

Systemic Implications of the Bail-In Design

Alissa M. Kleinnijenhuis^{1,3,5}, Charles Goodhart², and J. Doyne Farmer^{3,4,5}

¹Stanford Institute for Economic Policy Research, Stanford University, US

²London School of Economics, UK

³Institute for New Economic Thinking, University of Oxford, UK*

⁴Mathematical Institute, University of Oxford, UK

⁵Oxford Man Institute of Quantitative Finance, University of Oxford, UK

August 29, 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 disruptively. Bail-in has been put forward as an alternative that potentially addresses the too-big-to-fail and contagion risk problems 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 all five of the 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 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 the participants at internal 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, 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 Rudy Tanin for his excellent research assistance. All errors remain ours.

Contents

1	Introduction	1
2	Data	4
3	The Bail-In Design under the Bank Recovery and Resolution Directive	5
3.1	Failure Threshold	7
3.2	Ad-Hoc Exclusions	8
3.3	Recapitalisation Target	8
3.4	Debt-to-Equity Conversion Rates	9
3.5	A-Priori Debt Exclusions	10
3.6	Loss Absorption Requirements	10
3.7	The Uncertainty and Speed of Bail-In	11
4	Model	11
4.1	Modelling the Systemic Implications of the Bail-In Design	11
4.2	Modelling the Design of the Bail-In Mechanism	14
4.2.1	Failure Threshold	15
4.2.2	<i>Ad-Hoc</i> Debt Exclusions	16
4.2.3	Recapitalisation Target	16
4.2.4	Debt-to-Equity Conversion Rates	17
4.2.5	<i>A-Priori</i> Debt Exclusions	19
4.2.6	Loss Absorption Requirements	20
4.2.7	Uncertainty in the Bail-In Design	20
5	Results	21
5.1	Stability Impact of the Bank-Specific Bail-In Design	23
5.1.1	The Stability Impact of the Failure Threshold	24
5.1.2	The Stability Impact of the Recapitalisation Target	27

5.1.3	The Stability Impact of the Debt-to-Equity Conversion Rates	29
5.2	The Stability Impact of the Structural Bail-In Design	30
5.2.1	The Stability Impact of the Loss Absorbing Requirements & Debt Exclusions from Bail-In	31
5.3	Contagious Amplifications of the Bail-in Design	33
5.4	Sensitivity Analysis	35
6	Policy Implications: The Bail-In Design and its Political Economy	37
6.1	Failure Threshold	37
6.2	Recapitalisation Target	39
6.3	Debt-to-Equity Conversion Rates	40
6.4	Debt Exclusions and Loss Absorption Requirements	41
6.5	The Uncertainty and Speed of Bail-In	42
7	Conclusion	43

1 Introduction

The failure of Lehman Brothers was perhaps the defining event of the 2007-2008 financial crisis, bringing the financial system and the real economy to the brink of the abyss (Bernanke (2017)). Incumbent governments were forced to choose between the unattractive alternatives of either bailing-out a systemically important bank (SIB) 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 SIB failures (Bernanke (2017)). The Squam Lake Report recommended that authorities be given the necessary powers to affect the orderly resolution of systemically important institutions (French et al. (2010)). This recommendation has been adopted. The Bank Recovery and Resolution Directive (BRRD) establishes a common approach to the recovery and resolution in the European Union of banks and investment firms. In the United States, Title II, the Orderly Liquidation provision of the Dodd-Frank Act, provides a mechanism for efficiently liquidating a large, complex financial institution.

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

The efficacy of bail-in in serious crisis scenarios remains an open question. While bail-in has proved relatively successful in dealing with the failure of relatively minor SIBs, its ability to handle large bank failures and system-wide crises has not yet been tested. Bail-ins on comparatively small SIBs have been carried out successfully in Cyprus, the Netherlands, Italy, Denmark, the United Kingdom and Spain. 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)), the political influence of powerful banks and their unwillingness to take losses distorted policy. 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 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 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 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 the 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 a multi-layered network model of the European financial system developed by [Farmer et al. \(2020\)](#). In addition, we calibrate this model using *S & P Global Market Intelligence* data and the 2018 European Banking Authority (EBA) disclosures of the stress test results. The calibrated system-wide stress testing model captures 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 counterparty 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 price 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 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. These key design parameters are typically present in each jurisdiction's bail-in framework.

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, “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 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. 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.¹ Shareholders or creditors for whom this does not hold are entitled to a difference payment by the resolution financing arrangements.

We further show that the systemic footprint of the bail-in design will not be properly understood without taking account of multiple contagion mechanisms, which can substantially amplify initial bail-in losses, and the role of non-banks, who hold much of the bail-in debt (in part because cross-holdings of bail-in debt among banks are discouraged by the eligibility rules for the loss absorption requirements). Our findings are qualitatively robust to an extensive sensitivity analysis.² 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. This paper helps answer the long-standing question of how we can solve the too-big-to-fail problem. Our contribution adds to literature that seeks to address this issue (see e.g. [Stern and Feldman \(2004\)](#), [Morrison \(2011\)](#), [Bernanke \(2016\)](#)). More narrowly it adds to the nascent literature on the systemic effects of bail-in in a networked financial system. [Klimek et al. \(2015\)](#) employ an agent-based 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 systemic instability particular bail-in designs can generate.

Our paper also contributes qualitatively and quantitatively to a conceptual literature that

¹See: Article 74(2) of the BRRD.

²For 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 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.

critically evaluates aspects of the bail-in design. This literature includes the papers of [Eichengreen and Rühl \(2001\)](#), [Rutledge et al. \(2012\)](#), [Conlon and Cotter \(2014\)](#), [Sommer \(2014\)](#), [Persaud \(2014\)](#) and [Avgouleas and Goodhart \(2015\)](#).

Our approach to modelling contagion falls squarely within a widely used approach in the contagion literature ([Aymanns et al. \(2018\)](#)). 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. In times of distress, constraints more often bind ([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. In line with [Gai et al. \(2011\)](#), [Caccioli et al. \(2013\)](#), [Kok and Montagna \(2013\)](#), [Greenwood et al. \(2015\)](#), [Duarte and Eisenbach \(2015\)](#), [Cont and Schaanning \(2017\)](#) and [Farmer et al. \(2020\)](#), who 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 \(1976\)](#)). We also add “what-if” scenarios under which creditors cease to refinance a bank’s bail-in debt. Unlike earlier quantitative papers on the topic which describe the bail-in process largely verbally (see e.g. [Klimek et al. \(2015\)](#), [Hüser et al. \(2017\)](#)), 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. Finally, though bail-in has been designed with systemic considerations in mind,³ it is not enough just to assert its suitability on a system-wide scale as [Aymanns et al. \(2016\)](#) demonstrated. By modeling explicitly its systemic impacts we thus fill an 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

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

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 bank balance sheets using *S&P Global Market Intelligence* data. These 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 the 48 banks that participated in the 2018 EBA stress test. 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.⁴

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 banks and non-banks as a data input. 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. 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 \(2010\)](#), [Caccioli et al. \(2014\)](#) and [Paulin et al. \(2018\)](#)), we average our results across realisations of the reconstructed networks to remove dependency on the probabilistic realisation. In Online Appendix B, we describe data to construct the bipartite 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. Next, we set out the framework for resolving failing banks that applies in Europe, 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. 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 the failing-likely-to-

⁴We 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 $\lambda^{t0} := \frac{\hat{E}_i^{t0}}{A_i^{t0}}$. In our robustness checks, we vary parameters χ and λ^{t0} to investigate how they affect banking sector stability.

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. The bank-structural design consists of *a-priori* debt exclusions, loss absorption requirements, uncertainty in the bank-specific bail-in design and the speed to complete a bail-in.

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 prevented otherwise,⁵ 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.⁶ The resolution authority then has to decide which 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, then the next step involves determining what set of debt should be excluded on an *ad-hoc* basis from bail-in (step 2). 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. 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 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

⁵See: Article 59(1)a of the BRRD.

⁶See: Article 37(3)(4) of the BRRD.

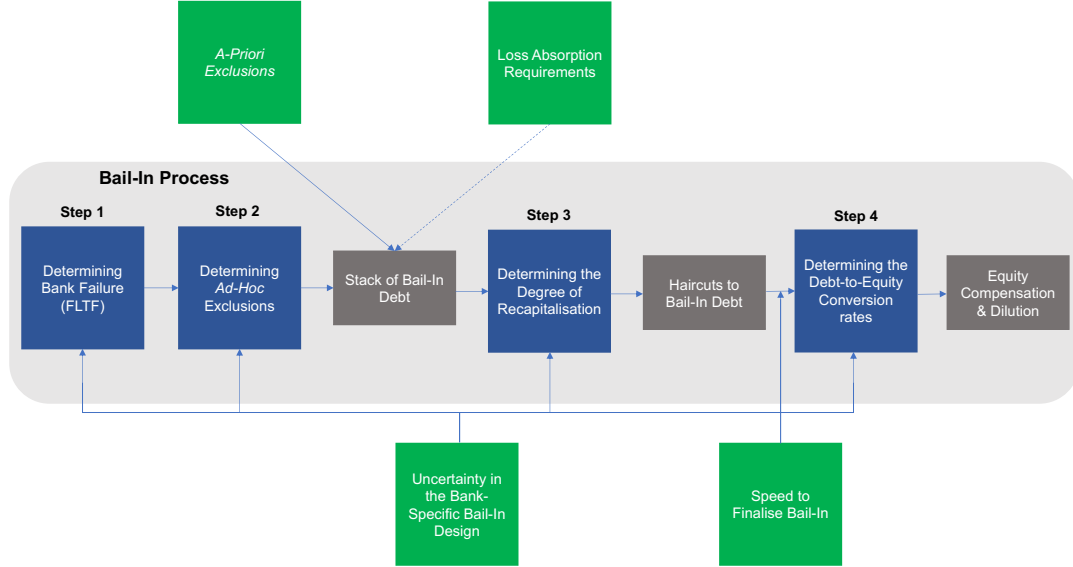


Figure 1: **Schematic overview of the bail-in mechanism.** The blue blocks denote the parameters of the bank-specific bail-in design and the green blocks the parameters of the structural bail-in design.

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 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.⁷

We will now further delineate the bail-in design under the BRRD. A bank will be resolved if three conditions are fulfilled.

3.1 Failure Threshold

A bank will be bailed-in if it is deemed to be *failing or likely to fail* (see 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);

⁷See: Article 52 of the BRRD.

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)).⁸ 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-in tool. The resolution authority may decide to apply *ad-hoc* debt exclusions if at least one of the following conditions is met:⁹

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.¹⁰ This arguably favors some creditors unfairly.

3.3 Recapitalisation Target

The recapitalisation target (see 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

⁸Valuation 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.

⁹See: Article 44(2) of the BRRD.

¹⁰Once 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 RFF bail-out amount is capped at a certain maximum, 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.

with the conditions for authorisation and sustains or regains market confidence.¹¹ 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 (Burrows et al. (2012)).

3.4 Debt-to-Equity Conversion Rates

We now discuss the debt-to-equity conversion rates (step 4 in Figure 1). 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 winding up the bank under normal insolvency proceedings.¹² 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).¹³
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 2.¹⁴ Valuation 3 is deployed by the independent valuer to evaluate the NCWO criteria (EBA (2017b)).¹⁵ We now proceed to discuss the parameters of the structural bail-in design, starting with *a-priori* debt-exclusions (see green box in Figure 1).

¹¹See: Article 43(2)a of the BRRD.

¹²See: Article 74(2) of the BRRD.

¹³See: Article 75 of the BRRD.

¹⁴See: Article 47(1)a of the BRRD.

¹⁵Valuation 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 and the market circumstances at the resolution date (EBA (2017b)).

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:¹⁶

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 are required to meet loss absorbing requirements (see 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),¹⁷ 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.¹⁸

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.

¹⁶See full list in Article 44(2) of the BRRD.

¹⁷See: Article 45 of the BRRD.

¹⁸Nevertheless, 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.

3.7 The Uncertainty and Speed of Bail-In

Bail-in debt can be priced only if the bank-specific design parameters (the four blue boxes in Figure 1) are clearly established in advance by the resolution authority (green box in Figure 1). If the resolution authority has not confirmed 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 debt and thus risk cannot be properly priced. The resolution authority has significant discretion in choosing the appropriate value of the bank-specific bail-in parameters, as the previous sections made clear. Hence, uncertainty in the bank-specific bail-in design currently prevails.

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.¹⁹

4 Model

In this Section, we present our model of the bail-in design’s components and its systemic implications. In the first subsection, we explain how we model the financial stability implications of the bail-in design. We model this using a multi-layered network representation of the financial system on which we conduct a system-wide stress test. This part spells out how bail-ins could induce contagion in the financial system and vice versa, and uses the model of the bail-in mechanism described in the second subsection of this Section. A summary of the notation used to specify the model is given in Table 2 and 3 of Online Appendix A.

4.1 Modelling the Systemic Implications of the Bail-In Design

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

¹⁹The 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.

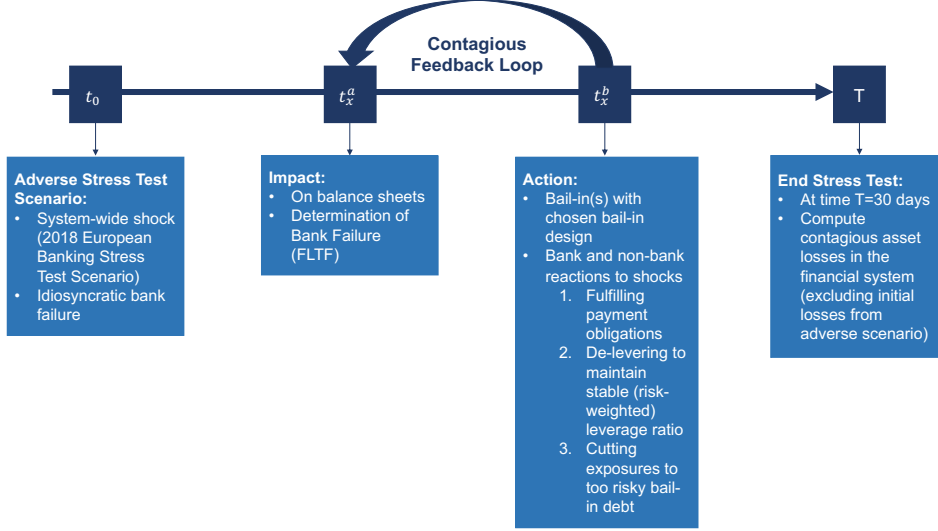


Figure 2: Time line of the system-wide stress test with bank bail-ins. The stress test starts at time t_0 with an exogenous crisis scenario and ends 30 days later at time $T = 30$. In the time steps t_x between contagion may take place, for $x = 1, \dots, 29$. At any sub time step t_x^a the impact of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

We will now walk you through the time steps of the stress test using Figure 2 at a high-level.²⁰ 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 some of the banks' and non-banks' balance sheets 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 the next sub step b of time t_1 using the applicable 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, as will be shown in the results in Section 5. 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 $t_{x=2}^a$ give rise to further 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 . Each of the possible reactions – action 1 (fulfilling payment obligations), action 2 (de-levering) and action 3 (cutting 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 , potentially resulting in more bail-ins and pre-cautionary actions at time t_2^b . All this can bring about another contagious feedback loop. Such a feedback loop might be exacerbated by downwards re-valuations of banks' bail-in debt, which result from asset losses that push up its probability of default. These feedback loops may continue for $x = 1, \dots, 30$ time steps. The system-wide stress test stops after $T = 30$

²⁰A more detailed account of the stress test model and its default settings (Table 4) are given in Online Appendix C.2.

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.

Let’s now discuss in more detail our motivation for modelling reactions to shocks at time t_x^2 in Figure 2. As noted before, we assume, in line with the literature, that the behaviour of institutions in times of crisis is driven by the desire to avoid failure. So in our model institutions act in response only to those shocks (felt at time t_x^a) that cause them to either nearly fail (would they not react) or fail completely. We proceed to explain technically how we model action 1, 2 and 3.

Action 1 (fulfilling payment obligations) & 2 (maintaining a stable capital ratio)

In practise, a bank may be prompted to de-lever (action 2) in response to asset losses in order to maintain a stable risk-weighted capital ratio ρ_i and stay away from their failure threshold, as well as avoid regulatory penalties (applicable if they use their regulatory buffers). Non-banks, likewise, may be forced to 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_i^T whenever their losses eat into over $u = 50\%$ of their Basel III “combined capital buffers”, ρ_i^{CB} , in line with [Farmer et al. \(2020\)](#). Non-banks in our model de-lever to their target λ_i^T whenever their leverage ratio falls below 90% of their initial leverage ratio $\lambda_i^{t_0}$, in line with [Cont and Schaanning \(2017\)](#).²¹ The assumption of leverage targeting is widely used in the contagion literature (see e.g. [Kok and Montagna \(2013\)](#), [Greenwood et al. \(2015\)](#), [Duarte and Eisenbach \(2015\)](#), [Cont and Schaanning \(2017\)](#)) 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. We assume that a bank returns to its internal target gradually: by no more than $\Lambda_i\%$ per time step. They thereby limit self-inflicted harm induced by high liquidation costs and contagious feedbacks ([Aymanns et al. \(2016\)](#) and [Caccioli et al. \(2014\)](#)).

The order in which assets are liquidated to de-lever is dictated by a “pecking order”, in line with [Halaj \(2018\)](#). This stipulates that the assets with the highest risk weights will be reduced first in order to most efficiently de-lever in risk-weighted terms.²² It further stipulates that the assets that are least costly to reduce will be reduced first if the institution wants to de-lever its non-risk weighted leverage ratio or wants to raise cash to meet obligations (action 1), in line with [Cifuentes et al. \(2005\)](#). Practically, this means that an institution will pull back maturing funding before liquidating securities, as the latter may have a price impact but the former has not.²³ 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 [Greenwood et al. \(2015\)](#), [Cont and Schaanning \(2017\)](#) and [Farmer et al. \(2020\)](#). In line with a common approach in the literature (see e.g. [Gai et al. \(2011\)](#)), we assume maturing funding is rolled over automatically, unless institutions decide explicitly to stop rolling over bail-in funds.

²¹Our findings are qualitatively robust to the choice of buffer threshold at which an institution starts to de-lever to return to its target capital ratio.

²²We conducted sensitivity analysis and found that choosing a different pecking order does not qualitatively change our results.

²³We 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.

Action 3 (reducing risky exposures)

Losses to a bank's assets at time t_x^a will result in a downwards re-valuation of its bail-in debt. Bail-in debt is valued as the expected discounted value of future payoffs, in line with [Black and Scholes \(1973\)](#) and [Merton \(1974\)](#). Online Appendix C.2.1 describes our novel method for pricing bail-in debt. Under our valuation method, bail-in debt becomes less valuable if the 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 certain 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. Creditors in our model will stop rolling over bail-in debt whose value has fallen sharply (action 3), since this indicates that the expected losses on holding this debt might be large. Section 4.2.7 motivates and defines the non-roll-over scenarios associated with action 3 in more detail. Each of these 3 actions could prompt further contagious feedback loops.

Interacting Contagion Mechanisms

So far, we have explained that an adverse scenario can result in bail-ins and contagious asset losses which may prompt further contagious losses and bail-ins. We have not yet discussed that the contagion mechanisms at play may amplify each other resulting in increased contagious losses, as demonstrated by [Kok and Montagna \(2013\)](#), [Caccioli et al. \(2013\)](#) and [Farmer et al. \(2020\)](#). Therefore, it is key to model 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 made.

4.2 Modelling the Design of the Bail-In Mechanism

In this section, we 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).

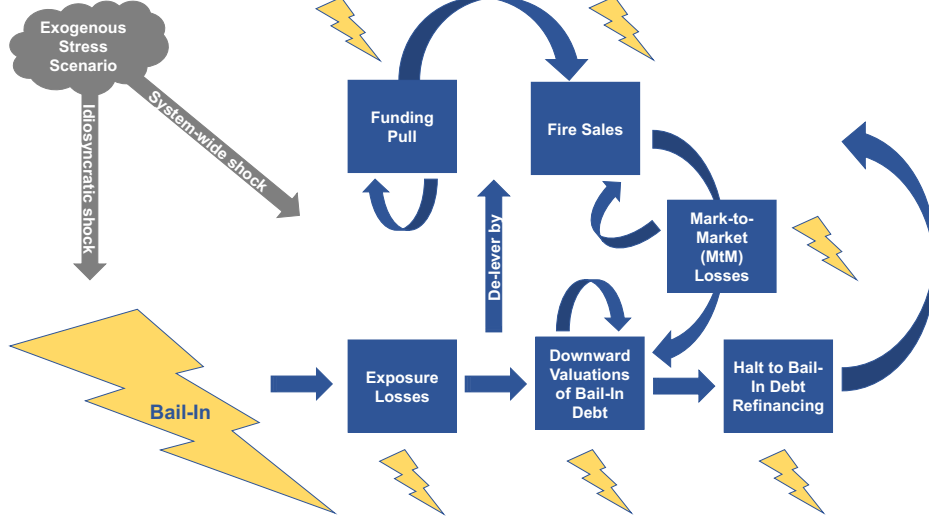


Figure 3: **Interaction among Contagion Mechanisms.** A bail-in (shown with a lightning bolt) may precipitate interactions between the five contagion mechanisms: exposure losses, revaluations of bail-in debt, halts to roll-overs of bail-in debt, funding contagion and firesale contagion. These forms of 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.

4.2.1 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).²⁴ 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 this assumption, a bank i is bailed-in at time τ_i given by

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

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 ²⁵ 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. 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.

²⁴We assume that the other resolution objectives, including that the resolution of the bank is in the public interest, are always true for a SIB, so do not evaluate these explicitly.

²⁵In 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 ?? in Online Appendix B.

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

$$B_i^{k\tau_e} = B_i^{k\tau} - U_i^{k\tau_e}, \quad (2)$$

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}

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

where $\mathcal{K} := \{2, \dots, 6\}$ 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 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.

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 having to change the model notation.

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-weighted assets Ω_i^τ , the total haircuts that must be imposed to recapitalise a bank to target is given by

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

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 is set higher ρ_i^{RT} . 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^τ 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, 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^τ : $h_i^\tau = \min\{\hat{h}_i^\tau, B_i^{\tau_e}\}$. If the bank cannot be recapitalised enough to meet its CET1 minimum capital requirements $\rho_i^M = 4.5\%$ applicable in Pillar I, then we assume it will be liquidated via the regular insolvency procedure, as is detailed in Online Appendix C.3.

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 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 proscribed in insolvency. The haircuts in each priority class $k \in \mathcal{K}$ will then be 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.²⁶ 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 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 phase Δ_{ia}^k and that applicable in the recapitalisation phase Δ_{ib}^k . Their relationship is given by

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

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

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

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.²⁷ We assume that creditors are compensated with equity at time τ_b immediately after haircuts have been imposed.²⁸

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

²⁶While bail-in debt must be treated equitably within a priority class, bail-in and non-bail-in debt does 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.

²⁷To 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 24 and 29 in Online Appendix C.1.1.

²⁸By 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.

by

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

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.²⁹

Conversion Rate Principles

The set of conversion rates must adhere to the NCWO and the preservation-of-hierarchy-of-claims principles. Therefore, conversion rates must 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 , as 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 or not 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. To adhere to the second principle, the conversion 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 55 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, are superficially fairer than "fair" conversion rates, they typically do not adhere to the conversion rate principles. "Zero" conversion rates or "pure write-downs", in which none of the bailed-in creditors receives compensation for its haircuts,

²⁹As 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.

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).³⁰ 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 in a hypothetical liquidation of the 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 13 in this Online Appendix illustrates how fair conversion rates work and Figure 14 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

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 .³¹ On the other hand, existing shareholders will be diluted if the bank is found to still be solvent at the start of bail-in.³² 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 4.³³³⁴

³⁰Conversion rates are frequently made unfair whenever the resolution authority applies *ad-hoc* exclusions.

³¹See: Article 47(1)a of the BRRD.

³²See: Article 47(1)b of the BRRD.

³³In 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 simplification 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.

³⁴A 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

B_i, Bail-In Debt	
$B_i^{k_6}$, Deposits (in excess of the Deposit Guarantee Scheme) (Excluding those with a TTM less than 7 days to institutions)	k_6
$B_i^{k_5}$, Interbank and Other Liabilities (Excluding those with a TTM less than 7 days to institutions)	k_5
$B_i^{k_4}$, Interbank and Other Liabilities (Excluding those with a TTM less than 7 days to institutions)	k_4
$B_i^{k_3}$, T2 Capital	k_3
$B_i^{k_2}$, AT1 Capital	k_2
$B_i^{k_1}$, CET1 Capital	k_1

Figure 4: Break-down of bail-in debt B_i^k per priority class $k = k_1, \dots, k_6$. 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.

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 $P_i^{k\tau}$:

$$B_i^{k\tau} := L_i^{k\tau} - P_i^{k\tau}, \quad (8)$$

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.

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

When there is uncertainty about the bail-in design, creditors cannot properly price the bail-in debt. Thus they cannot make roll-over decisions on bail-in debt based on its expected loss, as creditors do in our model when there is sufficient design certainty. We model the effect of uncertainty using two possible heuristics. One is that investors roll-over bail-in debt based on the distance-to-default, and the other is that they do it based on the similarity to a bank that has recently failed. The formulaic definition of these three criteria (which we refer to as expected, uncertainty and similarity runs) is given in Online Appendix C.2.2. We flag experiments with

paper, we do not include this debt class in this paper.

design uncertainty with $\nu = 1$ and those with sufficient design certainty with $\nu = 0$.

We now discuss how we model the financial stability implications of a bail-in design. The bail-in design of a bank in our model is given by $\mathcal{D}_i = (\rho_i^F, \mathbf{U}_i^{\tau_e}, \rho_i^{RT}, \Delta_i, \mathbf{P}_i, L_i^M, \nu)$. See table 1.

	Bail-in design parameter	Description
Bank-specific bail-in design	ρ_i^F	Failure threshold
	$\mathbf{U}_i^{\tau_e}$	Set of <i>ad-hoc</i> exclusions
	ρ_i^{RT}	Recapitalisation target
	Δ_i	Set of debt-to-equity conversion rates
Structural bail-in design	\mathbf{P}_i	Set of <i>a-priori</i> exclusions
	L_i^M	Minimum loss absorption requirement
	ν	Uncertainty in the bank-specific bail-in design

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

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 a shock to their risk-weighted capital so these drop $q = 4\%$ below the failure (F) threshold ρ_i^F .³⁵ The results are qualitatively robust to this choice.³⁶ If the threshold ρ_i^F is set at 0%, for instance, this means that 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.

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 shock.³⁷ The banks' average risk-weighted capital ratio before the system-wide shock takes place is 15% (with a standard deviation of 3.6%). For $s = 1$, after the adverse shock the average is 11.5%

³⁵We will show results for up to $n = 5$ idiosyncratic bank failures to investigate how systemic losses grow as the number of systemic bank failures goes up. In practise, we think it is unlikely that more than 3 idiosyncratic bank failures happen at a time.

³⁶A larger q simply means that more losses have to be absorbed with haircuts to recapitalise a bank to its recapitalisation target, resulting in a qualitatively bigger shock to the system.

³⁷Contrary to the EBA scenario, for simplicity we assume asset losses hit banks' external assets only, resulting in a diminished capital ratio, whereas in the EBA scenario all of the banks' assets suffer losses as a consequence of the adversely set 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 a risk-weighted capital ratio of each bank at the start of stress test equal to $\rho_i^{t_0} = \rho_i^{data} - s(\rho_i^{data} - \rho_i^{EBA, adverse})$. $\rho_i^{t_0}$ 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.

(with a standard deviation of 3.5%). Thus the original EBA shock corresponds to a loss of about 23% in the risk-weighted capital ratio relative to the pre-shock ratio.

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 “bad” 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 “bad” 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 also shows the results using the bad 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 bad bail-in design, and the other default settings used to generate the results are listed in Table 4 in Online Appendix A. The key parameters for the bail-in design are given in Figure 5. 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.

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 that our results are quantitatively precise. 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.

	Bail-In Design Parameters	Good Bail-In Design	Bad Bail-In Design
Bank-Specific Bail-In Design	Failing-Likely-To-Fail Threshold	Early Bail-In p^F equal to MCR	Late Bail-In $p^F = 0\%$ (equal to insolvency point)
	Recapitalisation Target	Strong Recapitalisation p^{RT} equal to MCR plus twice CBR	Weak Recapitalisation p^{RT} equal to MCR plus half CBR
	Debt-to-Equity Conversion Rates	Fair Conversion Rates See equation 39 & 40 in Appendix B.2.2	Unfair Conversion Rates See equation 46 in Appendix B.2.2
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 <i>Status Quo</i> to Counteract A-Priori Debt Exclusions i.e. debt exclusions with $7d < TTM < 1Y$	Current MREL and TLAC minimum loss absorption requirements See section 4.2.6
	Uncertainty in the Bank-Specific Bail-In Design	No/Limited Bail-In Debt can be priced; Roll-over decision based on expected loss; Run condition 1 in Appendix B.2.3	Yes Bail-In Debt cannot be properly priced; Roll-over decision based on rule of thumb; Run condition 2 in Appendix B.2.3

Figure 5: Shows the “good” and “bad” baseline configurations of the bail-in design. By default we show results for the good bank-specific design (highlighted in green) and the bad 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.

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.

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




Figure 6: The systemic response for idiosyncratic shocks and for the EBA shock. The blue line uses good and the orange line uses the bad bank-specific parameters. Both lines use the bad structural parameters. See Table 5 for the good and bad parameters.

Figure 6 shows how the financial system responds to shocks differently for a bank-specific bail-in design with the good and bad bank-specific parameters, keeping the structural parameters equal to their default (bad) 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 systemic loss to the banking system is measured in trillions of euro. The asset loss shown is only the contagious loss, i.e we do not include the exogenously imposed shocks in either case. We observe that when both the bad bank-specific and bad 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 22 trillion euro, so 6 trillion euros represents 27.3% 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 1.3% 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).³⁸ On the other hand, we observe that “bad” 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 19 in Online Appendix D shows. This is in line with our experience of bailing-in small SIBs.

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 “bad” 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 bad 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 The Stability Impact of the Failure Threshold




Figure 7: Systemic response varying the failure threshold. All other bank-specific parameters are set to their good values. As before, structural parameters are set to bad 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 the system-wide shock becomes sufficiently severe. Before that our analysis shows that the

³⁸We 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.

shock is not severe enough to trigger bail-ins, precautionary delevering or exposure reductions to 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, because of a phase transition.³⁹ 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 20 in Online Appendix D 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

³⁹We confirmed that asset losses keep slowly increasing 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).

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 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.⁴⁰

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 is 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 is 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.

⁴⁰Equation 55 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 The 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 bad 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 bad 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 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? Part of the answer is that a stronger recapitalisation means that a bank finds it easier to retain access to funding in our model (recall Section 4.2.7) and keep its funding cost low (latter not modelled). This part is well-understood in the literature (see e.g. Burrows et al. (2012)). Further, a stronger recapitalisation means that the bank in our model has less 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. 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 conversion rates, creditors in the recapitalisation phase of the bail-in can be compensated fully, resulting in net losses zero 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 The 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 have born the brunt of the losses. 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 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 The Stability Impact of the Structural Bail-In Design




Figure 10: The stability impact of altering the structural bail-in design relative to its *status quo* (i.e. bad baseline value). The blue (orange) line uses both the good (bad) 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 bad structural parameters unchanged, we now set the structural bank-specific parameters to their good value in the blue line, and leave it to their bad value in the orange line as before. 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 bad 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 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 bad 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.

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 bad 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 bad structural baseline, as per usual.) In the green line, the loss absorption requirements are increased relative to the bad 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. 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 Figure 22 in Online Appendix D shows. In Figure 22, which instead uses the bad bank-specific baseline, the differences between the three lines grow. The reason is that stability differences between structural designs are amplified by bad bank-specific parameters, since losses generated by multiple bad parameters amplify each other. Interestingly, how much a bail-in parameter matters thus depends on how the other bail-in parameters are set.

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.⁴¹ 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 vulnerable 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.⁴² 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.⁴³, 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.⁴⁴

On the other hand, short-term bail-in debt risks not only being reduced due to runs if it remains included, it is also reduced if short-term debt is decided to be excluded. Excluding short-term bail-in debt exerts a negative pressure on stability for same reason that a bail-in collapse is bad: it renders its successful recapitalisation less likely.

To alleviate the downsides 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 TTM 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 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.

⁴¹Recall 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.

⁴²Recall that in our model asset liquidations involve fire sales of securities and the retraction of maturing funding to counterparties.

⁴³Online Appendix C.3 explains how we model a liquidation.

⁴⁴Another 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.

We leave the inclusion of this channel for future research.

A simple way to gauge the trade-off is to measure the return on investment in terms of lowering contagious losses of increasing the long-term debt by x amount. Put differently, one can measure by how much long-term debt across the system will have to increase for contagious losses to drop from the green to the orange line. When the good bank-specific baseline is used as in Figure 11 the return on investment (drop in losses/increase in long-term debt) is approximately a meagre 0.2. When the bad bank-specific baseline is used as in Figure 22 in Online Appendix D the return is still a paltry, but slightly higher, 0.7 showing that the return on investment of increasing the loss absorption requirements depends on the chosen values of the other bail-in parameters. This simple exercise suggests that increasing the loss absorption requirements might be costly compared to the stability gains. A deeper study is required to confirm this. Next, we study how the contagion mechanisms at play amplify the systemic impact of bail-in designs.

5.3 Contagious Amplifications of the Bail-in Design




Figure 12: Contagious amplifications of the bad 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).

Figure 12 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 their bail-in debt losses. 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 bail-in debt can be priced. 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 priced. 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 Section 4.2.7 for a formulaic definition of these three scenarios.)

The results in the paper correspond to the baseline settings where all contagion mechanisms shown in the purple line are turned on, as Table 4 in Online Appendix A showed. “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 12 corresponds to the bad bank-specific and structural parameters in order to highlight how much bad bank-specific parameters amplify the interaction among contagion mechanisms, rather than for the good bank specific parameters as in the rest of the paper. We discuss how a different baseline affects contagious amplifications below.

The first thing we observe from Figure 12 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 EU financial system remains resilient in the face of severe system-wide shocks.⁴⁵ We observe that if instead, rightly, four more prevailing contagion mechanisms are taken into account (i.e. overlapping portfolio, funding, revaluation and various forms of run contagion) 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 (Kok and Montagna (2013), Caccioli et al. (2013), Wiersema et al. (2019), Farmer et al. (2020)). 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 12 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 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 its 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

⁴⁵The 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.

(see Figure 24 in Online Appendix D) and an ill-designed one (as in Figure 12 above). 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 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 cannot be priced.

Finally, we observe from Figure 12 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).⁴⁶ 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 system and the contagious shocks that may emerge.

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 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 “bad” 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.

⁴⁶With 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.

- (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 bad 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 bad bank-specific to a good bank-specific bail-in design, as well as from a bad 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 showed that our findings on the systemic implications of the bail-in design hold for different “what-if” scenarios (i.e. assumptions) regarding when creditors stop rolling over bail-in debt.
3. 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, but the qualitative impact of the choice of bail-in design remains unaltered.⁴⁷
4. We have conducted a host of other sensitivity analyses on model of the contagion dynamics in the financial system, which are summarised in Table 4 in Online 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. We have also asserted that price formation parameters (such as the price 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.

⁴⁷As 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. These results are not shown in the paper but can be provided upon request.

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 EU and US regulators are prone to bail-in a bank late. US regulators explicitly plan to bail-in a bank late, only when it is or nearly is insolvent ([McAndrews et al. \(2014\)](#)). EU 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. As noted before, one reason is that suffered losses are often recognised belatedly. Another is that a bank's asset value may fall sharply due to a sudden asset loss in a crisis. This suggests setting the FLTF threshold sufficiently high such that, even if unrecognised losses are discovered and asset value uncertainty prevails, a bank will typically be found to be solvent at the onset of bail-in. A final important reason for tending to bail-in a bank late is that it is in the short-term interest of bank's stakeholders ([Avgouleas and Goodhart \(2015\)](#)). The basic problem is that the incentive structure for almost all those closely involved leads them in the opposite direction from bailing-in early. Complementing the regulatory failure trigger with a market-based one may combat the tendency to bail-in too late. An alternative is a metric based on multiple requirements, see [Buckmann et al. \(2021\)](#). 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.⁴⁸ And it is this valuation that typically leads the regulatory authority to claim, at a late stage, that the failure threshold is triggered.

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

⁴⁸See: Article 36(1) and 36(4)a of the BRRD and [EBA \(2017b\)](#).

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 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 be 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 equated 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,⁴⁹ 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.⁵⁰ 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

⁴⁹In 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.

⁵⁰The 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.

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 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.⁵¹ For all the reasons set out in footnote 51, our tentative position is that, as a generality, it would be better to complete the process with Valuation 3 quickly rather than slowly; but it is a fairly 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 overstretched 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 the 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 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

⁵¹First, 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.

should alleviate concerns that bail-ins are unfit for dealing with failures in systemic crisis. And thereby instill regulators with the confidence to opt for well-designed bail-ins rather than bail-outs in the next systemic 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 parameters of bank-specific bail-in design. On the other hand, we found that a crisis-prone bail-in design consists of the opposite parameters. Our qualitative findings on the directional effect of bail-in design parameters on financial stability are 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.

Our evidence suggests that a possible pivot towards stability remains in the hands of policy-makers, 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 for 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 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 not just qualitative but also 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 compare EU and US); (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-ins are a credible alternative to bail-out. Our paper shows that *the credibility of bail-in critically depends on the bail-in design*.

References

- 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.
- Adrian, T. and Shin, H. S. (2010), ‘Liquidity and leverage’, Journal of financial intermediation **19**(3), 418–437.
- 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.
- Bernanke, B. (2016), ‘Ending “too big to fail”: What’s the right approach’, Brookings Institution’s Blog .
- 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’.
- Black, F. and Scholes, M. (1973), ‘The pricing of options and corporate liabilities’, Journal of Political Economy **81**(3), 637–654.
- BoE (2017), The Bank of England’s approach to resolution, Technical report, Bank of England.
- 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’.
- Burrows, O., Learmonth, D., McKeown, J. and Williams, R. (2012), ‘RAMSI: a top-down stress-testing model developed at the Bank of England’.
- 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.
- 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.
- Conlon, T. and Cotter, J. (2014), ‘Anatomy of a bail-in’, Journal of Financial Stability **15**, 257–263.
- Cont, R. and Schaanning, E. F. (2017), ‘Fire sales, indirect contagion and systemic stress testing’, Norges Bank Working Paper, 2017/2 .
- Duarte, F. and Eisenbach, T. M. (2015), ‘Fire-sale spillovers and systemic risk’, FRB of New York Staff Report 645 .
- 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.
- Eichengreen, B. and Ruehl, C. (2001), ‘The bail-in problem: systematic goals, ad hoc means’, Economic Systems **25**(1), 3–32.
- Farmer, J. D., Kleinnijenhuis, A. M., Nahai-Williamson, P. and Wetzter, T. (2020), ‘Foundations of system-wide stress testing with heterogeneous institutions’, Bank of England Staff Working Paper No. 861 .
- 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>.
- 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.
- 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.

- 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.
- 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’.
- He, Z. and Krishnamurthy, A. (2019), ‘A macroeconomic framework for quantifying systemic risk’, American Economic Journal: Macroeconomics **11**(4), 1–37.
- Hüser, A.-C., Halaj, G., Kok, C., Perales, C. and van der Kraaij, A. (2017), ‘The systemic implications of bail-in: a multi-layered network approach’, Journal of Financial Stability **38**, 81–97.
- Kleinnijenhuis, A. M., Kodres, L. E. and Wetzter, 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.
- 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.
- 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>.
- Paulin, J., Calinescu, A. and Wooldridge, M. (2018), ‘Understanding flash crash contagion and systemic risk: A micro-macro agent-based approach’, arXiv preprint arXiv:1805.08454 .

- Pennacchi, G. (2010), ‘A structural model of contingent bank capital’, FRB of Cleveland Working Paper 10-04 .
- Persaud, A. (2014), ‘Why bail-in securities are fool’s gold’, Available at SSRN 2643702 .
- 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’.
- Segoviano, M. B. and Goodhart, C. (2009), ‘Banking stability measures’, IMF working papers pp. 1–54.
- Sommer, J. H. (2014), ‘Why bail-in? and how!’, Economic Policy Review .
- Stern, G. H. and Feldman, R. J. (2004), Too big to fail: The hazards of bank bailouts, Brookings Institution Press.
- WBG (2017), Understanding bank recovery and resolution in the EU - A guidebook to the BRRD, Technical report, World Bank Group.
- Wiersema, G., Kleinnijenhuis, A. M., Wetzter, T. and Farmer, J. D. (2019), ‘Scenario-free analysis of financial stability with interacting contagion channels’.