though they carry a positive risk weight, will not be liquidated to improve the risk-weighted capital ratio, in line with Cont and Schaanning (2017). External assets consists largely of illiquid assets, such as mortgage loans with long-maturities, that are either not easily marketable or would be subject to a steep discount if they were to be sold during a stress scenario.

Given the ranking of risk-weights, the iterative method a bank uses to aim reaching its target ρ_i^T is as follows (we assume $\omega_{p_9} = 0$ here, but the iterative method can easily be generalised for any $\omega_{p_9} > 0$). It liquidates $\hat{r}_{ip_4}$ amount of asset type A_{ip_4}. It can never reduce more assets than it assets A_{ip_4} of this type. That is, $\hat{r}_{ip_4}$ is given by $\hat{r}_{ip_4} = \min\{r_{ip_4}, A_{ip_4}\}$, where r_{ip_4}

$$r_{ip_4} = \frac{1}{\omega_{p_4}}[\sum_{p\in\mathcal{P}} \omega_p A_{ip} - \frac{E_i}{\rho_i^T}], \tag{21}$$

and follows from

$$\rho_i^T = \frac{E_i}{\omega_{p_4}(A_{ip_4} - r_{ip_4}) + \sum_{p\in\mathcal{P}\setminus p_4} \omega_p A_{ip}}. \tag{22}$$

If $\hat{r}_{ip_4} < r_{ip_4}$ then the bank did not have enough assets A_{ip_4} 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

$$r_{ip_6} = \frac{1}{\omega_{p_6}}[\sum_{p\in\mathcal{P}} \omega_p A_{ip} - \sum_{p=p_4} \omega_p A_{ip} - \frac{E_i}{\rho_i^T}]\mathbb{1}_{\{\hat{r}_{ip_4} < r_{ip_4}\}} \tag{23}$$

We observe that the amount of 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 $\hat{r}_{ip_8} = \min\{r_{ip_8}, A_{ip8}\}$, where r_{ip_8} is given by

$$r_{ip_8} = \frac{1}{\omega_{p_8}}[\sum_{p\in\mathcal{P}} \omega_p A_{ip} - \sum_{p=p_4,p_6,p_5,p_7} \omega_p A_{ip} - \frac{E_i}{\rho_i^T}]\mathbb{1}_{\{\hat{r}_{ip_x} < r_{ip_x},\ \text{for } x=4,6,5,7\}} \tag{24}$$

[79] Our sensitivity analysis shows that adopting another strategy for reducing risk-weighted assets does not qualitatively change our results. Coen et al. (2019) show that it can be optimal for banks wishing to minimise liquidation losses to sell more liquid assets, rather than potentially less liquid assets with a higher risk weight as in our set-up.

If the following condition is true a bank cannot fully reach its target

$$\frac{E_i}{\sum_{p\in\mathcal{P}}\omega_p A_{ip} - \sum_{p=p_4,p_6,p_5,p_7,p_8,p_9}\omega_p A_{ip}} < \rho_i^T \tag{25}$$

by reducing its risk-weighted assets Ω_i^t. It will then resort to de-levering by using its cash to pay back its liabilities, as in equation 19.

(3) Reducing exposures to too risky bail-in debt

Since the conditions under which bail-in creditors will dash for cash are unknown to resolution authorities, we pose, as in the adverse scenario of a regulatory stress test (Aymanns et al. (2018)), three plausible run scenarios *in anticipation* of a bail-in or *following* the completion of bail-in. We note that the "public funding backstop mechanism", available to G-SIBs solely *during* the bail-in (FSB (2016)), is of no help in either of these cases.[80] Neither is the resolution authority's power to suspend certain payment obligations for one business day at the start of bail-in of any help here.[81][82] A run *during* a bail-in does not happen in our model since we assume the bail-in is completed fast, within one time step.[83] Below are the three scenarios we consider:

1. A creditor runs on a maturing bail-in debt contract B_{ji}^{kmt} whenever its expected loss Λ_{ji}^{kmt} exceeds a certain threshold ψ_j, that is if

$$\Lambda_{ji}^{kmt} := 1 - \frac{V_{ji}^{kmt}}{B_{ji}^{kmt}} > \psi_j, \tag{26}$$

   where the threshold ψ_j reflects investor j's risk tolerance. The first run criterion presupposes that a creditor can rely on the price of bail-in debt to inform discounted expected bail-in losses. (This will not be true if uncertainty about the applicable bail-in design prevails.) This enables a creditor to evaluate which of its bail-in debt contracts it should refinance. Since senior bail-in contracts are less prone to bear losses in a bail-in, senior contracts tend to revalue less than junior ones, and so are less vulnerable to run risk. The applicable conversion rates also matter for which debt contracts are expected to suffer the greatest net losses. With unfair conversion rates (which impose more losses on junior creditors than fair conversion rates do) junior contracts therefore tend to re-value downwards more sharply.

2. A creditor runs on a bank's maturing bail-in contracts whenever bank i's capital ratio ρ_i gets perilously close to its minimum capital ratio ρ_i^M, that is if

$$\rho_i < \rho_j^R := \rho_i^M + r_j, \tag{27}$$

   where the run threshold ρ_{ji}^R exceeds the failure threshold by a safety margin r_j reflecting

[80] The public funding backstop mechanism can be used if a (recapitalised) firm cannot maintain private sector access to refinance its liabilities as they fall due. The term of the funding is typically no longer than needed to maintain continuity of critical functions to achieve resolution, but sufficiently long to allow the G-SIB in resolution to regain access to private sources of funding (FSB (2016)). It is not clear whether D-SIBs and non-SIBs would also enjoy access to the public funding backstop mechanism. Hence, if they are also bailed in, a run on them during a bail-in could prove harmful.

[81] See: Article 69 of the BRRD.

[82] The suspension power does not apply to certain instruments, such as eligible deposits and payments to operating systems, presumably to ensure depositors retain access to cash and financial markets continue to operate.

[83] See a detailed explanation on the speed of bail-in in Sections 3.7 and 4.2.4

investor j's risk tolerance. The second run criteria acts as a crude rule of thumb for a refinancing decision in case the creditors cannot price bail-in debt due to uncertainty in the bail-in design.

3. A creditor will rein in its maturing exposure to a bank's bail-in debt if the bank looks similar to another bank that has just been bailed-in. We refer to these runs as "similarity runs". The set of banks $\mathcal{S}^t$ whose creditors will run out of similarity concerns at time t is given by

$$\mathcal{S}^t = \{i \in \mathcal{B} \setminus \mathcal{D}^t : \rho_i^t \leq \max(\rho_j^{t-2-u}, ..., \rho_j^{t-2}), \text{ for a } j \in \mathcal{D}^{t-1}\}, \tag{28}$$

where $\mathcal{D}^t$ is the set of bailed-in banks at time t. Equation 28 says that a bank i that is not being bailed-in at time t suffers a similarity run at time t if its capital ratio ρ_i^t is less than or equal to the capital ratio of any bank j in any of the u days running up to j's bail-in at time $t-1$. Similarity runs are more likely under profound uncertainty in the bail-in design, which inhibits a more accurate criterion to determine the refinancing of bail-in debt. Equation 28 proposes a simple measure of similarity. Many other, potentially more sophisticated ones, are of course conceivable. Two banks might, for instance, be perceived to be similar if the variance-covariance of their equity is alike.

Since non-bail-in debt is not subject to losses if a bank fails when the bail-in mechanism is applied, we assume non-bail-in debt is not subject to run risk. In practise, non-bail-in debt may be prone to some run risk, though presumably to a lesser extent than bail-in debt. One reason is that a bank that can not be recapitalised successfully with its available bail-in debt, may have to be liquidated if bail-out funds are not forthcoming. Creditors of the bank are then subject to a loss given default in a liquidation. Another reason is that investors may not know for certain that a failing SIB will pass the "public interest test". If it fails this test it will be liquidated via regular insolvency proceedings rather than resolved (via bail-in), in which case non-bail-in debt will also be at risk of losses. As noted in Section 3, we assume that the banks in our model, who are systemically important enough to be elected to participate in the 2018 EBA stress test, always satisfy the public interest test.

If a (subset of a) bank's maturing bail-in debt contracts are not rolled over then it will suffer a liquidity shock at time t_{x+1}^a.

Table 2 provides the default settings for model and figures in the result section. It also lists part of the sensitivity analysis we have done on each of the baseline settings.

Table 2: Default settings for model and figures in the result section.

Parameter Category	Default Settings	Detailed Default Settings	Brief Description & Motivation
Institutions	Banks turned on	Our research question focusses on stability in the banking sector.	
	Non-banks turned on	• Leveraged non-banks vs. non-leveraged non-banks: $\chi = 50\%$. • Initial leverage of leveraged non-banks $(\frac{E}{A})$: $\lambda_i^{t0} = 30\%$.	Turned on to take contagious feedback loops between the banks and non-banks into account. Moreover non-banks hold bail-in debt. We do not know what percentage of bail-in debt held by the non-banks is respectively held by leveraged vs. non-leveraged non-banks. We roughly know leverage of non-banks. (Sensitivity analysis available upon request.)
Contracts & Contagion Mechanisms	Bail-in induced exposure loss contagion, overlapping portfolio contagion & funding contagion turned on.	We include ('turn on') all relevant contagion channels, because modelling a subset of contagion channels may lead to an underestimation of systemic risk (see e.g. Kok and Montagna (2013), Caccioli et al. (2013), Farmer et al. (2020)). (Sensitivity analysis in Section 5.3 & Online Appendix D.2.)	
	Revaluation of bail-in debt turned on.	Jump process parameters: • Volatility $\sigma = 8\%$; • Jump volatility $\sigma_J = 2\%$; • Jump mean $\sigma_\mu = -2\%$; • Jump intensity $\lambda = 50$.	The jump process parameters are chosen in line with Chen et al. (2013), Pennacchi (2010), who model contingent convertibles (CoCos), the contractual analogue of a bail-in. (Sensitivity analysis in Online Appendix D.3.)
	Loss-concern-induced halts of rolling over bail-in debt ('bail-in runs')	Run scenario: • Runs based on expected losses (ψ_j=2.5%) turned off; • Runs based on 'uncertainty' ($r_j = 1\%$) turned on (consistent with poor structural design baseline); • Similarity runs ($u = 5$ days) turned off.	In the stress test excercise, we pose three 'what-if' scenarios for why creditors may decide to stop rolling-over bail-in debt. (Sensitivity analysis in results in Section 5.3 of paper & Online Appendix D.2.)
Constraints	Risk-weighted (rw) capital ratio	Regulatory stress tests focus on assessing whether the banks' rw capital ratios remain strong enough to survive an severely adverse scenario. We could also easily include the leverage ratio and liquidity coverage ratio, building forth upon Farmer et al. (2020).	
Market	Asset price fall is $x = 5\%$ if 5% of the market capitalisation has been sold.	This is in line with a standard assumption in the literature, see e.g. Schnabel and Shin (2004), Cifuentes et al. (2005) Gai and Kapadia (2010), Caccioli et al. (2014) and Farmer et al. (2020). (Sensitivity analysis in Farmer et al. (2020), on which this paper builds.)	
Behaviour	Seek to avoid default: Meet contractual obligations (CO) and regulatory constraints (RC).	• Fulfilling CO takes priority over complying with RC. • Pecking orders: liquidate most liquid assets first (for CO), liquidate assets with the highest risk-weight first (for the rw capital ratio). Internal rw capital buffer & target: • $\rho_i^B = \rho_i^M + (1-u_i)\rho_i^{CB}$, where $\rho_i^M = 4.5\%$ & $u_i = 50\%$; • $\tilde{\rho}_i^T = \min\{\rho_i^B + \Lambda_i\%, \rho_i^T\}$, where $\Lambda_i = 0.5\%$, $\rho_i^T = \rho_i + v_i\rho_i^{CB}$ & $v_i = 1$.	In line with Farmer et al. (2020) we assume banks are willing to use $u = 50\%$ of their combined regulatory buffer ρ_i^{CB}. Once they fall below this they will gradually seek to return to a stable rw capital ratio internal target ρ_i^T. (Sensitivity analysis in Farmer et al. (2020) on which this paper builds forth.)
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), Caccioli et al. (2014) and Farmer et al. (2020). (Sensitivity analysis in Farmer et al. (2020), on which this paper builds.)	
Failure Method & Design	Other bail-in parameter settings: • Equity valued at book value, $\sigma = B$; • Liquidation costs in hypothetical winddown of bank under normal insolvency procedures, $c_i^\tau = 20\%$ (i.e. recovery 80%); • Speedy bail-in completed over a resolution weekend, so bail-in is complete at time τ_b; • If bank does not have sufficient bail-in debt to be recapitalised to at least meet its minimum capital requirements it will be (disorderly) liquidated.	The hypothetical liquidation costs are in line with those estimated by Deloitte (2018), and the other settings correspond to the simplest choices the independent valuer and resolution authority could make. We show the impact of the bail-in design in absence of state aid.	