

Macroprudential Regulation: A Risk Management Approach*

Daniel Dimitrov[†] Sweder van Wijnbergen[‡]

First version: 8 Feb 2023; Current version: May 11, 2026

Abstract

We develop a credit-risk based framework for macroprudential policy that treats systemic risk as the credit exposure of a policymaker to a portfolio of supervised banks. Using a structural credit model estimated from Credit Default Swap (CDS) prices, we derive the socially optimal capital buffers per bank by balancing the benefits of lower systemic risk against the economic costs of higher capital requirements through reduced credit supply. Applying this framework to Europe's systemic banks, we find that the market-based optimal buffers are substantially higher than those currently in use and vary significantly across institutions. Conditional on current country averages, within-country buffers are aligned close to optimally; suboptimality mostly arises from substantial differences between country averages.

JEL codes: G01, G20, G18, G38

Keywords: systemic risk, capital buffers, financial institutions, regulation, CDS rates

*We thank Andreas Beyer (ECB), Thomas van den Berg (DNB), Diana Bonfim (Banco de Portugal), Paolo Bonomolo (DNB), Maurice Bun (DNB), Pascal Busch (ECB), Klaus Düllmann (ECB), Fabian Greimel (UvA), Jeroen Huiting (DNB), Franc Klaassen (UvA), Philipp König (Bundesbank), Kenny Martens (DNB), David Martinez-Miera (Universidad Carlos III), Dimitris Mokas (EBA), Dino Palazzo (Fed), Saskia de Vries (DNB); and seminar participants at the University of Amsterdam, De Nederlandsche Bank, ECB Bank Supervision Research, Bank of Finland, IRLC, IFABS, ESCB Research Cluster on Financial Stability, MPPG, RiskLab, ESRB Conference on AI and Systemic Risk Analytics, for the helpful comments and discussions. We thank Meilina Hoogland (DNB) for data updates and acknowledge financial support from DNB. All views expressed are ours and not necessarily supported by DNB.

[†]University of Amsterdam, De Nederlandsche Bank; D.K.dimitrov@uva.nl

[‡]University of Amsterdam, De Nederlandsche Bank, CEPR and Tinbergen Institute; S.J.G.vanwijnbergen@uva.nl.

1 Introduction

Regulators worldwide use macroprudential capital requirements, set as add-on buffers to the standard microprudential capital, to manage the risks of a systemic crisis. Increasing the loss-absorbing capacity of large, economically important, interconnected banks reduces the chances of their default in adverse circumstances, and thus curtails the possibility that they can trigger cascading distress of related institutions.

However, in contrast to the well-established methodologies which determine microprudential requirements based on banks' own risk, there exists little theoretical guidance on how to set add-on buffers tailored to the way individual banks interact with the financial system as a whole. Recent academic approaches to measuring individual banks' contributions to systemic risk have emerged that could resolve this gap (Adrian and Brunnermeier, 2016; Acharya et al., 2017; Brownlees and Engle, 2017). Their structural relation to capital requirements, however, is not yet clear. In addition, the use of equity returns in estimating these measures precludes the inclusion of systemic non-listed banks.

In this paper, we address this gap by providing a structural credit model behind the measures of a bank's systemic risk contribution. Then, we derive the socially optimal average buffer add-ons and the optimal mapping of individual banks' contributions to systemic risk into bank-specific macroprudential buffers. This is done by evaluating the trade-off between the social costs of higher buffers (reduced lending) against the social benefits (reduced expected systemic losses in a systemic crisis), as well as by determining the optimal allocation of capital buffer add-ons across the most systemic banks.

In our empirical application, we find that the welfare-optimizing level of total bank capitalization is substantially higher than what is currently observed in practice. In addition, when we recalibrate the optimization to match existing regulatory buffer averages, we find that at the European scale allocations diverge meaningfully from the model's optimal distribution, whereas under national calibration we do not reject alignment with optimality, consistent with domestic authorities allocating capital in line with country-specific systemic risk profiles. Taken together, these findings suggest that while national practices broadly reflect local systemic conditions, a harmonized and model-driven allocation at the European level would help internalize cross-border dependencies and enhance the

overall coherence and effectiveness of macroprudential capital policy.

Our work bridges three distinct strands of the literature. First, we contribute to the systemic risk measurement literature, which uses market-based indicators to identify systemic banks (Huang et al., 2012; Adrian and Brunnermeier, 2016; Acharya et al., 2017). While influential, this strand offers little guidance on the size of optimal capital buffers. We extend these approaches by showing how additional model structure can link systemic risk to buffer calibration.

Second, we build on the macro-level cost–benefit literature on optimal capital requirements (BCBS, 2010; Miles et al., 2013; Brooke et al., 2015; Cline, 2017). This work seeks to maximize net social welfare from higher steady-state capital ratios by weighing the benefits of fewer and milder systemic crises against the costs of higher lending spreads. This literature, however, has no way of taking into account bank heterogeneity in modeling bank capital requirements.

Finally, we connect to the policy-oriented literature on systemic risk buffer calibration (FRB, 2015; Passmore and von Hafften, 2019; Jiron et al., 2021; Geiger et al., 2022), which emphasizes distributional rules such as the Equal Expected Impact approach—equalizing the marginal systemic impact of large banks relative to non-systemic ones.

Our contribution is to integrate these three perspectives. We embed a structural model within a welfare-based capital optimization framework. Welfare analysis determines the regulator’s optimal buffer budget, and a system-wide optimization allocates this budget across banks in the most effective way. This approach combines the macro-level welfare perspective with micro-level systemic differentiation, offering a more consistent and risk-sensitive foundation for macroprudential policy. In line with arguments posed by Nagel and Purnanandam (2020) we estimate the model in a period encompassing significant stress in the banking sector (Covid crisis, the Russian invasion of Ukraine and consequent energy price shocks, the first signs of inflationary pressure and monetary tightening) to ensure that tail dependencies are in-sample.

As a result, our approach rests on three key mechanisms. First, we introduce a structural credit model to relate the default risk of a single bank to its capital requirements. Second, we use a bank’s CDS price to estimate its distance from the default threshold; the

structural model then evaluates the counterfactual reduction in a bank’s default probability when its capital requirement is raised. Finally, a multi-factor model is used to infer the common exposures between banks to asset shocks on their balance sheet. With this approach, we can evaluate the consequent reduction in systemic risk, or *Expected Systemic Shortfall*, when a single bank’s capital buffers are raised. Thus, we associate the policymaker’s risk with that of holding a structural portfolio of risky and correlated credits. Then we evaluate the Central Bank’s (or the supervisor’s) exposure to systemic risk stemming from direct and indirect contributions by individual banks. The policymaker can manage this risk by setting bank-specific minimum capital requirements in view of the contribution of each bank to the risk of the regulatory portfolio. This framework is our key contribution to the literature on macroprudential buffers (and in fact, to the practice of setting them).

We apply our credit-based empirical analysis and this novel optimization framework to the list of European systemic banks compiled by the European Banking Authority. We find substantial gaps between our market-price-based optimal buffers and the macroprudential buffers currently in use. There is also a considerable heterogeneity in the way national policymakers set these capital buffers. The lack of consistent methodology makes it not only difficult to assess the adequacy of the buffers in any given country but has apparently also led to buffers of widely diverging stringency across countries within the Eurozone, as each country applies the general guidelines differently. For want of a generally accepted basic framework, aligning these diverging approaches has proven to be difficult so far; this arguably leads to a distortion of competition.

Our framework is built in a partial equilibrium setting, taking in particular banks’ risk-taking decisions as given.

Overall, relying on market data rather than regulatory data offers an objective way to verify the adequacy of the currently assigned macroprudential buffers and eliminates the well-documented drawbacks of using regulatory scores, end-of-year transaction data, and balance sheet data to assess systemic importance. For example, a number of studies find that G-SIBs tend to reduce activities affecting their systemic scores at year-end ahead of the reporting date, without this affecting their economic exposures, presumably as a way

to lower their capital buffer surcharges (Garcia et al., 2023; Naylor et al., 2024).

We use CDS prices primarily because our sample includes a number of large and smaller unlisted banks, making equity-based measures commonly employed in systemic risk models (Adrian and Brunnermeier, 2016; Acharya et al., 2017) infeasible. As a result, we follow the suggestion of Dimitrov and van Wijnbergen (2023) and focus on a CDS-based systemic risk model. However, several important considerations need to be addressed regarding our reliance on CDS data.

First, this choice, raises questions about potential Subordinated instruments are typically excluded from rescue packages, so their CDS spreads are not necessarily affected ex ante by the implicit put option associated with bail-outs, and while Berndt et al. (2025) document such downward pressure on credit spreads of systemically important banks, they document it only in senior unsecured CDS spreads.

This paper continues as follows: Section 2 discusses the relation of our study to the wider literature; Section 3 outlines the regulatory framework for capital buffers and the systems in place; Section 4 discusses the structure and mechanics of the credit portfolio model behind our estimates of systemic risk, allowing us to go from observed CDS spreads to asset covariances and systemic risk relations; in Section 5 we outline the risk optimization problem underpinning the process of regulating systemic risk and provide a cost and benefit approach to calibrate the aggregate level of macroprudential buffers; Section 6 provides the empirical application of our optimization approach to deriving macroprudential buffers using CDS based data to a sample of European banks. Section 7 concludes. .

2 Related literature

This paper connects the literature on systemic risk measurement to that of macroprudential policy and optimal bank capital.

First, regarding the measurement of systemic risk Lehar (2005); Huang et al. (2012); Adrian and Brunnermeier (2016); Brownlees and Engle (2017); Acharya et al. (2017); Engle (2018) all imply risk through asset price co-movements. The CoVaR approach of Adrian and Brunnermeier (2016) for example, relies on quantifying the tail loss of the

system given that a single bank is in the tail of its equity returns distribution. While intuitively appealing, the quantile nature of their measure makes it difficult to decompose the measure as it does not add up to a total systemic risk figure. In contrast, the MES approach by Acharya et al. (2017) and the DIP measure by Huang et al. (2012) define co-dependency as the expected loss of a bank given that the system is in its tail. The additivity of the expectation terms embedded in these measures provides an intuitive aggregation of these measures, which is why we follow their definitions of systemic dependence.¹

Second, there is extensive academic literature on macroprudential policy, but its focus is almost exclusively on interventions aimed at limiting household leverage in the mortgage market and on breaking the leverage-credit-housing build-up of systemic risk (Acharya et al., 2022). Instead, we look at regulators' ability to set additional buffer requirements on top of systemic banks' microprudential capital as a way of internalizing the implicit costs these institutions pose on the financial system through their contribution to aggregate systemic risk. Of course, funding of housing booms makes banks more prone to joint distress. We aim to capture this by allowing for systematic factors to drive banks' asset portfolio correlations and consequently the probability of their joint distress.

Furthermore, we build on the securitization literature in modeling the clustering of defaults in a credit portfolio (Vasicek, 1987; Tarashev and Zhu, 2006). Conceptually, the universe of banks in an economy can be viewed as a portfolio of defaultable loans, where the liability size of each bank represents the exposure that stands at default, and the potential default losses are indicative of the relative cost each institution will impose on the economy if it becomes distressed. Systemic risk then relates to the potential for multiple large defaults to occur, and the role of the regulator is to manage the credit risk of that portfolio by setting macroprudential add-on buffers.

Finally, some have argued for an approach based on a macroeconomic model in setting requirements by equalizing at the margin the costs in terms of reduced bank lending to the non-financial sector and the benefits in terms of lower expected costs of defaults: see

¹The main difference between MES and DIP is that the former defines the tail of the distribution of systemic scenarios as a quantile of the portfolio's distribution, while the latter sets it as losses above a fixed threshold. The two measures can be made identical by choosing the threshold appropriately.

for example Miles et al. (2013); BCBS (2010); Firestone et al. (2017); Cline (2017). To do this convincingly one needs a full-fledged macro-finance model. Such models have been developed but they either tend to abstract from risk contagion between banks by modeling the financial sector as a single large bank, like in Cline (2017), or as a continuum of ex-ante homogeneous banks. Both approaches make it impossible to incorporate asymmetric spillovers between banks, a key element in the mechanisms behind systemic risk (see for example Adrian and Brunnermeier (2016)). Therefore we follow the academic research on systemic risk, referred earlier, by putting heterogeneity among banks at the center of our analysis. At the same time, we simplify the macro side of the analysis into a single sensitivity parameter linking the average aggregate capital ratio to GDP directly.

Currently regulators use accounting data and heuristic estimates to adjust buffers for each bank until their estimated contribution to systemic risk equals the contribution of a (presumably non-systemic) reference bank. This is the so-called *Equal Expected Impact approach* (EEI).² Gauthier et al. (2012) points out a consistency problem in this approach, as the very risk spill-overs that necessitate macroprudential policy cause a fixed point problem for a regulator applying the EEI approach: lowering the initial systemic contribution of a bank with high systemic scores affects the systemic contribution of other banks (presumably that of the reference bank, as well), thereby necessitating another round of buffer setting and so on. We address this problem not by applying the EEI approach iteratively as Gauthier et al. (2012) suggests, but by defining an explicit welfare function which incorporates a trade-off between the social costs and benefits of capital buffers, and then optimizing it over the buffer add-on requirements for all banks simultaneously. In this way, we derive the optimal macroprudential buffers by directly internalizing the positive safety spillover externality highlighted by Gauthier et al. (2012).

²Cf. FRB (2015); Passmore and von Hafften (2019); Jiron et al. (2021); Geiger et al. (2022). In a companion paper we show that the EEI approach can also be applied within a credit-based framework, similar to the one shown here, instead of basing it on the heuristic rankings currently in use by policymakers.

3 A review of the regulatory framework

In practice there are two parallel systems of add-on capital buffers applied to systemically important institutions, the G-SII framework which addresses capital buffers for banks which are systemic in a global context, and the specifically European O-SII framework for banks found to be systemically important domestically. In both systems, central banks rely on a set of assigned regulatory scores of systemic impact to identify important banks. Banks are ranked according to size, interconnectedness, substitutability and importance in the lending market, complexity, and cross-border activity which are then weighted and aggregated into an overall score. Based on each bank's overall weighted score, regulators determine the size of the add-on requirements on top of the minimum loss-absorbing capacity that banks have to hold (the microprudential buffers). This is done either through approaches aiming to equalize the default-probability weighted scores between systemic and non-systemic institutions (the EEI approach), or through direct mapping from the systemic score of a bank to its capital requirement.

In mapping from systemic scores to buffers, the G-SII framework relies explicitly on the EEI approach. In Europe, apart from the EEI methodology, countries are also allowed to use an approach where banks with similar systemic (O-SII) scores are allocated to buckets and then the same buffer requirement is assigned across buckets. However, the number of buckets and the methods for their classification differ across countries. Also, even the assumptions and parameter choices behind the EEI methodology have been very diverse across countries (cf. ESRB (2017)). This inevitably affects the calibration of the buffers, allowing national regulators discretion over the size of the buffer requirements, with substantial differences across countries as a result, differences that are not in any way related to differences in systemic impact. The discretion national authorities have over the mapping from scores to buffers obviously stands in the way of a level playing field across the EU.

International discussions on the absence of a level playing field have been hampered by the absence of a common formal framework allowing for a systematic comparison of banks in different countries to Europe-wide systemic risk. Providing such a framework is precisely what we do in this paper.

4 Systemic shortfall and bank capitalization

Now, we provide a basic structure of the financial system with multiple banks subject to default risk, and the consequent aggregation of systemic losses. First, we consider banks' capital regulation, consisting of microprudential and macroprudential buffers. Then we look at bank default risk in isolation from other players in the system; then we look at the policymaker's objective in setting macroprudential buffer add-ons. This determines the underlying structure behind the policymaker's risk measurement problem defined in the subsequent sections. Finally, we define what drives distress correlations and look at how CDS data can be used to obtain the relevant empirical estimates.

4.1 Modeling the link from bank capital to bank risk

Micro- and macroprudential capital. Assume that there are n banks in the financial system indexed by $i \in (1, \dots, n)$. Let k_i denote the capitalization ratio of bank i observed at time t , defined as the ratio of common equity tier 1 (CET1) to risk-weighted assets (RWA).³ In a stylized representation bank capital can be split into the following components:

$$k_i = k_i^{mi} + k_i^{ma} + k_i^{headr} \quad (1)$$

The microprudential component k_i^{mi} is the minimum regulatory requirement that banks need to satisfy in view of their own creditworthiness.⁴

The macroprudential requirement, k_i^{ma} , which is the focus of this study, represents an additional buffer applied on top of the microprudential requirement. It is calibrated for each bank in light of its contribution to systemic risk and its role within the financial network. The purpose of this buffer is to enhance the resilience of the system as a

³While the regulatory framework on capital requirements in Europe is fairly complex (cf. Annex A.2 for details), we abstract from the various other types of capital and assume that CET1 is a good representation of a bank's equity with the view that it is the main going concern capital ratio in the Basel III framework.

⁴This includes, to a large extent, capital ratio requirements that are fixed at the same level for all banks, such as the minimum capital requirement (MCR) and the capital conservation buffer (CCB). In our empirical application, we allow also for a bank-specific micro component in order to capture Pillar 2 Requirements (P2R).

whole and safeguard overall financial stability. Finally, k_i^{headr} is the capital headroom a bank voluntarily holds to prevent breaching the total regulatory capital imposed by regulation (Gornicka and van Wijnbergen, 2013; Juselius et al., 2025). Going forward, we assume that the microprudential and the headroom components are exogenous to the macroprudential policy.

Bank default model. Now, assume that an exogenous stochastic latent variable $U_i \sim N(0,1)$ governs the risks to a bank’s creditworthiness over the coming year. Higher realizations of U_i indicate better states of nature. Assume that default occurs if the latent variable U_i falls below a threshold X_i . Up next we will show that this threshold is a function of the bank’s capitalization at time t given some basic structure of the assets and liabilities of the bank.⁵ This leads to the following default indicator function:

$$\mathbb{1}_i \equiv \begin{cases} 1 & \text{if } U_i \leq X_i \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

And we can write the default probability of bank i as the expectation that the latent variable will fall below the capitalization threshold:

$$\mathbb{P}(X_i) \equiv \mathbb{E}(U_i \leq X_i) \quad (3)$$

where $\mathbb{P}(\cdot)$ is defined over the probability distribution of the latent variable.

Importantly, given the observed (or inferred) default probability PD_i this formulation allows us to invert the relationship in (3) to obtain the threshold level directly from observed default probabilities:

$$X_i = \mathbb{P}^{-1}(PD_i) \quad (4)$$

This provides the foundation for the model components introduced next, as:

⁵The current baseline structure is part of the class of default threshold models commonly utilized in the credit risk literature and in practice. Bolder (2018); McNeil and Embrechts (2005) provide an extensive discussion.

- The default probability PD_i (given the actual bank capitalization) is inferred from the observed price on a credit default swap (cf. Equation (9));
- The Merton model is used to provide structure on a bank's default risk, such that the latent variable follows a standard Gaussian distribution, fixing $\Phi(X_i) \equiv \mathbb{P}(X_i)$ (cf. Equation (6)), allowing us to back out (1) the bank-specific default threshold implicit in the CDS price levels; (2) the relation between bank capital and bank default risk; (3) and the implied bank asset correlations, implicit in the banks' credit spread co-movements.

Link between a banks default probabilities and capital ratios. Next, we put further structure on the default process in line with the Merton model (Merton, 1974). Assume that the (unobserved) market value of bank i 's aggregate assets $V_{i,t}$ follows Merton's dynamics in continuous time under the risk-neutral distribution:

$$d \ln V_{i,t} = \left(r - \frac{1}{2} \sigma_i^2 \right) dt + \sigma_i dW_{i,t} \quad (5)$$

where, assuming market completeness, $r = \mu_i - \theta_i \sigma_i$ is the risk-free rate, with σ_i is the standard deviation of the bank's assets and θ_i is the price of risk consistent with bank i , and $dW_{i,t}$ is a Brownian motion.

As a result, non-viability (default), if it occurs, does so at maturity (time T) if the bank's assets fall below a fixed default threshold $D_{i,T}$ at that time. Defining the discounted value of debt as $D_{i,t} = D_{i,T} e^{-r(t-T)}$ we can then write the (risk-neutral) default probability for the bank as

$$\begin{aligned} PD_{i,t} &= \mathbb{P}(V_{i,T} \leq D_{i,T}) = \mathbb{P}(V_{i,T} e^{-r(t-T)} \leq D_{i,t}) \\ &= \mathbb{P} \left(V_{i,t} \exp \left(-\frac{\sigma_i^2}{2} (T-t) + \sigma_i (W_{i,T} - W_{i,t}) \right) \leq D_{i,t} \right) \\ &= \mathbb{P} \left(\frac{W_{i,T} - W_{i,t}}{\sqrt{T-t}} \leq \frac{-\ln \frac{V_{i,t}}{D_{i,t}} + \frac{\sigma_i^2}{2} (T-t)}{\sigma_i \sqrt{T-t}} \right) \end{aligned} \quad (6)$$

This relates very closely to the broad structure we provided in Equation (3). To see that, denote $U_i = \frac{W_{i,T} - W_{i,t}}{\sqrt{T-t}}$ and set the horizon $T - t = 1$. This is a stochastic

latent variable with a standard normal distribution that governs the risks to a bank's creditworthiness over the coming one-year period. A higher realization of U_i indicates a better state of nature and consequently a lower default probability. The Merton model also provides the default threshold as $X_i = \frac{-\ln \frac{V_{i,t}}{D_{i,t}} + \frac{\sigma_i^2}{2}(T-t)}{\sigma_i \sqrt{T-t}}$.

In order to relate this to the bank's capitalization, assume that v_i is the bank's RWA intensity, i.e. the ratio of risk-weighted assets to total assets, assumed fixed over the examination period; and that all bank debt consists of short-term deposits, such that equity is the asset value net of debt, $E_{i,t} = V_{i,t} - D_{i,t}$. The capital ratio as a percent of the bank's risk-weighted assets then is:

$$k_{i,t} = \frac{E_{i,t}}{v_i V_{i,t}} = \frac{V_{i,t} - D_{i,t}}{v_i V_{i,t}}$$

From this, it follows that

$$D_{i,t}/V_{i,t} = 1 - v_i k_{i,t}.$$

Inserting this into the right-hand-side of equation (6) and suppressing time notation, this leads to an expression linking the default probability and the CET1 ratio of a bank:

$$PD(k_i) = \Phi \left(\frac{\ln(1 - v_i k_i)}{\sigma_i} + \frac{1}{2} \sigma_i \right) \quad (7)$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution.

This relationship provides the default probability which a regulator targets when setting capital requirements. This is the key mechanism through which the regulator is able to reduce the contribution to systemic risk of a particular institution (cf. Section 5.2).

In order to use equation (7) the regulator needs to know σ_i . Note that here we do not have to fully calibrate the Merton model in order to evaluate this relationship. Section 4.2 discusses how σ_i can be estimated based on the bank's observed CDS price and current CET1 capital ratio.

Note that the application of the Merton model to banks has been criticized by Nagel and Purnanandam (2020) because of its log-normality assumption on the asset side. Banks' assets in reality consist of loans to firms and households, which indeed have a

skewed payoff as the bank does not participate fully in the upside of the underlying; our focus on downside risk implicit in the use of market prices of debt via the observed CDS pricing, rather than equity, arguably makes that a relatively minor problem.

Multi-variate setting. So far we have modeled a single bank in isolation from the others. The next step is to specify a process that drives the correlations between different banks' latent variables. In our approach, this is done through a set of common unobserved factors. This common component then influences the probability of multiple defaults happening at the same time. The exposure of each bank to these factors is determined by observing co-variations in the default probabilities of different banks.⁶ Formally, we can write:

$$U_i = \boldsymbol{\rho}_i \mathbf{M} + \sqrt{1 - \boldsymbol{\rho}_i \boldsymbol{\rho}_i'} \mathbf{Z}_i \quad (8)$$


where $\mathbf{M} = [m_1, \dots, m_f]'$ is the vector of f common latent factors, $\mathbf{Z}_i$ is the bank-specific factor, $\boldsymbol{\rho}_i = [\rho_{i,1}, \dots, \rho_{i,f}]$ is a vector of factor loadings, such that $\boldsymbol{\rho}_i \boldsymbol{\rho}_i' \leq 1$. Without loss of generality, all factors are selected to be mutually independent with zero mean and a standard deviation of one.⁷ In our baseline model, we use the standard Gaussian Copula framework, where all factors M and Z_i are assumed to be generated by standard normal distributions. Furthermore, since we operate under the assumption of the Merton model, the covariance matrix is fixed and we do not presume the existence of correlation premia.

Note that one gets the well-known Vasicek loan portfolio model as a special case from Equation (8) by assuming a single common factor and the same factor exposure across all banks. Furthermore, note the process in (8) is constructed to have a zero mean and unit variance, thus ensuring consistency on a univariate level with the definition of U_i derived in (cf. Equation (6)).

⁶This approach is statistical in nature and we do not aim to find a direct interpretation of the factors; however, in the literature, they have commonly been associated with market, industry, and geographically specific risk drivers (Pascual et al., 2006).

⁷The use of latent factors estimated from the common time variation in asset prices appears often in the systemic risk literature. See for example Pelger (2020) who uses five-minute tick data from the NYSE to identify a statistical factor model accounting for systemic risk and finds economic interpretation for the factors; and Billio et al. (2012) who use Principal Component factors to evaluate the evolution of systemic risk in the context of dynamic network interlinkages.

Figure 1: Conditional Default Probabilities



Note. Calibrated for $r = 0.;$ $\sigma = 3\%$.

4.2 Credit model estimation

Estimation of Default Probabilities from CDS spreads: Note that we do not estimate the Merton model to derive default probabilities. Instead, we use the observed CDS market prices to infer the distance-to-default of the bank based on the default probability embedded in the observed CDS price. We first have to infer this default probability. For this purpose, we use the pricing from Duffie (1999) (also refined in Tarashev and Zhu (2006)):

$$PD_i = \frac{aCDS_i}{aLGD_i + bCDS_i} \quad (9)$$

where a and b are known constants, CDS_i is the spread on the CDS contract written on bank i .⁸ Figure 2a plots the relationship.

There is ongoing debate in the literature on the extent to which CDS prices are affected by shocks to protection sellers that are unrelated to the protection buyer. Siriwardane

⁸Annex B.1 provides the details behind this CDS pricing formula, such as the definition of the constants a and b . We use a recover rate $RR_i = 1 - LGD_i$ assumption of 20% to extract the PDs, in line with the prospectus of the CDS contracts provided in Bloomberg. Since we use subordinate debt CDSs on financial institutions, this standard RR assumption is fairly conservative. The systemic literature employs a wide range of assumptions on the RR term. It is particularly difficult to set this parameter as defaults of systemic institutions and subsequent valuation of the collateral are very rare. Puzanova and Düllmann (2013) set the RR conservatively at 0%; Kaserer and Klein (2019) calibrate it based on the liability composition of the systemic institutions, assuming that deposits have a higher recovery rate than long-term corporate debt; Huang et al. (2009) set it to 55%; Huang et al. (2012) calibrate the RRs to Markit survey data and show that they exhibit very little time variation.

(2019) in particular finds evidence that shocks to the sellers’ balance sheets can distort CDS prices; on the other hand Augustin and Izhakian (2020) find that the marginal investor sets the price and is a protection buyer. We remain agnostic in this debate and sweep any potential liquidity price effects into the estimated default probabilities. To assess robustness, Annex A.8 compares systemic risk contribution estimates when asset correlations are derived from equity returns versus CDS spreads for the subsample of listed banks. The two approaches exhibit strong alignment, suggesting that CDS-based measures provide a reliable basis for systemic risk calibration while enabling coverage of non-listed institutions.

Estimation of banks’ asset variances: Equation 9 allows us to derive the implied volatility of a bank’s RWAs given its current capital ratio k_i and the CDS-implied default probability by solving numerically for σ_i ⁹. By inverting relationship (7) we can write the implied volatility derived from (7) as a function of the capital ratio and the PD_i we derived earlier:


$$\sigma_i = \hat{\sigma}_i(k_{i,obs}, PD_i) \quad (10)$$

where $k_{i,obs}$ is the currently observed CET1 ratio of the bank, and PD_i is the current default probability on which the bank’s CDS trades. Table 5 in Annex A shows the resulting implied asset volatilities for the sample of European banks used in the subsequent empirical analysis.

Figure 2b illustrates the relationship between $\hat{\sigma}_i(k_i, PD_i)$ and k_i for three different levels of the default probability PD_i . Each upward-sloping line traces out $\hat{\sigma}_i(k_i)$ as a function of k_i for a specific value of PD_i . Clearly, increases in the observed default probability produce upward shifts in the curve. The figure shows, quite intuitively, that if we observe a highly capitalized bank whose debt protection is priced at the same level as that of a low capitalization bank, it follows that the market perceives the assets of the first bank to be riskier than that of the second bank.

⁹We use the Python implementation of Brent’s root-finding algorithm, the ‘*brentq*’ method, to solve numerically for the variance in (7) given PD_i and $k_{i,obs}$.

Figure 2: Estimation of PD and Asset Variance from CDS prices



Note.

Estimation of asset correlations and factor loadings: The parameters of (8) can then be estimated by minimizing the sum of squared differences between a set of target correlations a_{ij} and the correlations implied from the factor loadings:

$$\min_{\rho_i, \dots, \rho_j} \sum_{i=2}^N \sum_{j=1}^N (a_{ij} - \rho_i \rho_j')^2 \quad (11)$$

where the optimization problem is typically solved numerically by the efficient algorithm provided by Andersen and Basu (2003). The Merton model allows us to link the target correlations to (time changes in) the default probabilities estimated from the observed CDS prices

$$a_{i,j} \equiv \text{Corr}(\Delta \Phi^{-1}(PD_{i,t}), \Delta \Phi^{-1}(PD_{j,t})) \quad (12)$$

This approach is standard in the securitization (Tarashev and Zhu, 2006) as well as in the systemic risk literature (Huang et al., 2012). We refer the reader to Annex B.2 for details. Table 5 in the Annex shows the fitted model exposures for the sample of banks that we analyze further. This will be one of the key inputs in our subsequent systemic risk analysis.

4.3 The systemic risk model

Relating to the default model of Equation (2), we define bank default losses as:

$$L_i = \begin{cases} LGD & \text{if } U_i \leq \Phi^{-1}(PD(k_i)) \\ 0 & \text{otherwise} \end{cases} \quad \text{for } i = 1, \dots, n \quad (13)$$

Systemic risk. We define aggregate systemic losses in the financial sector at a horizon of one year as a weighted sum of institution-level losses:

$$L_{sys} = \sum_{i=1}^n w_i L_i, \quad (14)$$

where weights w_i reflect institutions' relative size in the system (normalized to sum up to 1), with size measured by total bank liabilities. A systemic financial crisis occurs when aggregate losses exceed a critical threshold $\bar{L}$, such that the probability of a systemic event is

$$P_{sys} = \mathbb{P}(L_{sys} > \bar{L}). \quad (15)$$

This framework makes $\bar{L}$ a central policy parameter: it encodes the regulator's tolerance for aggregate losses before the system is deemed to be in crisis. A low $\bar{L}$ implies a more precautionary stance, triggering interventions for relatively modest losses; a high $\bar{L}$ implies that only very large losses qualify as systemic, allowing more losses to be absorbed privately.

Define the Marginal Expected Shortfall (MES) of a bank as its expected average loss conditional on total systemic losses being above a threshold $\bar{L}$:

$$MES_i = \mathbb{E}(L_i | L_{sys} > \bar{L}) \quad (16)$$


We can also quantify aggregate systemic risk (Expected Systemic Shortfall, ESS) as the potential default loss on a portfolio containing all banks in the financial system for

which the supervisor is accountable given that that aggregate loss exceeds $\bar{L}$:

$$ESS \equiv \mathbb{E}(L_{sys} | L_{sys} > \bar{L}) = \sum_i w_i MES_i \quad (17)$$

where the last line follows from the additivity property of expectations and the fact that all MES_i terms are conditioned on $L_{sys} > \bar{L}$. This provides a useful interpretation of the MES: a bank's weighted MES represents the portion of total systemic risk that it brings in. Lowering the bank's MES by imposing higher capital buffers thus will lower overall systemic risk.

Figure 3: Simulated Example of Macroprudential Capital Increases



Note. The simulation is parameterized with symmetric assumptions where each bank has an asset volatility of $\sigma_i = 2\%$, equal weights $w_i = 0.5$, $LGD_i = 0.8$, $\nu_i = 0.3$, base microprudential buffers of 7%, $\bar{L} = 0$, $\rho = 0.5$.

Simulated Examples. With the underlying structure of the model, we proceed to evaluate the basic properties of the systemic risk measures defined so far in a simulated example involving two symmetric, equally sized banks with correlated assets. The plots in Figure 3, illustrate the monotonic reductions in a bank's PD, MES, and in the aggregate measures of systemic risk, as its macroprudential capital requirements are raised.

The results are intuitive: increasing the macroprudential buffer of Bank 1 reduces its probability of default, as shown in panel (a). From a systemic perspective, this leads to a decline in its marginal expected shortfall (MES), while the MES of Bank 2 increases (panel b), reflecting the fact that as Bank 1 becomes more resilient, the banks become jointly distressed only under very severe scenarios, i.e. Bank 2 must now be exposed to more extreme shocks to be distressed alongside bank 1. Nevertheless, the aggregate effect is beneficial: both the expected systemic shortfall (ESS, panel (c)) and the probability of a systemic crisis (P_{sys} , panel (d)) decline. Finally, the 3D surface in panel (e) shows that a regulator targeting a specific ESS can afford to be less strict with Bank 1's buffer if Bank 2 is already well-capitalized. From that point of view, in some cases it is more effective to manage systemic risk by allocating buffers across both banks, rather than suppressing a single one, as the tradeoff in capital buffers between Bank 1 and Bank 2 can be utilized to achieve a targeted level of systemic risk reduction.

5 The Policymaker's optimization problem

5.1 Optimal capital relative to $\bar{k}$

We model the policymaker's task of maintaining financial stability as a two-step process: first, the socially optimal average macroprudential buffer $\bar{k}$ is determined; second, this average is allocated across individual banks to minimize systemic risk. Formally, the allocation problem can be expressed as:

$$\begin{aligned} \min_{k_1^{ma}, \dots, k_n^{ma}} & \left\{ \mathbb{E} (L_{sys} \mid L_{sys} > \bar{L}) \mathbb{P}(L_{sys} > \bar{L}) \right\} \\ \text{s.t.} \quad & \sum_{i=1}^n w_i k_i^{ma} = \bar{k}; \quad k_i \geq 0, i = 1, \dots, n \end{aligned} \tag{18}$$

where $\bar{L}$ is the systemic loss tolerance threshold, as before, and $\bar{k}$ is the target average macroprudential add-on. This objective then ensures that the macroprudential buffers are allocated across banks in a way that optimizes the combined effect of two components: the probability of a systemic crisis and the expected severity of losses conditional on such a crisis. While $\bar{k}$ may vary over time in reality —e.g., due to countercyclical capital

buffers— we treat it as constant, focusing solely on the allocation of capital buffers in the cross-section in line with the O-SII and G-SII regulation. Also, the microprudential requirements here are fixed and exogenous to the macroprudential decision, again in line with the policy practice.

Solving numerically the optimization problem in (18) is not straightforward. As shown in Figure 3, both the Expected Systemic Shortfall (ESS) and the probability of a systemic crisis decrease monotonically with higher capital buffers. However, the objective function surface exhibits a broad flat region (see Figure 3e), which makes constrained multidimensional optimization challenging.

To address this, we proceed in three steps. First, we eliminate the budget constraint on $\bar{k}$ not by adding it as a separate restriction in the numerical maximization process but by incorporating it directly via a SoftMax transformation (see Annex B.3). This procedure guarantees that the average buffer budget is always fully exhausted while the non-negativity constraints on the individual macroprudential buffers are always satisfied. This allows for the use of more efficient unconstrained optimization algorithms. Second, to improve the likelihood of finding a global optimum, we initialize the optimizer from multiple random starting points. For each starting point, we first obtain a coarse solution using Powell’s method and then refine it with a more precise but computationally intensive algorithm (Nelder–Mead). Finally, we compare all candidate solutions and select the one that achieves the lowest objective value.


The SoftMax transformation ensures that the macroprudential buffer budget is always optimally allocated, such that an unconstrained optimization of the transformed objective directly provides the minimum feasible Expected Systemic Shortfall $ESS(\bar{k})$ and the minimum feasible probability of a systemic crisis, $P(\bar{k})$ as a function of a given budget $\bar{k}$. The final question then is: what is the socially optimal level of this budget. We address that issue in the next section.

Simulated Model Comparative Statics. Before solving for the optimal $\bar{k}$ in the next section, we illustrate through several simulated examples how, given a budget $\bar{k}$, optimal macroprudential buffers respond to changes in a banks’ asset risk, correlation, and size. Each example begins with a system at its optimized state; we then vary one parameter

and compare the resulting unoptimized scenario with the new optimized counterfactual.

First, Figure 4a shows a two-bank system where Bank 1's asset volatility rises from 2% to 15%, while Bank 2 remains unchanged. With an average buffer of 10% to be allocated between the two banks, the policymaker reallocates buffers toward Bank 1 as its risk increases (Panel a), reducing systemic crisis probability compared to the counterfactual of no intervention (Panel b). Bank 1's default probability (PD) still rises but stays below its baseline, compared to the no-intervention case. Note that at the new optimum Bank 2's PD increases slightly as it becomes optimal for the policymaker to relocate buffers away from Bank 2 and towards Bank 1, reflecting the systemic perspective in the a trade-off between the two banks (Panel c).


Figure 4: Simulated Example of Optimal Macroprudential Buffers with σ_1 increasing



Note. The simulation is parameterized with symmetric assumptions where each bank has equal weights $w_i = 0.5$, $LGD_i = 0.8$, $\nu_i = 0.3$, base microprudential buffers of 7%, $\bar{L} = 0$, $\bar{k} = 10\%$ and a single systematic factor with bank exposures of $\rho_1 = \rho_2 = 0.5$.

Figure 5, then, examines correlation effects in a three-bank system, . As Banks 1 and 2 become more correlated (their common factor exposure ρ goes from 0 to 0.8), buffers shift toward them, while Bank 3, designed to bear only idiosyncratic risk, receives less. This allocation curbs joint failure risk among interconnected banks, even if Bank 3's PD rises. Systemic crisis probability still grows with correlation (the dashed curve, Panel (c)), but optimal allocation mitigates much of the risk relative to the case that is not optimized (the solid black curve). Now, the policymaker is willing to tolerate more default risk for the idiosyncratic bank in order to mitigate with higher buffers a potential joint default of banks 1 and 2 (Panel (b)). Overall, as interdependence between two banks intensifies, macroprudential policy concentrates resources on the most interconnected institutions to contain systemic risk.

Figure 5: Simulated Example of Optimal Macprudential Buffers with $\rho_1 = \rho_2$ increasing; ρ_3 fixed




Note. In each case, we fix the exposure of Bank 3 $\rho_3 = 0$ and vary the exposures to the systematic factor of Bank 1 and Bank 2 such that $\rho_1 = \rho_2$. The simulation is parameterized with symmetric assumptions where each bank has equal weights $w_i = 1/3$, $LGD_i = 0.8$, $\nu_i = 0.3$, base microprudential buffers of 7%, $\bar{L} = LGD * 1/3$, $\bar{k} = 2\%$ and a single systematic factor.

Finally, Figure 6 explores size effects. In a three-bank system, again, with fixed average buffers at 2%, Bank 1's weight increases relative to the others. The systemic threshold $\bar{L}$ is set such that a systemic crisis is defined by the default of at least one-third of the sector. What we observe is that below one-third of the total, Bank 1 receives no buffer (solid black curve in panel (a)); beyond this threshold, however, the optimal buffers start concentrating on Bank 1, approaching the full 2% budget as the bank dominates systemic risk. This prioritization lowers systemic risk despite tolerating higher PDs for smaller banks (panel (b)), underscoring the policy trade-off between systemic stability and uniform risk reduction. Below one-third, Bank 1 is effectively non-systemic; above two-thirds, it dominates systemic risk, and the optimization concentrates buffers accordingly.

5.2 The optimal $\bar{k}$

Social Welfare and Capital Buffers. In setting the average buffer add-on in the system, the regulator balances the costs and benefits of holding higher buffers. We look at the policymaker's problem as one of maintaining a healthy supply of credit in the economy while managing the risk of a systemic crisis from materializing. Higher capital requirements reduce the probability of a financial crisis, but they also bring the risk of inducing a reduction in lending if banks struggle to satisfy the stricter regulatory constraints. The policymaker thus needs to balance the two costs for the economy. We

Figure 6: Simulated Example of Optimal Macroprudential Buffers with w_1 increasing



Note. This figure illustrates the optimal macroprudential capital buffers and the risk evaluated at the optimum for a system of three banks. The simulation is parameterized with symmetric assumptions: $LGD_i = 0.8$, $\nu_i = 0.3$, $\rho_1 = \rho_2 = \rho_3 = .5$ base microprudential buffers of 7%, $\bar{L} = LGD * 1/3$, $\bar{k} = 2\%$ and a single systematic factor; the policy weight of Bank 1 is increased while that of banks 2 and 3 is decreased pro rata.

formalize this choice by, for a start, defining a welfare cost function measuring the social cost of holding capital buffers:

Then, we can define in more detail the welfare cost function as

$$W_c(\bar{k}) = \underbrace{\lambda \cdot ESS(\bar{k})}_{\text{GDP Loss in Crisis}} \cdot \underbrace{P(\bar{k})}_{\text{Probab. of crisis}} + \underbrace{\eta \cdot \bar{k}}_{\text{Social cost of buffers}} \cdot \underbrace{(1 - P(\bar{k}))}_{\text{Probab. of no crisis}} \quad (19)$$

To keep the model tractable, we assume a linear relationship between GDP decline in a financial crisis and the monetary losses in the banking sector given a financial crisis. Then, λ is defined as a multiplier, which translates the expected tail losses in the financial sector to losses for the wider economy, calibrated to a particular scenario (more on this in Section 5.3).¹⁰

On the other hand, if no systemic distress occurs, the public has to bear the cost to the economy of imposing capital buffers of $\bar{k}$ (on average across banks) above their microprudential requirements. We quantify this cost through the response of aggregate output (Y) to an increase in capital requirement. This response runs via reduced credit lending in the economy, as banks need to satisfy the higher capital requirements by accessing the possibly more expensive source of financing that common equity imposes, or

¹⁰Note that even though the λ here is fixed, the capital buffers still affect the ESS in a non-linear way: a lower buffers would imply higher default rates, and this higher default losses. The buffer costs in a crisis are counted via λ as well.

by reduced risk-shifting incentives.¹¹ Our goal here, however, is to quantify this response as a relative rate of change term, but without specifying a full-blown macro model, so we define:

$$\eta = -\frac{dY/d\bar{k}}{Y} = - \underbrace{\left(\frac{dY}{dC} \cdot \frac{C}{Y}\right)}_{\text{Elasticity: GDP to Credit Supply}} \underbrace{\left(\frac{dC}{d\bar{k}} \cdot \frac{1}{C}\right)}_{\text{(semi-)Elasticity: Credit to Capital Buffer}} \quad (20)$$

where C is the total equilibrium level of credit in the economy. η then represents the percentage drop in GDP for a one percentage point increase in the macroprudential ratio requirement.

Given (18) it is then straightforward to show that:

$$\frac{\partial P(\bar{k})}{\partial \bar{k}} < 0; \frac{\partial ESS(\bar{k})}{\partial \bar{k}} < 0$$

where the partial derivative reflects what we already established earlier, that higher overall buffer levels lower the probability of individual and systemic distress and the realized systemic losses.

Finally, the policymaker chooses the optimal $\bar{k}$ that minimizes (19). The first-order condition implies that at the optimum $\bar{k}$ the expected benefits (in terms of reduced expected crisis losses) of a marginal increase in the macroprudential capital just compensates for the marginal increase in the macroeconomic costs (associated with lower credit growth):

$$\lambda \left(P \frac{\partial ESS}{\partial \bar{k}} + \frac{\partial P}{\partial \bar{k}} ESS \right) = \eta \left(\bar{k} \frac{\partial P}{\partial \bar{k}} - (1 - P) \right)$$

Note that $\bar{L}$ determines the intervention threshold for the regulator, and implicitly can be linked to the regulator's loss tolerance. Below $\bar{L}$ the regulator provides Lender of Last Resort (LoLR) funding to stave off a crisis, but when losses exceed $\bar{L}$ a financial crisis becomes unavoidable. In line with the previous section, we treat all losses on a

¹¹Jakucionyte and van Wijnbergen (2018) for example discuss the macro effects of higher capitalization of the banking system with risk-shifting and debt overhang problems through the latter's impact on aggregate credit supply.

relative scale, i.e. as a percent of the size of the system measured by the outstanding liabilities of the financial sector.

5.3 Welfare calibration

Consider now the parametrization of the optimization problem (19) which will allow us to set the socially optimal average buffer size $\bar{k}$. The core assumption behind our approach is that increased capitalization requirements will lead to a negative lending shock; low average capitalization, on the other hand, will lead to an increased probability of a financial crisis and increased welfare loss. We refrain from making explicit assumptions about the channel through which lending is reduced, this would require a full banking model, which is outside of the scope of this paper. Also, we refrain from the approach used in part of the literature of quantifying specific Modigliani-Miller deviations and of estimating the subsequent pass-through of the associated higher financing costs to the public (see Cline (2017); Brooke et al. (2015); Miles et al. (2013)). Empirically, the two questions have been a matter of debate: see for example Dick-Nielsen et al. (2022), who cast doubt on the claim that deviations from MM would increase significantly the cost of capital to banks.

Instead, we focus on studies quantifying the relationships between output and lending shocks on one hand, and lending decline due to capital requirements on the other. By combining the two, we hope to capture the overall causal effect from increased capital requirements to output losses.

Calibrating the cost of higher buffers (η). Empirically, macroprudential buffers have been found to constrain the supply of credit for the individual banks that are targeted. Using regulatory data Cappelletti et al. (2019) find that in the short run banks identified as O-SII cut the credit supply to households and the financial sector, even though in the medium run this tendency is diffused. In a diff-in-diff setting Behn and Schramm (2021) do not find a significant effect on the overall lending activity of G-SIB designated companies, but find a significant shift towards lending to less risky counterparties. Degryse et al. (2020) find a more pronounced effect by focusing on a narrower

time window and unexpected G-SIB designations. Favara et al. (2021) look at the US and find that banks designated as G-SIBs do reduce their credit supply but the aggregate effect is muted as firms switch to non-G-SIB banks. In their estimate, a one percentage point increase in macroprudential capital surcharges leads to loan commitments by GSIBs banks to fall by 3–4% on average relative to other banks. We use their upper estimate as the most conservative figure on the negative impact of capital requirements on the supply of credit in the economy in the short run. This gives us an estimate of the term $(dC/d\bar{k})(1/C)$ in 20.

On quantifying the impact of reduced credit supply to lower GDP growth, we rely on Barauskaitė et al. (2022), who in a BVAR framework identified with sign and inequality restrictions, determine that a 1% reduction in loan supply would result in a worst-case scenario in about .6% reduction in GDP growth.¹² This provides an estimate of $(dY/dC)(C/Y)$. Combining these two numbers (4% and 0.6%) gets us a baseline figure for the η in equation (20) of .024.

Calibrating the threshold of a systemic crisis ($\bar{L}$). Next, we need to calibrate $\bar{L}$, the threshold loss parameter of government involvement to prevent a financial crisis from becoming systemic. To do that, we refer to Laeven and Valencia (2020) who evaluate bank restructuring fiscal costs associated with systemic crises to at least 3% of GDP. This number has also been used as a minimum bank capital shortfall rate in policy exercises, as in Brooke et al. (2015).¹³ Within the Euro-area for the period under consideration, total bank liabilities are above three times the value of GDP. As a result, we set $\bar{L}$ to 10%.

Calibrating and interpreting the cost of a financial crisis.

¹²Refer in particular to Figure 3 in Barauskaitė et al. (2022) outlining the impulse-response functions from a credit supply shock.

¹³Refer to Eurostat’s database to source the value of deposit holding institution’s liabilities relative to GDP for the countries in the Euro-area, and weight by the relative size of each country.

Table 1: Model Calibration with Parameter Interpretation

Meaning	Variable	Value	Source
Systemic loss threshold triggering crisis	$\bar{L}$	0.1	Implied from Laeven and Valencia (2020) 3% of GDP restructuring costs, matched by the size of the European banking sector
Elasticity of GDP to credit supply	$\eta_{Y,C} = (dY/dC)(C/Y)$	0.006	Barauskaitė et al. (2022)
Elasticity of credit to macropr. buffer	$\eta_{C,\bar{k}} = (dC/d\bar{k})(1/C)$	-0.04	Favara et al. (2021)
GDP cost of higher buffers (per 1pp increase)	η	0.024	$\eta_{Y,C} \times \eta_{C,\bar{k}}$
Pass-through from systemic losses to GDP (severe crisis)	λ_{sev}	0.2552	Implied from Reinhart and Rogoff (2009), 9% GDP loss
Pass-through from systemic losses to GDP (moderate crisis)	λ_{mod}	0.1692	Implied from Romer and Romer (2017), 6% GDP loss
Pass-through from systemic losses to GDP (mild crisis), 6% GDP loss	λ_{mild}	0.1128	Implied from Romer and Romer (2017), 4% GDP loss

Note. This table adds an interpretation column clarifying the economic role of each parameter in the welfare optimization problem.

6 Empirical Results

In Section 6.1 we apply our framework to the entire sample of European banks in order to fully capture risk spillovers across jurisdictions in Europe. We calibrate and optimize the social welfare function and evaluate the optimal capitalization level of individual banks in the sample. We find substantial variation in the deviations between the optimal buffers so derived and the currently prevailing buffers (see Figure 8). To see whether these deviations are due to different optimal averages across countries or to differences in the way the averages are allocated over individual bank within and across jurisdictions, in Section 6.2 we apply our optimization framework taking as given the capital budget. We compare the resulting optimal buffers to the current O-SII rates, designed to safeguard systemic spillovers from the banking sector (see Section 3 for a discussion of O-SII buffers). The banks that are also designated as G-SIBs are subject to an addition G-SII buffer, but only the larger of the O-SII and G-SII buffers applies, and since in our sample the G-SII buffer is never larger, comparing to the O-SII buffers is appropriate.

Data Overview. We now turn to the empirical application of the framework established in the previous sections. Our data sample consist of the 26 European banks designated in 2022 as systemically important by the European Banking Authority.¹⁴ We have dropped four banks (BPCE, Banque Postale, DNB ASA, Nykredit) from the original list of 30 banks due to insufficient CDS data, and have added a smaller Dutch domestically significant bank, Volksbank (VB), at the request of the Dutch Central bank. VB is too small to bias the systemic measures or the weighted buffer budgets. We discuss the data in further detail in Annex A.

Liquidity bias.

6.1 The socially optimal average buffer $\bar{k}$


We find the socially optimal capital buffers through a dual optimization problem. First, given an average capital $\bar{k}$ to be allocated across banks, the policymaker optimizes over its allocation across banks; and second, the policymaker decides on the optimal $\bar{k}$ by minimizing the welfare cost function specified in Section 5.2, given an optimal allocation of the $\bar{k}$ budget. Eventually, this allows us to simultaneously find the optimal $\bar{k}$ while also ensuring that it is optimally allocated across banks based on their systemic risk profile (see Equation (19)).

Figure 7 shows respectively the estimated probability of a financial crisis occurring (defined as earlier as bank default losses exceeding 3% of GDP) as a function of the average macroprudential buffers; and the welfare cost function with the three parametrization scenarios on the severity of the financial crisis. A policymaker picks the level of $\bar{k}$ which minimizes the expected welfare cost.

Note the strong convexity of the cost function in Figure (15b) for each of the parametrizations: as macroprudential buffers are initially introduced, we see a rapid decrease in the welfare costs, as the initially high probability of a financial crisis in a system with only microprudential capital rapidly declines when bank-specific macroprudential buffers are added. As the buffers increase, the probability of a financial crisis decreases pushing down

¹⁴For the EBA list see <https://www.eba.europa.eu/risk-and-data-analysis/risk-analysis/risk-monitoring/global-systemically-important-institutions-g-siis>

Figure 7:



Note. This set of figures shows the calibration of the social disutility function and the resulting optimal $\bar{k}$. Chart (a) shows the risk-neutral probability that a financial crisis occurs when it is defined as bank losses greater than 3% of GDP. Chart (b) shows the resulting welfare cost of introducing macroprudential buffers. The model is calibrated for three cases: Severe (9%), moderate (6%), and mild (4%) GDP loss triggered by a financial recession. The grey lines show the optimal macro capital buffers in each case, evaluated at respectively.


the overall welfare costs. At some point the marginal benefits of higher buffers become smaller, and more weight in the cost function is put on the economic cost of maintaining those higher buffers. Then, the decrease in expected costs slows down and is consequently reversed.

From Figure 15b, we can identify the average macroprudential buffer at which the welfare cost function W_c reaches its minimum. A higher expected GDP decline from a systemic crisis—captured by the pass-through rate λ linking financial losses to real economic losses—requires a more conservative stance, meaning the policymaker must lower the probability of such a crisis by setting a higher optimal capital ratio. Under the severe-crisis scenario (solid black line), welfare is minimized at average macroprudential buffers of 11.58%. For moderate and mild scenarios, the optimal values are 10.52% and 9.47%, respectively. The associated GDP cost of maintaining these buffers is higher in the severe-crisis case because the welfare function places greater weight on avoiding large systemic losses, while higher buffers also entail larger credit supply reductions for banks required to hold them.

Figure 8 zooms in on the distribution of capital buffers across individual banks. The chart compares current CET1 capitalization ratios (k_i) with model-based total capital requirements (k_i^*) under the three pass-through calibration scenarios (mild, moderate,

severe). This comparison highlights how far actual capitalization diverges from what the welfare-based optimization suggests.

Figure 8:



Note. This figure shows the current CET1 capitalization ratios (k_i) vs. the model-based total capital requirements (k_i^*), i.e. the current micro (MCR, CCB, and P2R) plus the model macro requirement, at the socially optimal $\bar{k}$.

Several patterns stand out. First, the model implies substantial heterogeneity in optimal buffers across banks, even under a common average $\bar{k}$. For some institutions, such as CAIX, SWED, and SEB, the welfare-optimal allocation remains close to their microprudential baseline (Mild scenario around 8%), indicating that their contribution to systemic risk is relatively small. In contrast, banks like DANK, BAY, UNIC, and RABO receive significantly higher buffers under the model (Mild scenario above 21%), reflecting their size, interconnectedness, or risk profile. These differences clearly relate to the systemic perspective of the optimization: buffers are concentrated where they most reduce aggregate risk rather than distributed uniformly.

Second, the gap between actual and model-based capitalization varies widely across countries. Table 2 illustrates these differences clearly by comparing the weighted average optimal macroprudential buffer (above the microprudential requirement of 7% and Pillar 2 Requirement) against the total actual macroprudential CET1 capitalization. In some jurisdictions, such as Sweden, current CET1 ratios of individual banks are well above the model's optimal levels, resulting in negative gaps under all scenarios (e.g. -7 percentage

points under the mild case). By contrast, most other countries fall short—sometimes by several percentage points: Germany and Italy exhibit gaps of roughly 6 – 9 percentage points in the range between the mild and the severe scenarios, reflecting the systemic importance and interconnectedness of their largest banks. France shows intermediate gaps (around 4 – 6 percentage points), while the Netherlands displays only modest shortfalls (about 1 – 4 percentage points). Spain remains on the lower end with gaps near 2 – 4 percentage points, and Sweden stands out with actual buffers exceeding model based recommendations across all scenarios.

These patterns suggest heterogeneous macroprudential practices across Europe, the absence of a harmonized methodology. Adopting a model-driven approach could thus improve consistency and risk sensitivity in buffer allocation, reducing systemic risk without imposing unnecessary costs on low-risk institutions.

Table 2:


Country	Actual Buffer	Model Mild	Model Moderate	Model Severe
France	5.48	8.99 (3.51)	10.07 (4.59)	11.44 (5.96)
Germany	6.02	11.69 (5.67)	12.86 (6.84)	14.03 (8.00)
Italy	6.48	13.13 (6.65)	14.19 (7.71)	16.06 (9.58)
Netherlands	8.60	10.28 (1.68)	11.31 (2.71)	12.13 (3.53)
Spain	4.58	6.83 (2.25)	7.73 (3.15)	8.10 (3.52)
Sweden	11.17	4.13 (-7.04)	4.75 (-6.42)	4.67 (-6.50)

Note. This table shows weighted average additional capital buffers (beyond microprudential requirements of 7% plus Pillar 2 Requirement) per country compared to model-based optimal buffers under three systemic risk scenarios (mild, moderate, severe). Values in brackets indicate the difference between model-based and actual buffers (Model minus Actual). Positive values imply model recommends higher buffers than currently observed; negative values indicate actual buffers exceed model-based levels.

Drivers of Optimal Capital Buffers.

Optimal Capital Compared to Previous Studies. Even though our approach differs fundamentally from the methodologies used in prior studies, most of which rely on a single representative bank macroeconomic models to drive the cost-benefit analyses at an aggregate level (BCBS, 2010; Miles et al., 2013; Firestone et al., 2017; Cline, 2017), our results are still broadly consistent with the literature on optimal capital buffers. We find a weighted average capitalization in the range of 17.5% to 19.6%, depending on assump-

Figure 9:




Note.

tions about the severity of systemic crises, compared to the current average of 14.3%. These figures are close for example to the upper bound of estimates reported by Miles et al. (2013), who suggest optimal Tier 1 ratios of 16–20%, and align with the midpoint of the range proposed by Firestone et al. (2017), which spans 13–26%. They also fall within the range where the Basel Committee’s long-term impact study finds net benefits remain positive (15–19%, see BCBS (2010)). The key distinction is that our framework targets systemic importance explicitly and differentiates between individual banks rather than prescribing a uniform sector-wide ratio. By embedding systemic risk contributions into the optimization, our model concentrates buffers where they most reduce aggregate risk—on large, highly interconnected institutions—while avoiding unnecessary surcharges for smaller or less correlated banks. This contrasts with aggregate approaches that often abstract from heterogeneity and cross-border spillovers, and underscores the value of a model-driven allocation mechanism for macroprudential policy.

Robustness of the optimal buffers.

Figure 10: Macroprudential Buffers. Sensitivities



Note. This figure reports welfare sensitivities of the optimal average macroprudential buffer $\bar{k}$. Panel (a) shows the crisis probability $P(\bar{k}) = \Pr(L_{sys} > \bar{L})$ for alternative systemic-loss thresholds $\bar{L} \in \{5\%, 10\%, 15\%\}$. Panel (b) shows the corresponding welfare function $W_c(\bar{k})$; vertical lines mark the minimizing $\bar{k}$ in each case. Panel (c) varies the output cost of buffers, $\eta \in \{0.014, 0.024, 0.034\}$, holding other parameters fixed. Crisis-severity scenarios (mild/moderate/severe) are implemented by adjusting the crisis-cost mapping λ to match assumed GDP losses in a no-macro-buffer benchmark, i.e. $\lambda_g = g/ESS(0)$.

6.2 Comparison between current O-SII buffers and model-based buffers


Next, we examine whether the suboptimality documented in the previous section results from between-country differences or from differences in the way national supervisors allocate given averages.

Table 3:

	France	Germany	Italy	Netherlands	Spain	Sweden	Europe
Actual	1.13	1.36	0.87	2.10	0.76	1.00	1.25
Model	1.18	1.72	2.28	0.95	0.98	0.00	1.25
Actual – Model	- 0.05	- 0.36	- 1.42	1.15	-0.23	1.00	0.00

Note. Current weighted averages of O-SII buffers per country and model-based rates based on an optimization relative to the European average. Values in the last row indicate difference from O-SII Buffer (Model minus Actual). The actual and the model on European level coincide by construction.

Figure 11:



Note. This figure shows optimal buffers against the regulatory set O-SII buffers. The optimal buffers are determined relative (1) to the Euro-wide average O-SII buffer rate as $\bar{k}$ and the full European sample (dark grey bubbles), and (2) to the country-wide O-SII average and the country-wide O-SII correlations only (light blue).

6.3 Model vs. Actual: Distributional Tests

On a Euro scale. Up until now we have compared model outcomes and actual buffers in an informal way. We now proceed to formally testing for differences by applying a two-sample Kolmogorov–Smirnov (KS) test to the full cross-section of banks ($N = 27$). Comparing the optimal buffers from a Euro perspective (see Table 6 in the Annex) to the actual *O-SII buffers* (see Table 5) yields a KS statistic of $D = 0.370$ and a p -value

of 0.048. We therefore reject the null hypothesis that the two samples are drawn from the same distribution at the 5% level. The results are confirmed also after we bin the model-allocated buffers to buckets of size .5% in the range of $[0; 3.5)$, typically employed by policymakers, and conduct a Pearson Chi-Squared test, indicating again that the two allocation rules are statistically different. Table 10 in the Annex provides the binned data. Since we have two distributions and seven full bins, the chi-squared test is of degree $(2 - 1) \times (7 - 1)$. The test then yields $\chi^2(6) = 15.711$ with $p = 0.015$. Thus, at conventional levels we reject the null that the two histograms are drawn from the same underlying distribution, in line with the KS evidence.

Consistent with this result, the empirical distribution of model-based buffers places more mass in the very low ($< 0.5\%$) and higher ($\geq 1.5\%$) ranges, while actual O-SII buffers are concentrated primarily between in the mid ranges between 0.5% and 1.5%, implying a statistically and economically different allocation pattern. This is consistent with our earlier observation, the optimization identifies banks that are dependent on the key systemic factors driving common default dependency, and allocates them with the majority of the buffer budget available to the policymaker, while being more relaxed on banks that are more idiosyncratic. See Annex A.6 for details.

On a local scale. We repeat the exercise at the local scale by comparing the national *O-SII buffer* rates to the model’s locally optimized buffers. For the set of banks with both values available ($N = 23$), the two-sample KS test yields $D = 0.261$ with a p -value of 0.422, so we cannot reject the hypothesis that the two distributions are the same. Hence, conditional on operating within each jurisdiction’s average buffer budget, the local model allocation produces a distribution of bank-level buffers that is statistically indistinguishable from the distribution of actual O-SII rates. As a result, we cannot reject the hypothesis that policymakers are setting O-SII buffers at a national level that are optimal with respect to the local system of banks only. This is again confirmed by discretizing the model-produced buffers (see Table 11 in the Annex) and comparing to the actual O-SII rates via the Pearson Chi-Squared test. Since again we have two distributions and six independent bins, the chi-squared test is of degree $(2 - 1) \times (6 - 1)$, yielding $\chi^2(5) = 9.807$ with $p = 0.081$. Hence, we do not reject equality of the two binned

distributions at the 5% level (borderline at 10%).

7 Conclusions

In this paper we address the problem of how to calibrate macroprudential capital buffers of banks. To that end, we developed a novel optimizing framework that links overall systemic risk to the size of banks' capital buffer requirements. The model works as follows.

First, we set up a credit portfolio model with endogenous default thresholds. The credit portfolio represents a Central Bank's exposure to commercial bank liabilities in case of default of one or more of those banks. This model can be used to quantify the optimal macroprudential requirements for each bank in the universe considered. Second, we define a tail-risk-based measure of the expected cost of default which captures the externality costs imposed by systemic institutions by explicitly incorporating tail-risk dependencies. We then embed this credit model in a portfolio risk framework and take a two step approach to optimizing add-on macroprudential buffers: in the first step we solve a risk minimization exercise subject to an average capital buffers target. The optimal level of systemic risk and the probability of a systemic crisis then become a function of that average. In the second step we solve for the optimal average buffer size on which the distribution of buffer sizes is conditioned. In our final application, the two steps are performed simultaneously providing at the same time the socially optimal average level of capital buffers, as well as the allocation of this average across individual banks.

We apply this framework to a universe of 27 large European banks. We use CDS prices rather than equity returns to infer the default probabilities, asset variances, and default correlations between the different institutions; using CDS prices instead of equity returns allows us to integrate into the analysis systemic banks that are not traded on equity markets.

We then set up the optimization-based cost/benefit analysis of macroprudential capital requirements, a trade off between the benefits in terms of reduced contributions to systemic risks against the costs of higher capital requirements in terms of reduced credit availability. Solving for each individual bank's macroprudential buffer is done in two

steps: in the first step we optimize individual banks' macroprudential ratios subject to (and as a function of) the average capitalization buffer. In step two we derive that optimal average add-on macroprudential buffer, again trading off the benefits of reducing systemic risk against the benefits of higher credit availability.

Once the sample wide average rate is allocated across all the individual European banks in our sample, we see a notable heterogeneity between countries in the gap between their current capitalization and the model-based prescription. This leads to the obvious question of whether these discrepancies are due to differences between country averages or to the way national regulators allocate their country average over individual banks.

To answer that question we next apply our optimization framework on an individual country basis, allowing for differences in average buffer size between countries and optimizing the within country allocation given the country specific buffer average. We show that the way individual countries allocate an average buffer size over individual banks under their remit is not the source of the Europe wide differences. So any differences between the allocation in the country-constrained and the Euro-wide optimization arise mostly from differences in the cross-country average, not in the way any average is allocated across individual banks in each country.

References

- Acharya, V. V. and Johnson, T. C. 2007. Insider trading in credit derivatives. *Journal of Financial Economics*, 84(1):110–141.
- Acharya, V. V., Pedersen, L. H., Philippon, T., and Richardson, M. 2017. Measuring systemic risk. *Review of Financial Studies*, 30(1):2–47.
- Acharya, V. V., Bergant, K., Crossignani, M., Eisert, T., and Mccan, F. 2022. The anatomy of the transmission of macroprudential policies. *Journal of Finance*, LXXVII (5):2533–2575.
- Adrian, T. and Brunnermeier, M. K. 2016. CoVaR. *American Economic Review*, 106(7): 1705–1741.
- Andersen, J., Leif, Sidenius and Basu, S. 2003. All your hedges in one basket. *Risk*, 16: 67–72.
- Augustin, P. and Izhakian, Y. 2020. Ambiguity, volatility, and credit risk. *The Review of Financial Studies*, 33(4):1618–1672.
- Baraukaitė, K., Nguyen, A. D., Rousová, L., and Cappiello, L. 2022. The impact of credit supply shocks in the euro area: market-based financing versus loans.

- BCBS. 2010. An assessment of the long-term economic impact of stronger capital and liquidity requirements. Technical report, Basel: Bank for International Settlements.
- Behn, M. and Schramm, A. 2021. The impact of G-SIB identification on bank lending: Evidence from syndicated loans. *Journal of Financial Stability*, 57:100930.
- Berndt, A., Duffie, D., and Zhu, Y. 2025. The decline of too big to fail. *American Economic Review*, 115(3):945–974.
- Billio, M., Getmansky, M., Lo, A. W., and Pelizzon, L. 2012. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535–559.
- Bolder, D. J. 2018, *Credit-Risk Modelling: Theoretical Foundations, Diagnostic Tools, Practical Examples, and Numerical Recipes in Python*. Springer. ISBN 978-3-319-94687-0.
- Brooke, M., Bush, O., Edwards, R., Ellis, J., Francis, B., Harimohan, R., Neiss, K., and Siegert, C. 2015. Measuring the macroeconomic costs and benefits of higher UK bank capital requirements. *Bank of England Financial Stability Paper*, 35(35):1–55.
- Brownlees, C. and Engle, R. F. 2017. SRISK: A conditional capital shortfall measure of systemic risk. *Review of Financial Studies*, 30(1):48–79.
- Cappelletti, G., Ponte Marques, A., Varraso, P., Budrys, Ž., and Peeters, J. 2019. Impact of higher capital buffers on banks’ lending and risk-taking: evidence from the euro area experiments. ECB working paper.
- Christensen, J. 2006. Joint estimation of default and recovery risk: A simulation study. Technical report.
- Cline, W. R. 2017, *The Right Balance for Banks: Theory and Evidence on Optimal Capital Requirements*. Princeton University Press. ISBN 9780881327229.
- Degryse, H., Mariathasan, M., and Tang, T. H. 2020. GSIB status and corporate lending: An international analysis. *Available at SSRN 3642702*.
- Dick-Nielsen, J., Gyntelberg, J., and Thimsen, C. 2022. The cost of capital for banks: Evidence from analyst earnings forecasts. *The Journal of Finance*, 77(5):2577–2611.
- Dimitrov, D. and van Wijnbergen, S. 2023. Quantifying systemic risk in the presence of unlisted banks: Application to the European banking sector. *De Nederlandsche Bank Working Paper*.
- Duffie, D. 1999. Credit swap valuation. *Financial Analysts Journal*, 55.
- Duffie, D. and Singleton, K. J. 1999. Modeling term structures of defaultable bonds. *Review of Financial Studies*, 12(4):687–720.
- Engle, R. 2018. Systemic risk 10 years later. *Annual Review of Financial Economics*, 10(1):125–152.
- ESRB. 2017. Final report on the use of structural macroprudential instruments in the EU. Policy report, European Systemic Risk Board.
- Favara, G., Ivanov, I., and Rezende, M. 2021. GSIB surcharges and bank lending: Evidence from US corporate loan data. *Journal of Financial Economics*, 142(3):1426–1443.
- Firestone, S., Lorenc, A., and Ranish, B. 2017. An empirical economic assessment of the costs and benefits of bank capital in the US.
- FRB. 2015. Calibrating the GSIB surcharge. Technical report, Federal Reserve Board.

- Garcia, L., Lewrick, U., and Sečnik, T. 2023. Window dressing and the designation of global systemically important banks. *Journal of Financial Services Research*, 64(2): 231–264.
- Gauthier, C., Lehar, A., and Souissi, M. 2012. Macroprudential capital requirements and systemic risk. *Journal of Financial Intermediation*, 21(4):594–618.
- Geiger, S., Heires, M., Krüger, U., Ludwig, J., and Vogel, U. 2022. Calibrating capital buffers for other systemically important institutions (O-SIIs) in germany—utilising the equal expected impact approach. Technical report, Deutsche Bundesbank Technical Paper.
- Gornicka, L. and van Wijnbergen, S. 2013. Financial Frictions and the Credit Transmission Channel: Capital Requirements and Bank Capital. Tinbergen Institute Discussion Papers 13-013/VI/DSF50, Tinbergen Institute.
- Huang, X., Zhou, H., and Zhu, H. 2009. A framework for assessing the systemic risk of major financial institutions. *Journal of Banking & Finance*, 33(11):2036–2049.
- Huang, X., Zhou, H., and Zhu, H. 2012. Systemic risk contributions. *Journal of Financial Services Research*, 42(1):55–83.
- Hull, J. C. 2018, *Risk Management and Financial Institutions*. John Wiley & Sons. ISBN 978-1-119-44811-2.
- Jakucionyte, E. and van Wijnbergen, S. 2018. Unclogging the credit channel: On the macroeconomics of banking frictions.
- Jiron, A., Passmore, W., and Werman, A. 2021. An empirical foundation for calibrating the G-SIB surcharge. *BIS Working Papers* 935.
- Juselius, M., Marques, A. P., and Tarashev, N. A. 2025. Banks’ regulatory risk tolerance. Technical report, Bank of Finland Research Discussion Papers.
- Kaserer, C. and Klein, C. 2019. Systemic risk in financial markets: How systemically important are insurers? *Journal of Risk & Insurance*, 86(3):729–759.
- Laeven, L. and Valencia, F. 2020. Systemic banking crises database ii. *IMF Economic Review*, 68:307–361.
- Lehar, A. 2005. Measuring systemic risk: A risk management approach. *Journal of Banking & Finance*, 29(10):2577–2603.
- McNeil, F. R., Alexander J. and Embrechts, P. 2005, *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press. ISBN 0-691-12255-5.
- Merton, R. C. 1974. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29(2):449–470.
- Miles, D., Yang, J., and Marcheggiano, G. 2013. Optimal bank capital. *The Economic Journal*, 123(567):1–37.
- Nagel, S. and Purnanandam, A. 2020. Banks’ risk dynamics and distance to default. *The Review of Financial Studies*, 33(6):2421–2467.
- Naylor, M., Corrias, R., and Welz, P. 2024. Banks’ window-dressing of the g-sib framework: causal evidence from a quantitative impact study.
- Pan, J. and Singleton, K. J. 2008. Default and recovery implicit in the term structure of sovereign CDS spreads. *Journal of Finance*, 63(5):2345–2384.
- Pascual, A. G., Avesani, M. R. G., and Li, M. J. 2006. A new risk indicator and stress testing tool: A multifactor nth-to-default CDS basket. IMF Working Papers 2006/105,

- International Monetary Fund.
- Passmore, W. and von Hafften, A. H. 2019. Are Basel’s capital surcharges for global systemically important banks too small? *International Journal of Central Banking*.
- Pelger, M. 2020. Understanding systematic risk: a high-frequency approach. *The Journal of Finance*, 75(4):2179–2220.
- Puzanova, N. and Düllmann, K. 2013. Systemic risk contributions: A credit portfolio approach. *Journal of Banking & Finance*, 37(4):1243–1257.
- Reinhart, C. M. and Rogoff, K. S. 2009. The aftermath of financial crises. *American Economic Review*, 99(2):466–72.
- Reverdy, P. and Leonard, N. E. 2015. Parameter estimation in softmax decision-making models with linear objective functions. *IEEE Transactions on Automation Science and Engineering*, 13(1):54–67.
- Romer, C. D. and Romer, D. H. 2017. New evidence on the aftermath of financial crises in advanced countries. *American Economic Review*, 107(10):3072–3118.
- Siriwardane, E. N. 2019. Limited investment capital and credit spreads. *The Journal of Finance*, 74(5):2303–2347.
- Tarashev, N. A. and Zhu, H. 2006. The pricing of portfolio credit risk. BIS Working Papers 214, Bank for International Settlements.
- Vasicek, O. 1987. Probability of loss on loan portfolio. Working papers, KVM Corporation.

A Data Annex

A.1 Sample selection

Our data sample consist of the European banks designated in 2022 as systemically important by the European Banking Authority.¹⁵ The list consists of 30 banks, out of which we drop four with insufficient CDS data: BPCE (France) accounting for 6.8% of the systemic exposure calculated by EBA; Banque Postale (France); DNB ASA (Denmark); Nykredit (Norway), each accounting for less than 1.5% of the systemic exposure calculated by EBA. We add the Netherlands' Volksbank (VB), designated as a D-SIB in the Netherlands, at the request of the Dutch Central Bank. VB is too small to have a material influence on the overall analysis; adding it, however, allows us to evaluate its optimal buffers as well. Table (4) provides the details.

We collect weekly prices for ISDA2014-compliant CDS contracts with a fixed five-year maturity from Bloomberg. Annual end-of-year balance sheet liability size figures and current CET1 capitalization ratios are retrieved from FactSet, BankFocus and from publicly available financial statements whenever the data providers have a gap. Annual balance sheet numbers are interpolated to weekly values using a cubic spline to avoid jumps at fiscal year-end due to accounting conventions. Bank-specific P2R is retrieved from supervisory sources.¹⁶ In line with the rest of the buffers in this study, we only consider the CET1 portion of the P2R framework.

Approximately one-third of the banks in our sample are not publicly listed. This group includes cooperative institutions such as France's CRMU, Germany's DZ, BAY, LBBW, and HESL; the Dutch state-owned VB; as well as privately held banks like RABO and INGB. While ABN was partially re-listed in 2015 following earlier government intervention, only a minority stake is currently traded on public markets. The use of CDS rates enables us to incorporate these institutions into the analysis, thereby offering a more comprehensive view of the financial system. Importantly, we rely on the CDS rate for

¹⁵See <https://www.eba.europa.eu/risk-and-data-analysis/risk-analysis/risk-monitoring/global-systemically-important-institutions-g-siis>. O-SII buffers for 2022 are available in EBA's interactive dashboard: <https://www.eba.europa.eu/risk-and-data-analysis/risk-analysis/risk-monitoring/other-systemically-important-institutions-o-siis>

¹⁶See <https://www.bankingsupervision.europa.eu/banking/srep/html/p2r.en.html>

ING Bank (INGB), a subsidiary of the publicly traded ING Group, due to its relevance for European regulatory oversight. Since ING Group does not carry a formal guarantee for ING Bank’s liabilities, the availability of this CDS contract allows for a more precise assessment of the risks associated with the bank’s European operations. For several banks (ERST, DZ, BAY, LBBW and HESLN) only senior debt CDSs are available. To ensure that the default probabilities inferred from these banks are on the same level-playing field as those derived for the rest of the sample, we add to their CDS prices the average spread between senior and subordinated CDS prices for core area banks in Europe.

Table 4: Data Description


Short Code	Country	Bank Full Name	Public	CDS Type	G-SIB	D-SIB
ERST	Austria	Erste Group	Y	SR	N	Y
KBCB	Belgium	KBC	Y	SUB	N	Y
DANK	Denmark	Danske Bank	Y	SUB	N	Y
NORD	Finland	Nordea	Y	SUB	N	Y
BNP	France	BNP Paribas	Y	SUB	Y	Y
CRAG	France	Credit Agricole	Y	SUB	Y	Y
CRMU	France	Credit Mutuel	N	SUB	N	Y
SOCG	France	Societe Generale	Y	SUB	Y	Y
COMZ	Germany	Commerzbank	Y	SUB	N	Y
DB	Germany	Deutsche Bank	Y	SUB	Y	Y
DZ	Germany	DZ Bank	N	SR	N	Y
BAY	Germany	Bayern LB	N	SR	N	Y
LBBW	Germany	LBBW	N	SR	N	Y
HESLN	Germany	Helaba	N	SR	N	Y
INTE	Italy	Intesa Sanpaolo	Y	SUB	N	Y
UNIC	Italy	Unicredit	Y	SUB	Y	Y
RABO	Netherlands	Rabobank	N	SUB	N	Y
ABN	Netherlands	ABN Amro	Y	SUB	N	Y
INGB	Netherlands	ING Bank	N	SUB	Y	Y
VB	Netherlands	Volksbank	N	SUB	N	Y
CAIX	Spain	Caixabank	Y	SUB	N	Y
SAB	Spain	Sabadell	Y	SUB	N	Y
SANT	Spain	Santander	Y	SUB	Y	Y
BBVA	Spain	BBVA	Y	SUB	Y	Y
SWEN	Sweden	Handelsbanken	Y	SUB	N	Y
SEB	Sweden	Skandinaviska Enskilda Banken	Y	SUB	N	Y
SWED	Sweden	Swedbank	Y	SUB	N	Y

Note. This table summarizes the key characteristics of the dataset: the country of origin for each bank, the short code used throughout the paper, the status of the bank’s equity—whether it is publicly traded (‘Y’) or privately held (‘N’)—and the type of CDS spreads employed in the analysis, either on senior debt (SR) or subordinate debt (SUB). For senior debt CDSs, we adjust the spreads by the median difference between senior and subordinate debt to ensure comparability.

The data encompasses several major tail events for the European economy - the initial Covid shock as of the beginning of 2020 and the first lockdowns, the start of the war in Ukraine with fears of gas shortages as of the beginning of 2022, and the inflation spikes and interest rate tightening by the Fed and the ECB. Figure 12 shows the evolution of

the CDS spreads in our sample for the evaluation period.

Figure 12: CDS Prices (bps)




(a) Median Rates per Country

Throughout, we use an LGD assumption of 80% (i.e. RR of 20%) in line with the prospectus provided in Bloomberg for these CDS on subordinated debt.

We use the time period August 2019 to August 2022. A three-factor latent model is then estimated based on weekly CDS data. The use of weekly data allows us to reduce the estimation noise from daily trading and calibrate the model on longer-term spread co-movements. The number of factors were chosen from the scree plot of the principal components of the data (see Figure 13) indicating that that three latent factors capture more than 80% of the co-variation in the movements of the CDS spreads.

Figure 13: Share of Total Explained Variance



Note. This figure shows the cumulative share of explained variance (the solid curve) and the incremental share (dashed curve) over the CDS price log-changes in our sample.

Table 5 then summarizes the full input data into the optimization model, including

the data estimated from CDS prices and correlations.

Table 5:

Bank	Weights	O-SII (%)	P2R (%)	RWA Int.	CDS (bps)	(%)	(%)		ρ_1	ρ_2	ρ_3
ABN	1.99	1.50	1.13	30.38	104.46	1.100	0.165	2.423	0.731	0.006	-0.265
BAY	1.34	0.50	1.13	21.93	163.10	1.789	0.152	2.070	0.910	-0.078	0.022
BBVA	3.22	0.75	0.84	46.04	230.76	2.557	0.134	3.330	0.937	0.168	-0.018
BNP	13.24	1.50	0.74	24.82	163.10	1.823	0.117	1.786	0.957	0.201	0.058
CAIX	3.38	0.38	0.93	31.05	225.64	2.410	0.225	2.341	0.193	-0.089	-0.506
COMZ	2.33	1.25	1.13	31.94	317.91	3.474	0.142	2.702	0.949	0.170	-0.015
CRAG	10.51	1.00	0.84	23.26	156.92	1.753	0.117	1.529	0.949	0.224	0.074
CRMU	4.15	0.50	0.98	29.58	206.82	2.003	0.426	3.008	0.510	0.081	-0.050
DANK	2.66	3.00	1.01	19.70	266.43	2.938	0.137	2.098	0.952	0.103	0.099
DB	6.64	2.00	1.41	24.95	328.06	3.576	0.145	2.123	0.916	0.145	-0.084
DZ	3.14	1.00	0.96	22.09	148.81	1.513	0.265	1.810	0.852	0.001	0.096
ERST	1.51	1.00	0.98	41.12	178.66	1.946	0.170	3.194	0.925	0.017	0.030
HESLN	1.07	0.50	0.98	29.46	168.19	1.897	0.101	2.302	0.916	-0.072	0.084
INGB	4.71	2.50	0.98	31.93	70.71	0.699	0.166	2.311	0.747	-0.072	0.101
INTE	5.28	0.75	1.01	31.51	323.84	3.523	0.155	2.750	0.913	0.142	0.053
KBCB	1.67	1.50	1.05	32.00	214.03	1.943	0.565	2.687	0.162	0.136	-0.142
LBBW	1.41	0.75	1.03	29.54	150.82	1.630	0.170	2.285	0.905	-0.034	0.081
NORD	2.82	2.00	0.98	24.06	131.16	1.331	0.244	2.099	0.598	-0.696	0.192
RABO	3.15	2.00	1.07	36.86	157.35	1.739	0.136	3.351	0.949	0.147	0.068
SAB	1.25	0.25	1.21	31.59	365.34	3.460	0.640	2.426	0.310	-0.099	-0.645
SANT	7.87	1.00	0.84	33.70	214.60	2.379	0.136	2.344	0.956	0.162	-0.003
SEB	1.59	1.00	1.01	20.25	139.54	1.428	0.244	2.077	0.636	-0.713	0.031
SOCG	7.32	1.00	1.19	22.93	192.76	2.159	0.114	1.819	0.937	0.179	0.064
SWED	1.32	1.00	1.01	24.96	164.63	1.706	0.251	2.429	0.653	-0.378	-0.253
SWEN	1.61	1.00	1.01	22.17	133.98	1.348	0.260	2.202	0.674	-0.625	0.063
UNIC	4.49	1.00	0.98	32.22	362.50	3.913	0.159	3.076	0.917	0.123	0.022
VB	0.34	1.00	1.69	20.21	147.71	1.206	0.559	2.292	0.650	0.099	-0.201

Note: Full sample of 27 European systemic banks. Columns include the liability weight in the sample (w), observed macroprudential buffers (O-SII), Pillar 2 requirements (P2R), observed CDS spreads (bps), implied default probabilities (pure credit-based PD_c), liquidity component of the initial default probability (PD_l), asset volatility (σ), and factor loadings (ρ_1, ρ_2, ρ_3) estimated from CDS data. The model is calibrated on the pure credit-based PD, which corrects for CDS illiquidity premia, the size of the correction (in default probability terms) is indicated by PD_l .

Table 6:

Bank	Country	Welfare Optimal $\bar{k}$			O-SII Optional $\bar{k}$	
		Mild	Moderate	Severe	Euro Perspective	Local Perspective
Austria	ERST	18.86	19.51	20.39	1.59	–
Belgium	KBCB	8.89	9.89	9.08	0.00	–
Denmark	DANK	24.25	25.90	27.10	2.37	–
Finland	NORD	16.77	18.07	19.20	0.03	–
France	BNP	17.38	18.93	19.52	1.58	1.36
France	CRAG	16.23	17.17	18.17	0.88	0.97
France	CRMU	13.50	12.72	17.97	0.00	0.00
France	SOCG	18.83	20.32	21.46	1.57	1.57
Germany	BAY	21.50	22.83	24.54	2.53	0.00
Germany	COMZ	20.27	21.41	22.61	1.78	1.93
Germany	DB	20.47	21.53	22.50	1.73	1.42
Germany	DZ	18.79	20.01	21.37	1.16	1.83
Germany	HESLN	18.57	19.88	21.08	1.78	1.13
Germany	LBBW	18.38	19.72	20.77	2.00	0.49
Italy	INTE	20.69	21.94	23.43	2.26	1.02
Italy	UNIC	21.63	22.49	24.80	2.31	0.68
Netherlands	ABN	17.19	19.57	19.52	0.01	2.16
Netherlands	INGB	17.21	17.63	18.68	0.58	1.64
Netherlands	RABO	21.80	22.99	24.09	2.21	2.97
Netherlands	VB	8.76	8.72	8.69	0.00	0.00
Spain	BBVA	17.50	18.38	19.66	1.81	0.70
Spain	CAIX	8.13	8.08	8.26	0.00	1.43
Spain	SAB	10.59	12.51	8.22	0.00	0.00
Spain	SANT	17.07	18.22	19.05	1.23	0.60
Sweden	SEB	8.02	8.02	8.02	0.00	0.62
Sweden	SWED	8.05	8.01	8.01	0.01	1.56
Sweden	SWEN	19.55	21.31	21.08	0.00	0.92

Note. This table reports optimization model outputs for systemic banks under three systemic risk scenarios (Mild, Moderate, Severe) and two macroprudential perspectives: Euro area-wide and local. Values represent optimal additional capital buffer levels derived from the model.

A.2 Regulatory Capital Requirements

Here we provide a short overview of the different regulatory capital requirements as summarized in Table 7¹⁷

Minimum Capital Requirement (Pillar 1 and Pillar 2):

- Pillar 1: (i) Common equity Tier 1 capital (CET1) has to be at least 4.5% of risk-

¹⁷For a general discussion see Hull (2018); and for details on the latest implementations and regulatory debates see EBA's guidance on the O-SII framework; ESRB's systemic reports; BIS's guidance on the G-SIB framework.

weighted assets (RWA); (ii) Total Tier 1 capital at least 6% of RWA; (iii) Total (Tier 1 and Tier 2) capital of at least 8% of RWA; (iv) Leverage Ratio (Tier 1 Capital/Total Exposure) at least 3%. Total exposure includes on- and off-balance sheet exposure, derivatives exposure, and securities financing transaction exposures. No risk-weighting is applied.

- Pillar 2 (P2R): Bank-specific microprudential capital requirement that aims to cover risks that are not (fully) covered by the Pillar 1. Partially satisfied with CET1 capital. Pillar 2 Guidance (P2G) is not legally binding but may have implications on the distributional capacity to shareholders.

Combined Buffer Requirement (CBR)

- Capital Conservation Buffer (CCB) (Basel III): To be maintained in normal times. If levels fall below requirements, banks restrain dividends and bonus payments until capital has been replenished. CET1 add-on of 2.5% of RWA.
- Countercyclical Capital Buffer (CCyB): Applied country-wide, and similar to CCB but at the discretion of national authorities. It is a CET1 add-on of between 0% and 2.5% of the bank's domestic RWAs.
- Systemic risk buffer (SyBR): Additive to other buffers. Designed to address risk spillover from the economy (from the system) to individual banks. At the discretion of national authorities aiming to address risks that are not covered by the CCyB or the G-SII/O-SII buffers. It may apply to all banks, particular individual banks, and across a subset of exposures (e.g. on the residential exposure of the RWAs as in e.g. Belgium, Germany, etc.).
- Global Systemically Important Institution (G-SII) buffers on Globally Systemically Important Banks (G-SIBs): Part of the combined buffer requirement. Typically ranging between 0% and 2.5%. Designed to address negative spillovers from individual banks to the global economy. Framework and assessment methodology set by the Basel Committee on Banking Supervision (BCBS) and applies to banks glob-

ally (ranking in categories based on size, complexity, cross-jurisdictional activity, interconnectedness, substitutability of activities) Enforced by national authorities.

- **Other Systemically Important Institution (O-SII) buffers:** Part of the combined buffer requirement. Typically ranging between 0% and 3%, where the maximum of G-SII and O-SII applies. National authorities have the discretion on the size of the buffer surcharge. Designed to address negative spillovers from individual banks to the national economy. Guidelines set by European Banking Authority (EBA) (ranking in categories based on size, importance, complexity, and interconnectedness). Measurements and enforcement by national authorities on a “comply or explain” basis.

Table 7: Overview of Capital Buffers under the EU Regulatory Framework

Buffer	CET1 Requirement	Scope	Authority
CCB	2.5%	All banks	Basel III / National
CCyB	0–2.5%	Country-wide	National
SyRB	Variable	All or specific banks	National
G-SII Buffer	0–2.5%	G-SIBs	BCBS / National
O-SII Buffer	0–3%	O-SIIs	EBA / National

A.3 Stale price analysis

Let $CDS_{i,t}$ denote the weekly CDS level of bank i at week t . Define the week-over-week change $\Delta CDS_{i,t} = CDS_{i,t} - CDS_{i,t-1}$. For a tolerance parameter $\tau \geq 0$, we say week t is unchanged for bank i if $|\Delta CDS_{i,t}| \leq \tau$. We report the following summary statistics per bank:

1. **Fraction of unchanged weeks:** $\text{UnchgFrac}_i(\tau) = \frac{\#\{t : |\Delta CDS_{i,t}| \leq \tau\}}{T_i - 1}$, where T_i is the number of weeks available for bank i .
2. **Longest plateau (weeks):** the maximum length of a consecutive sequence of weeks for which $CDS_{i,t}$ is exactly constant (i.e., $\tau = 0$).
3. **Mean consecutive unchanged length:** the average run length of consecutive unchanged weeks, computed over all runs $\{t : |\Delta CDS_{i,t}| \leq \tau\}$.

4. **Share of weeks in plateaus ≥ 2 :** the fraction of total weeks that are members of a plateau of length at least two weeks.

The metric is model-free and aims to detect vendor-quote staleness or illiquidity. We present results for the exact-equality case ($\tau = 0$) and for a small tolerance ($\tau = 0.01$), which helps account for rounding or micro-jitter in some feeds.

Table 8: Stale-price metrics (weekly CDS) — tolerance $\tau = 0.01$ (alphabetical by bank)

Bank	n	pct	mean consec	max plateau
ABN	15	9.62%	1.15	3
BAY	35	22.44%	2.50	6
BBVA	0	0.00%	0.00	1
BNP	0	0.00%	0.00	1
CAIX	12	7.69%	1.71	6
COMZ	0	0.00%	0.00	1
CRAG	0	0.00%	0.00	1
CRMU	1	0.64%	1.00	2
DANK	0	0.00%	0.00	1
DB	0	0.00%	0.00	1
DZ	8	5.13%	1.14	3
ERST	1	0.64%	1.00	2
HESLN	32	20.51%	2.46	6
INGB	0	0.00%	0.00	1
INTE	0	0.00%	0.00	1
KBCB	3	1.92%	1.50	3
LBBW	14	8.97%	2.00	7
NORD	8	5.13%	1.60	3
RABO	2	1.28%	1.00	2
SAB	2	1.28%	1.00	2
SANT	0	0.00%	0.00	1
SEB	3	1.92%	1.00	2
SOCG	0	0.00%	0.00	1
SWED	5	3.21%	1.00	2
SWEN	5	3.21%	1.25	3
UNIC	0	0.00%	0.00	1
VB	0	0.00%	0.00	1

Table 9: Summary Statistics

	N	Mean	Std.Dev	Median	Min	Max
lnCDS	19298	4.893	0.893	4.967	1.705	7.377
BA	19199	0.125	0.104	0.095	0.002	0.993
CET1	19298	13.342	4.175	12.815	5.339	37.100
LIQ	13500	152.180	47.872	141.000	95.000	455.000
NI_TA	19298	0.003	0.003	0.003	-0.018	0.013
DtoL	19298	0.449	0.141	0.430	0.183	0.840

Note. This table summarizes the data used for the diagnostic panel regression.

A.4 Testing for Sovereign (Too-Big-To-Fail) Effects

A.5 Summary Statistics: Liquidity regression

A.6 O-SII buffer buckets: Optimal vs. Model

Table 10: Binned histograms (0.5 pp bins) — Euro-wide model vs. O-SII buffer

Bin ([0, 3.5))	Model: Count	Model: Share (%)	O-SII: Count	O-SII: Share (%)
[0.0, 0.5)	10	37.0	2	7.4
[0.5, 1.0)	2	7.4	6	22.2
[1.0, 1.5)	3	11.1	11	40.7
[1.5, 2.0)	8	29.6	3	11.1
[2.0, 2.5)	2	7.4	3	11.1
[2.5, 3.0)	2	7.4	1	3.7
[3.0, 3.5)	0	0.0	1	3.7
Total (N)	27	100.0	27	100.0

Notes: Bins are left-closed and right-open. Shares are within-series and may not sum to 100% due to rounding. The table uses banks with non-missing values in both series (Euro-wide sample).

A.7 Robustness to the choice of factor model estimation window

A.8 Equity-based vs. CDS-based target correlations


A natural question is whether our findings remain robust when systemic risk is assessed using the more conventional equity-based approach. Prior work, such as Tarashev and Zhu (2006), has raised concerns that equity and CDS markets may diverge in their assessment of asset dependencies, particularly during periods of market stress, while Siriwardane (2019) has pointed out the CDS prices a subject to intermediation shocks. To

Table 11: Binned histograms (0.5 pp bins) — Local model vs. O-SII buffer

Bin ($[0, 3.5)$)	Local Model: Count	Local Model: Share (%)	O-SII: Count	O-SII: Share (%)
$[0.0, 0.5)$	5	21.7	2	8.7
$[0.5, 1.0)$	7	30.4	6	26.1
$[1.0, 1.5)$	2	8.7	10	43.5
$[1.5, 2.0)$	7	30.4	2	8.7
$[2.0, 2.5)$	1	4.3	2	8.7
$[2.5, 3.0)$	1	4.3	1	4.3
$[3.0, 3.5)$	0	0.0	0	0.0
Total (N)	23	100.0	23	100.0

Notes: Bins are left-closed and right-open. Shares are within-series and may not sum to 100% due to rounding. The table uses banks with non-missing values in both columns and ignores NaN entries (local comparison).

Figure 14: 3-Year Rolling Window Principal Component Analysis




Note. This figures shows the share of explained co-variation in our CDS sample based on a three year rolling window analysis.

investigate whether segmentation or intermediation frictions in the CDS market might distort our results, we isolate the subset of banks for which both CDS and equity data are available and compare the resulting MES estimates. Specifically, we re-estimate the bank target correlations—originally derived from CDS co-movements (see equations (11) and (12))—using equity return correlations instead.

Figure 16 presents a comparison of MES values obtained under both specifications. In panel (b) we present also the weighted version of the MES as a share of the total Expected Systemic Shortfall of the sector, i.e. its Percentage Contribution to Expected Shortfall. The results show a strong alignment, with a correlation exceeding 80% for the MES and exceeding 90% after weighting it measure, suggesting that the CDS-based approach yields systemic risk estimates broadly consistent with those derived from equity data. We therefore proceed with the CDS-based correlations, which offer the added advantage of


Figure 15: Macroprudential Welfare Function with Different Correlation Calibrations



Note. This set of figures shows the calibration of the social disutility function on recent correlations and on correlations based on data extending back to 2008.

incorporating non-listed but systemically relevant banks into the analysis.

Figure 16: Equity-based vs CDS-Based Risk Attribution



Note. This set of figures shows comparisons in risk attribution between the credit-based approach estimated with CDS-based correlations, and with equity-based correlations (for the listed banks only).

B Analytical Annex

B.1 Inferring PDs from observed CDSs

We infer the banks' default probabilities from single-name CDS close prices using the approach outlined by Duffie (1999). It is based on the simplifying assumption that recovery rates (RR) are known and constant over the horizon of the contract.¹⁸

With this in mind, we can proceed with identifying the equation for pricing a CDS contract. By market convention, at the initiation date t of the contract the spread CDS_t is set to ensure that the value of the protection leg and the premium leg of the contract are equal, such that the contract has a zero value:

$$\underbrace{CDS_t \int_t^{T_{cds}} e^{-r_\tau \tau} \Gamma_\tau d\tau}_{\text{PV of CDS premia}} = \underbrace{(1 - RR_t) \int_t^{T_{cds}} e^{-r_\tau \tau} q_\tau d\tau}_{\text{PV of protection payment}} \quad (21)$$

where T_{cds} is the maturity date of the CDS contract, $\tau > t$ is future time after the initiation of the contract, r_τ is the annualized instantaneous risk-free rate, CDS_t is the observed CDS spread for the day, q_τ is the annualized instantaneous risk-neutral default probability, $\Gamma_\tau = 1 - \int_t^\tau q_s ds$ is the risk-neutral survival probability until time τ , and RR_t is the expected recovery rate in case of default, assumed to be constant over time.

For simplicity, we assume that the yield and the probability default curves are flat over the lifetime of the CDS contract once the contract is established. Then, we can set $r_\tau = r_t$ and $q_\tau = q_t$ over the lifetime of the contract initiated at time t . Then the default probability q at time t follows from equation (21):

$$q_t = \frac{a CDS_t}{a(1 - RR_t) + b CDS_t} \quad (22)$$

with $a = \int_t^{T_{cds}} e^{-r\tau} d\tau$ and $b = \int_t^{T_{cds}} \tau e^{-r\tau} d\tau$. Setting $T_{cds} - t = 5$ to capture 5-year CDS contracts, we can imply the annualized default probabilities.

¹⁸We do not try to identify expected recovery rates separately from the observed CDS data. There are alternative and more sophisticated approaches (cf Pan and Singleton (2008); Christensen (2006); Acharya and Johnson (2007); Duffie and Singleton (1999)). However, given the identification challenge between PDs and RRs, the simplifying assumption we employ in estimation is widely used in the literature and is difficult to improve.

B.2 Latent Factor Model Estimation

On a given day we don't observe the market value of a bank's assets, but we observe how far it is from the default threshold. This is implied by the default probability associated with the market price of a CDS contract traded, such that

$$PD_{i,t} = \mathbb{P}(V_{i,T} \leq D_i) = \mathbb{P}(U_i \leq -DD_{i,t}) = \Phi(-DD_{i,t})$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution, and t are periodic observations of the default probability, in our case with a weekly frequency. This implies the DD measure at the end of the trading week

$$DD_{i,t} = -\Phi^{-1}(PD_{i,t}) \quad (23)$$

Observed changes in the default probability can then be linked to the changes in the asset value by first-differencing (23) and relating it to Merton's DD:

$$\Delta\Phi^{-1}(-PD_{i,t}) = \Delta DD_{i,t} = \frac{\ln V_{i,t} - \ln V_{i,t-1}}{\sigma_i \sqrt{T-t}} \quad (24)$$

which is defined as

$$DD_{i,t} = \frac{\ln \frac{V_{i,t}}{D_i} + \left(r - \frac{\sigma_i^2}{2}\right) T}{\sigma_i \sqrt{T}} \quad (25)$$

More importantly, this allows us to infer the correlation structure between the latent default variables. For example, the correlation between banks i and j can be written as

$$\text{Corr}(U_{i,t}, U_{j,t}) = \text{Corr}(\Delta DD_{i,t}, \Delta DD_{j,t}) \equiv a_{i,j} \quad (26)$$

Based on these observed correlations, we can construct the target correlation matrix

towards which to fit the latent factor model of (8) as

$$\Sigma \equiv \begin{bmatrix} 1 & a_{12} & a_{13} & \dots & a_{1N} \\ a_{21} & 1 & a_{23} & \dots & a_{2N} \\ \dots & \dots & \dots & \dots & \dots \\ a_{N1} & a_{N2} & a_{N3} & \dots & 1 \end{bmatrix} \quad (27)$$

B.3 Softmax Transformation for the Optimization Problem

We have the optimization problem

$$\begin{aligned} \min_{k_1^{ma}, \dots, k_n^{ma}} f(k_1, k_2, \dots, k_n) \\ s.t. \quad \sum_{i=1}^n w_i k_i^{ma} = \bar{k}; \quad k_i \geq 0, i = 1, \dots, n \end{aligned} \quad (28)$$

where the objective function follows from the underlying risk model: $f(k_1, k_2, \dots, k_n) = \mathbb{E}(L_{sys} \mid L_{sys} > \bar{L}) \mathbb{P}(L_{sys} > \bar{L})$ with the components of the function, defined in (15) and (17).

Introduce an auxiliary variable $y_i \in \mathbb{R}$ for each $k_i, i \in n$. Then we use a softmax transformation i to compute the weight out of the budget $\bar{k}$ associated with each auxiliary variable:

$$p(y_i) = \frac{\exp(y_i)}{\sum_j \exp(y_j)}$$

This transformation preserves differentiability, which may be important for gradient-based methods.

Then we have the mapping from k_i to y_i as

$$k_i \equiv \bar{k} \frac{p(y_i)}{\sum_j w_j p(y_j)}$$


By construction this set-up implies that budget constraint holds with $\sum_i w_i k_i = \bar{k}$; and the non-negativity constraints hold with $k_i^{ma} \geq 0$ for each $y_i \in R$ for each i . In that case, we can transform the original constrained optimization over $(k_1, \dots, k_n)$ to an

unconstrained optimization over $(y_1, \dots, y_n)$ as

$$\min_{y_1^{ma}, \dots, y_n^{ma}} f(k_i(y)), \quad y_i \in \mathbb{R}$$

The Softmax transformation then acts as a smooth mapping from $\mathbb{R}$ to the positive simplex. For further discussion of the approach see Reverdy and Leonard (2015).

Figure 17: Objective Function Surface



Note. These figure shows the objective function in the domain of the softmax transformed variables y_1 and y_2 and in the domain of the macroprudential capitalization of bank 1 along the feasible line. The feasible capitalization of bank 2 follows from constraint $k_2 = 2\bar{k} - k_1$. The results are generated by the baseline case with two symmetric banks.