Uncertain credit-loss phases and bank capital*

Mikael Juselius[a], Nikola Tarashev[b]

[a] *Bank of Finland, Snellmaninaukio, PL 160, 00101 Helsinki, email: mikael.juselius@bof.fi*

[b] **Corresponding author.** *Bank for International Settlements, Centralbahnplatz 2, 4051 Basel, Switzerland, email: nikola.tarashev@bis.org*

Abstract

To account for difficult-to-forecast turning points in credit risk, we incorporate uncertainty about transitions between low- and high-loss phases in a well-known portfolio credit-risk model. We then study how the adequacy of bank capital depends on two macro risk factors, driving respectively default clustering within a phase and phase switches. For a bank with *lower* exposure to the within-phase macro risk, ignoring phase uncertainty is more detrimental and improving phase-switch forecasts brings larger solvency benefits. To investigate the practical relevance of our results, we design a novel empirical method for comparing the degree of within-phase macro-risk exposure across credit portfolios.

JEL Codes: G21; G28; G32

Keywords: Expected loss provisioning; Bank capital; Unexpected losses; Credit cycles; Portfolio credit risk

*This paper reflects the views of its authors, not necessarily the views of the Bank of Finland or the Bank for International Settlements. We would like to thank Michael Gordy for several insightful discussions as well as Stijn Claessens, Esa Jokivuolle, Hyun Song Shin and seminar participants at the Bank of Finland, the BIS, the International Risk Management Conference, 2021, a workshop at Her Majesty's Treasury, the United Kingdom, a research cluster conference of the European System of Central Banks and the Bank of England for valuable suggestions, and Anamaria Illes for excellent help with the data. Declaration of competing interest: none.

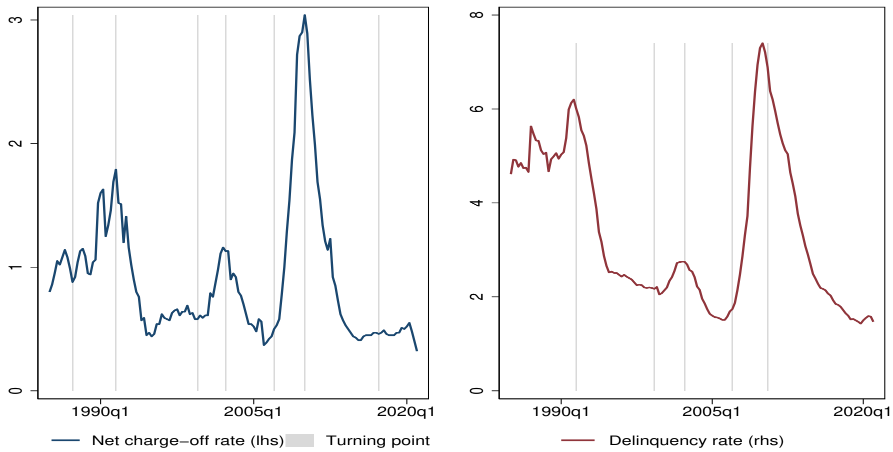

Figure 1: ***Abrupt turning points and credit-loss phases.*** Quarterly net charge-off and delinquency rates on loans by US banks to the non-financial private sectors. Turning points are identified with the Harding and Pagan (2002) business cycle dating algorithm, applied to the 3-year ahead change in each series (window = 2Y, phase length > 2Y and cycle length > 3Y).

1 Introduction

Turning points in credit risk and the distinct loss phases that they generate have had important financial stability implications. Abrupt transitions between low- and high-loss phases transpire from delinquency and loan charge-off rates (*Figure 1*) as well as from defaults on bonds issued by non-financial firms (Krüger et al. (2018)). On the back of such transitions, the old "incurred loss" provisioning framework systematically failed to support banks' resilience and had to be eventually overhauled (FSF (2009)).

Even though credit-risk forecasts are counted on for the timely accumulation of loss-absorbing resources that protect lenders' solvency – not least by the new provisioning framework (IASB (2014) and FASB (2016)) – there is an ongoing debate whether they can live up to the task in practice. Covas and Nelson (2018), Abad and Suárez (2017), Chae et al. (2018), Krüger et al. (2018), Goncharenko and Rauf (2020), and Loudis and Ranish (2019) are sceptical about the possibility to forecast credit-risk turning pints in real time (i.e. out of sample). And while Harris et al. (2018), Lu and Nikolaev (2022) and Juselius and Tarashev (2020) seek to improve the accuracy of real-time forecasts of such turning points – respectively on the basis of cross-sectional analysis, a high-dimensional dynamic factor model and indicators of the build-up of financial imbalances – much uncertainty remains.

Wary of the low precision of forecasts around turning points in credit risk, prudential authorities have increasingly de-emphasised point estimates as drivers of regulation. For instance, European authorities have enshrined in law the need to account for estimation error through conservative capital requirements (European Union (2013), Article 179). In another expression of conservatism, the US Federal Reserve assesses large US banks' regulatory capital via stress tests that do not refer to the likelihood of the underlying scenarios (e.g. Federal Reserve Board (2022)). Recently, the Basel Committee on Banking Supervision (BCBS) has supported the build-up of "cycle neutral" buffers that would be in place as long as there is no systemic crisis, thus abstracting from financial-risk estimates outside such an event (BCBS (2022); see also Bank of England (2016)). Yet, these and similar other measures to handle uncertainty are costly. Hence, there is value in studying under which conditions they provide particularly high financial-stability benefits.

With this as a motivation, we investigate bank capital adequacy under uncertainty about the evolution of credit risk. Building on a parsimonious setup, in which a bank sets capital to target the same solvency objective as its regulator, we introduce a persistent risk factor to model uncertainty about turning points in probabilities of default (PDs). We then study how benefits of accounting for or reducing such uncertainty depend on exposure to a second, short-lived macro risk factor that drives default clustering for given PDs. Finally, we derive an empirical method for comparing portfolios in terms of their exposure to the default-clustering factor.

Throughout, we focus on two aspects of the probability distribution of credit-loss rates. One is the expected loss (EL), which determines banks' new forward-looking provisioning and thus drives their loan-loss reserves (IASB (2014) and FASB (2016)). In a portfolio of exposures with homogeneous risk parameters – which we assume throughout the paper, in order to focus on the *evolution* of credit risk – EL is the product of loss-given-default (LGD) and probability of default (PD). The other aspect is the unexpected loss (UL) – or the difference between some high percentile of the default-rate distribution and EL – which drives banks' regulatory capital requirements (BCBS (2017)). Rising uncertainty as to whether the turning point is around the corner – e.g. an increasing similarity of the perceived likelihoods of increasingly distinct phases – need not change EL but is sure to raise UL all else the same. Below, when we discuss such a decoupling between EL and UL, we will think of it as driven by a change in the general credit outlook, rather than a change in a portfolio's composition.

We build our analysis on a well-known model of portfolio credit risk (Vasicek (1991) and Gordy (2003)), which provides the backbone of banks' internal ratings-based regulatory capital requirements (Basel III: BCBS (2017), pp 62-3). This model assumes that the relevant phase is known, i.e. that banks estimate perfectly the PD and LGD – and thus the EL – of the borrowers in their portfolio. It also assumes that any deviation of default rates from ELs stems from a transient macro risk factor that affects simultaneously the creditworthiness of all borrowers, i.e. a "default-clustering" factor. All else equal, a bank with a higher exposure to (or "loading" on) this factor faces a higher UL. For a given loading on the default-clustering factor, the model implies that PD and LGD are the only drivers of *any* change in the probability distribution of credit losses and, in particular, that EL and UL do not decouple.[1]

We extend this model by introducing a second macro risk factor that follows a two-state Markov switching process – driving transitions between low- and high-PD phases. Admittedly, such a factor cannot account fully for the richness of the time series properties of credit losses, e.g. it does not distinguish between turning points and phase switches and leaves all loss variation within a phase to be explained by the default-clustering factor. Importantly, however, it is designed to capture abrupt transitions between extreme loss rates – whereby recent incurred losses are uninformative about future ones around a phase switch – and implies that loss rates are persistent within a phase, similar to when a deterioration in the macro environment reduces borrowers' credit quality over several periods. We find that the Markov process describes well data on US loan charge-off rates from 1985 to 2021. Again, all else equal, a bank with a higher exposure to the macro risk factor within a phase faces a higher UL. However, changes to the likelihood of each phase and the magnitude of phase-contingent PDs can now lead UL to decouple from EL. Since time variation in LGD does not generate decoupling, we keep this parameter fixed in the background.

We then cast three banks that differ with respect to their information sets and approaches to uncertainty. The first bank is "informed", as it can genuinely anticipate phase switches and estimate exactly the relevant PD level. This is in line with the traditional Vasicek model. The second bank is "uninformed": it faces uncertainty about the phase – i.e. it knows only the parameters of the Markov process, not the relevant

[1] Because changes to the loading on the default-clustering factor have been extremely difficult to validate empirically, the literature has tended to treat this loading as constant over time (see Section 3.4.2). For its part, Basel III assumes that the loading is a deterministic decreasing function of the PD (see Appendix C). Under this assumption, EL and UL still do not decouple over realistic PD levels. Throughout the paper, we abstract from regulatory adjustments in relation to exposures' maturity (BCBS (2017)), as these would overburden the exposition without affecting the upshot of the analysis.

PD level – but accounts accurately for this uncertainty in assessing its EL and UL. The third bank is "naive": it has the same information set as the uninformed bank but assumes falsely that it can anticipate accurately phase switches, and that the PD it estimates pertains to the relevant phase. In other words, this bank applies the Vasicek model, even though its information set does not warrant it. All three banks set their loss-absorbing resources equal to the sum of perceived EL and UL and fail if their losses exceed these resources.

Keeping all else the same across the three banks, and in particular the probability of own failure that they target, we derive analytical expressions to answer the following two questions. First, what would be the shortfall of loss-absorbing resources – i.e. the difference between the level that generates a target failure probability and the actual level – if a bank ignored that there is uncertainty about the PD phase? The answer stems from a comparison between the naive and uninformed banks. Second, how would loss-absorbing resources change if a bank, aware of phase uncertainty, improved its capacity to forecast phase switches? For this question, we compare the informed and uninformed banks. With the phase-contingent risk factors assumed to be Gaussian, we need the following parameter values to perform each comparison: borrowers' PD in each phase, the phase-switching probabilities and the loading of borrowers' assets on the default-clustering factor.

The novel conceptual insight of our analysis stems from an investigation of how the answers to the above two questions depend on the asset portfolio's exposure to the transient default-clustering factor, ie to within-phase macro risk. First, we find that the naive bank faces a higher failure probability from the perspective of the uninformed bank when the loading on the default clustering factor is *lower*. That is, even though a bank faces a higher UL when its exposure to within-phase macro risk is higher (and should thus have higher capital), ignoring phase uncertainty is more detrimental for a bank with *lower* exposure to such risk. Second, lower exposure to within-phase macro risk also raises the importance of the quality of the information set. Namely, immediately after a switch from the low- to the high-PD phase, the informed bank perceives a higher failure probability of the uninformed bank when the loading on the default-clustering factor is lower. The overarching intuition is as follows. A lower loading on the default clustering factor means that benign realisations of this factor are less likely to compensate for wrongly abstracting from turning points in credit risk (the naive bank) or failing to anticipate such a turning point (the uninformed bank).

Motivated by the importance of loadings on the default-clustering factor in the context of default phase uncertainty, we develop a novel method for comparing these loadings across portfolios. The underlying intuition is that a smaller exposure to within-phase macro risk would imply – all else the same – less dispersion of phase-contingent losses. And when there are several phases, the modes associated with each one of them would be more distinct. This leads us to conjecture that one should be less likely to reject a multi-modality of the unconditional loss distribution when the portfolio loads less strongly on the default-clustering factor. We confirm this conjecture with three standard modality tests (developed by Cheng and Hall (1998), Hall and York (2001), and Ameijeiras-Alonso et al. (2019)) in a Monte Carlo environment. Finally, using credit-loss data for two portfolio types, we find that the business loan portfolio is less exposed to within-phase macro risk than the real-estate loan portfolio.

In sum, we argue that the regulatory framework should address uncertainty about turning points in credit risk with both universal and bank-specific measures. In general, the regulator should not be satisfied only with an accurate EL estimate but should seek to align the loss-absorbing resources with a loss-rate probability distribution that incorporates estimates of phase-switch probabilities and phase-contingent PDs. The regulator should encourage greater accuracy in the forecasts of time-varying loss rates, but especially in the case of banks with a *low* exposure to within-phase macro risk. In addition, when building loss-absorbing resources, it is more important for such banks to account for their uncertainty about turning points in credit risk. We show the relevance of this takeaway with a new empirical method that rank-orders banks' loadings on a within-phase macro risk factor by comparing the modality of their unconditional loss distributions.

Roadmap. The rest of this paper is organised as follows. Section 2 presents the risk environment in which we conduct the analysis. Then, Section 3 introduces three banks and discusses how risk parameters affect their loss-absorbing resources. Sections 4-5 conduct cross-bank comparisons to first determine the implications of accounting for or reducing phase-switch uncertainty and then dig into these implications. Section 6 develops an empirical method for comparing portfolios with respect to their within-phase exposure to macro factors. Section 7 concludes.

2 Risk environment

We develop the risk environment in three steps. First, we outline a simplified version of the credit-risk model used for international bank regulatory standards (Vasicek (1991), Gordy (2003), BCBS (2017)), henceforth "the Vasicek model". The simplifying assumption is that the credit portfolio comprises homogeneous exposures. Second, we present a novel extension of the Vasicek model that introduces uncertainty about turning points in credit risk. The assumed homogeneity facilitates the analysis of this uncertainty. Third, we derive the probability distribution of portfolio losses.

2.1 Simplified Vasicek model

At the beginning of each year t, the portfolio is composed of n homogeneous one-year loans and is asymptotic. The loans are *homogeneous* in the sense that each one is of size $1/n$ and is extended to a (non-defaulted) borrower, with all borrowers sharing the same one-year probability of default (PD_t) at the beginning of t. The portfolio is *asymptotic* because $n \to \infty$.

A representative borrower j, defaults in year t if the value of its assets falls below its debt, $A_{j,t} < D_t$. The only stochastic elements are the macro and borrower-specific risk factors – G_t and $Z_{j,t}$, respectively – that drive the asset value: $A_{j,t} = \mu_{A,t} + \sigma_{G,t} G_t + \sigma_{Z,t} Z_{j,t}$. Concretely, $G_t \sim N(0,1)$, $Z_{j,t} \sim N(0,1)$ and the two are mutually and serially independent and inherently *unpredictable*, with $Z_{j,t}$ independent across j. The volatility parameters – $\sigma_{G,t} > 0$ and $\sigma_{Z,t} > 0$ – and the expected asset value – $\mu_{A,t}$ – are *known*. The default condition is typically rewritten in terms a default threshold, B_t:

$$\begin{gathered} \rho_t G_t + \sqrt{1-\rho_t^2} Z_{j,t} < B_t = \Phi^{-1}(PD_t), \text{ where:} \\ \rho_t \equiv \frac{1}{\sqrt{1+\sigma_{Z,t}^2/\sigma_{G,t}^2}} \in (0,1) \text{ and } B_t \equiv \frac{D_t - \mu_{A,t}}{\sqrt{\sigma_{G,t}^2 + \sigma_{Z,t}^2}}, \end{gathered} \tag{1}$$

Φ is the standard-normal CDF and the equality on the first line of (1) follows from the left-hand side being equal to a standard normal variable. Thus, PD_t increases in the debt level but decreases in the expected asset value. And, since debt needs to be lower than the expected asset value for realistically low levels of PD_t (see below) – i.e. $D_t - \mu_{A,t} < 0$ – PD_t also increases in the overall asset volatility, $\sigma_{G,t}^2 + \sigma_{Z,t}^2$,

The key difference between the two risk factors surfaces at the portfolio level. While $Z_{j,t}$ is fully diversified at that level, G_t is not and generates default clustering. This clustering is stronger for a higher "loading parameter" ρ_t, which is equivalent to a higher volatility of the "default clustering" factor *relative* to that of idiosyncratic factors. Since the differentiation across borrowers stems only from the i.i.d. factor $Z_{j,t}$, the surviving borrowers in each period remain homogeneous.

2.2 Extending the model: uncertainty about the PD phase

To build on the Vasicek model, we note that the evolution of credit losses in Figure 1 has specific implications for the time series properties of PD_t. This evolution features turning points that lead to *persistent* low- or high-loss phases. In the light of expression (1), given that G_t and Z_t are i.i.d. over time, persistent shocks to credit losses cannot be stemming from changes in ρ_t for a fixed PD_t. Such changes can drive only period-to-period deviations of the loss rate around the PD_t. By contrast, the observed time profile of default rates indicates serial dependence in PD_t.

Thus, we need to introduce a persistent macro risk factor that drives PD_t, alongside the i.i.d. macro factor that drives default clustering for a given PD_t. By (1), the former risk factor would be related to components of the default point B_t – i.e. the level of debt, the expected asset value, or the asset volatility. That said, we will continue assuming that the relative volatility of the two i.i.d factors is known so that ρ_t is also known.

Our specific modelling decision is guided by the analytical complexities that can arise in the Vasicek model in the presence of a second risk factor (Gordy (2003)). To avoid these complexities, we opt for an extremely simple process that leads PD_t to take on one of two possible values. While such a process will not fully account for the richness of loss rates in the data – eg it leaves all within phase variability to be explained by the default-clustering factor – it is designed to capture in a stylised fashion abrupt transitions between extreme loss phases as well as within-phase persistence.

Concretely, we let PD_t follow a two-phase Markov process. We assume that this process is independent of the other risk factors and has realisations $PD_t \in \left\{PD_t^l, PD_t^h\right\}$, where the levels of PD_t^l and PD_t^h are known at the beginning of period t and $PD_t^l < PD_t^h$. In addition, $\pi_t^x \equiv \Pr\left(PD_t = PD_t^x | PD_{t-1}^x\right)$ is the probability that phase $x \in \{l, h\}$ materialises in period t, conditional on the same phase being in place in period $t-1$. This is the phase-*continuation* probability.[2] Since this specification does not

[2]A similar Markov switching setup underpins studies of the effects of ratings-sensitive capital requirements on banks' capital buffers over the business cycle (see Peura and Jokivuolle (2004) and references therein).

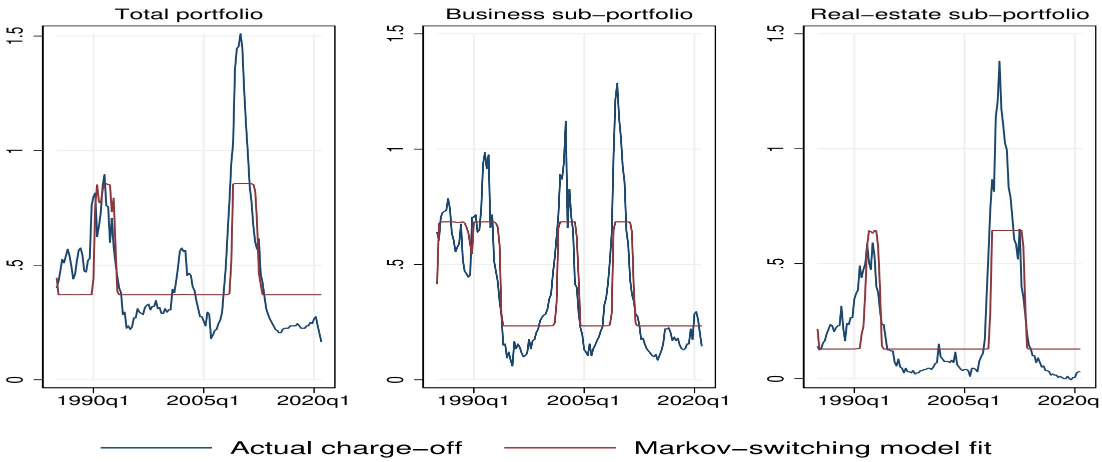

Figure 2: ***Fitting a Markov-switching model.*** Blue lines: Quarterly charge-off rates on loans by US banks to the non-financial private sectors (first panel) or sub-sectors (second and third panels). Red lines: the fitted values of estimating a two-regime Markov switching model.

distinguish between turning points and phase switches, we refer only to the latter in what follows.

A two-state Markov model – with constant phase-continuation probabilities – fits well the large abrupt swings in US loan loss rates (see Figure 2 and, for further detail, Appendix A). For the total loan portfolio, the model captures the high-loss phases associated with the banking crisis in the early 1990's and the great financial crisis of 2008. The persistence of the different phases surfaces as high estimates of phase-continuation probabilities, $\pi^l = 94\%$ and $\pi^h = 70\%$, respectively. By allowing for only two phases, the model misses the relatively mild increase in total portfolio loss rates that is associated with the burst of the dot-com bubble (hence, it misses the corresponding turning points identified in Figure 1). That said, the model does capture this increase (and the corresponding turning points) within the business loan sub-portfolio.

Remarks on the phase-continuation probabilities are in order. The constant probabilities, which we estimate as part of the Markov model, do not reflect any information other than the current phase. In practice, the macro-financial environment may generate uncertainty, driving the perceived continuation probability of the current phase towards 1/2 (i.e. towards a setting in which the forecaster's information set is as poor as that contained in a coin toss). We consider such a scenario in Section 3.4.2. Alternatively, if forecasters refine their estimates of a phase-contingent probability, these estimates would decline towards 0 in the run-up to an actual phase switch. We revisit this point in Section 5.

2.3 Loss distribution

In each period t, the portfolio's stochastic default rate is equal to[3] $\lim_{n\to\infty}\sum_{j=1}^{n}\frac{1}{n}\Pr_Z\left(\rho_t G_t+\sqrt{1-\rho_t^2}Z_{j,t}<\Phi^{-1}\left(PD_t\right)\right)=\Pr_Z\left(Z_{j,t}<\frac{\Phi^{-1}(PD_t)-\rho_t G_t}{\sqrt{1-\rho_t^2}}\right)$. Rearranging and assuming without loss of generality that LGD is 100%, we obtain that the portfolio's loss rate is:

$$Loss\left(G_t, PD_t;\rho_t\right)=\Phi\left(\frac{\Phi^{-1}\left(PD_t\right)-\rho_t G_t}{\sqrt{1-\rho_t^2}}\right). \tag{2}$$

The loss-rate probability distribution has three drivers. First, the default clustering factor G_t. Second, the known and potentially time-varying loading on this factor, ρ_t. Third, the persistent phase-switching factor PD_t.

3 Three banks: setting loss-absorbing resources

We now study three banks, for which the only source of potential losses are their loan portfolios. While all the banks operate in the above risk environment, they differ with respect to their perceptions, which reflect: (i) capacity to forecast the relevant loan PD, and (ii) the extent to which this capacity is taken into account when building loss-absorbing resources (LAR). At the beginning of year t, each bank sets its LAR equal to the $(1-\alpha)$-quantile of the *perceived* probability distribution of the random variable in (2). In case the actual losses exceed a bank's LAR, this bank fails and an identical bank replaces it, facing the same loss distribution that the failed one would have faced. Ultimately, while each bank perceives its one-year failure probability to be equal to α – which would tend to be a very low number (see Section 3.4.1 below) – this probability would typically not be equal to α from the perspective of the other banks.

In practice, LAR are broken into two parts. First, a bank's loan-loss reserves (the accumulated provisions net of write-offs) reflect the estimated expected loss (EL), equivalently the estimated loans' PD in our setting. Second, the difference between the $(1-\alpha)$ percentile of the perceived loss distribution and EL is denoted by unexpected

[3]See Gordy (2003) for the theoretical underpinning of an asymptotic portfolio's loss rate under weak conditions on the shape, continuity and differentiability of the idiosyncratic and default-clustering factors, Z_j and G. Our setup satisfies these conditions by assuming normality of these two factors.

losses (UL), which is covered by capital. With accounting and prudential authorities governing respectively provisioning and capital requirements (IASB (2014) and BCBS (2017)), there are differences between the time horizons underpinning the corresponding loss distribution. We abstract from these differences, which implies that the sum of reserves and capital equals LAR.

For each bank, we next derive EL and UL and study their properties.

3.1 Informed bank

The *informed* bank has exact knowledge of the relevant risk parameters and the relevant loan-loss phase. Denoting its LAR by the generic Λ, expression (2) implies that its failure probability is

$$\begin{aligned} FP^I\left(PD_t, \rho_t; \Lambda\right) &\equiv FP_t^I = \Pr\left(\Lambda < \Phi\left(\frac{\Phi^{-1}\left(PD_t\right) - \rho_t G_t}{\sqrt{1-\rho_t^2}}\right)\right) \\ &= \Phi\left(\frac{\Phi^{-1}\left(PD_t\right) - \sqrt{1-\rho_t^2}\Phi^{-1}\left(\Lambda\right)}{\rho_t}\right) \end{aligned}$$

In turn, targeting a failure probability of α, the informed banks sets its LAR to:

$$\Lambda^I\left(PD_t, \rho_t; \alpha\right) = \Phi\left(\frac{\Phi^{-1}\left(PD_t\right) - \rho_t \Phi^{-1}\left(\alpha\right)}{\sqrt{1-\rho_t^2}}\right). \tag{3}$$

Finally, we record this bank's EL and UL:

$$EL_t^I = PD_t \text{ and } UL^I\left(b_t, \rho_t; \alpha\right) \equiv UL_t^I = \Lambda^I\left(PD_t, \rho_t; \alpha\right) - PD_t, \tag{4}$$

3.2 Uninformed bank

While the *uninformed* bank knows the default clustering factor, ρ_t, and the phase-contingent levels of loans' probability of default, PD_t^l and PD_t^h, it faces uncertainty as to whether the period-$t-1$ phase $x \in \{l, h\}$ will continue in period t, $\pi_t^x \in (0,1)$. Taking this uncertainty into account, this bank attains its target failure probability by setting its loss-absorbing resources, $\Lambda^U\left(\pi_t^x, PD_t^x, PD_t^{\tilde{x}}; \rho_t, \alpha\right)$, at the (unique) solution of the following equation in terms of Λ:

$$
\begin{aligned}
FP^{U}\left(\Lambda;\pi_t^x,PD_t^x,PD_t^{\tilde{x}},\rho_t\right) &\equiv FP_t^U \\
&= \pi_t^x\Phi\left(\frac{\Phi^{-1}\left(PD_t^x\right)-\sqrt{1-\rho_t^2}\Phi^{-1}\left(\Lambda\right)}{\rho_t}\right) \\
&+\left(1-\pi_t^x\right)\Phi\left(\frac{\Phi^{-1}\left(PD_t^{\tilde{x}}\right)-\sqrt{1-\rho_t^2}\Phi^{-1}\left(\Lambda\right)}{\rho_t}\right)=\alpha
\end{aligned}
\tag{5}
$$

where $\widetilde{x}$ is the phase that did *not* materialize in year $t-1$. The uninformed bank's LAR is sandwiched between the corresponding phase-contingent levels for the informed bank (see Appendix *B.1* for a proof):

$$
\Lambda^{U}\left(\pi_t^x,PD_t^x,PD_t^{\tilde{x}};\rho_t,\alpha\right)\equiv\Lambda_t^U\in\left(\Lambda^I\left(PD_t^l;\rho_t,\alpha\right),\Lambda^I\left(PD_t^h;\rho_t,\alpha\right)\right). \tag{6}
$$

Ultimately, the uninformed bank perceives the following EL and UL:

$$
EL^{U}\left(\pi_t^x,PD_t^x,PD_t^{\tilde{x}}\right) \equiv EL_t^U=\pi_t^xPD_t^x+\left(1-\pi_t^x\right)PD_t^{\tilde{x}}, \tag{7}
$$

$$
UL^{U}\left(\pi_t^x,PD_t^x,PD_t^{\tilde{x}};\rho_t,\alpha\right) \equiv UL_t^U=\Lambda_t^U-EL_t^U. \tag{8}
$$

3.3 Naive bank

The *naive* bank has the same forecasting capacity as the uninformed bank but ignores the uncertainty it is subject to. The naive bank shares important features with a bank that adopts blindly the credit-risk model stipulated in global regulatory standards.[4] Concretely, if the current phase is x_t, the naive bank estimates $PD_t^N=\pi_t^xPD_t^x+\left(1-\pi_t^x\right)PD_t^{\tilde{x}}$ but perceives its failure probability to be equal to:

$$
FP^{N}\left(PD_t^N;\rho_t,\Lambda\right)=FP_t^N=\Phi\left(\frac{\Phi^{-1}\left(PD_t^N\right)-\sqrt{1-\rho_t^2}\Phi^{-1}\left(\Lambda\right)}{\rho_t}\right). \tag{9}
$$

[4] For a discussion of the "regulatory" bank, see Appendix C.

In turn, this leads it to set the following LAR, EL and UL:

$$\Lambda^N\left(PD_t^N;\rho_t,\alpha\right) \equiv \Lambda_t^N=\Phi\left(\frac{\Phi^{-1}\left(PD_t^N\right)-\rho_t\Phi^{-1}(\alpha)}{\sqrt{1-\rho_t^2}}\right), \tag{10}$$

$$EL_t^N = PD_t^N, \tag{11}$$

$$UL^N\left(PD_t^N;\rho,\alpha\right) \equiv UL_t^N=\Lambda_t^N-EL_t^N. \tag{12}$$

3.4 Discussion: Comparative statics and EL-UL decoupling

In this section, we conduct comparative statics from each bank's own perspective. Since the discussions of the informed and naive banks would be qualitatively identical in this setting, we consider only the former bank.

3.4.1 Parameter restrictions

For the discussion and proofs, we impose three parameter restrictions as maintained assumptions, which turn out to be borne one by the data.

First, we assume that, even in the low-loss phase, the loan PD is larger than the bank's target failure probability: $PD_t^l>\alpha$. One would expect this condition to hold in practice, as otherwise banks would have a higher cost of funding than their borrowers. In turn, this would make banks' intermediation model non-viable. Indeed, given that $\alpha=0.1\%$ in Basel III (BCBS (2017)), the condition is satisfied in our data on aggregate loss rates – the lowest one-year loan PD estimate reported in Appendix A is 0.44%.

Second, we assume that a borrower services its debt in the absence of shocks ($G_t = Z_{j,t}=0$). By expression (1), this requires $B_t<0$, which implies that $PD_t=\Phi^{-1}(B_t)<50\%$. The aggregate-loss data corroborates this condition, as the highest PD estimate reported in Appendix A is below 3%.

Third, we limit the persistence of the low-loss regime from above and below with the assumption, $PD_t^h<1-\pi_t^l<50\%$, which is also in line with our empirical estimates (Appendix A).

Combining all assumptions, we impose:

$$\alpha<PD_t^l<PD_t^h<1-\pi_t^l<50\%. \tag{13}$$

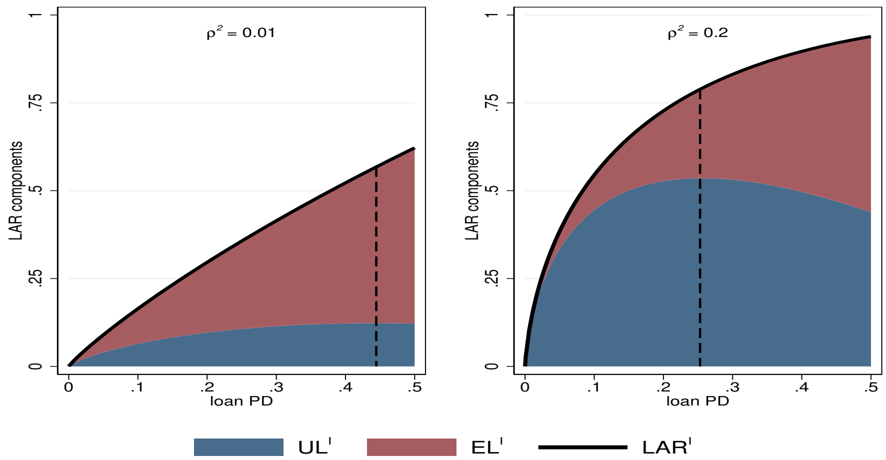

Figure 3: ***EL and UL: joined at the hip for realistic PDs.*** Loss-absorbing resources (LAR) and their components – expected loss (EL) and unexpected loss (UL) – from the perspective of the informed bank when the target one-year probability of bank failure is 0.1%. The dashed line indicates maximum UL.

3.4.2 Comparative statics

In Figure 3, we illustrate the impact of loans' PD and loading on the default clustering factor on the informed bank's EL, UL and LAR.[5] Trivially, EL^I increases in PD_t and is insensitive to ρ_t. In turn, it follows directly from (3) that Λ^I increases in both PD_t and ρ_t. While UL^I increases in ρ_t, the impact of loans' PD on UL^I is ambiguous. Note that UL is continuous in PD_t, zero at $PD_t = EL_t \in \{0, 1\}$ and positive for intermediate levels of PD_t. Thus, UL first increases but eventually decreases with PD_t. In the light of the sandwiching in (6) and the restrictions in (13), the picture is qualitatively the same for the uninformed bank.

That said, the relationships are less ambiguous over realistic parameter values. Given that the levels of ρ in Figure 3 span the range of this parameter's estimates in the literature (Düllmann et al. (2007), Zhang et al. (2008) and Bams et al. (2012)) and average PDs at the portfolio level tend to be way below 10% (Section 3.4.1), UL increases monotonically with PD_t in practice. Thus, when PDs changes over time, EL and UL either increase or decrease together.

[5]See Appendix B.2 for proofs of statements in Section 3.4.2 and for how these proofs make use of the assumptions sated in Section 3.4.1.

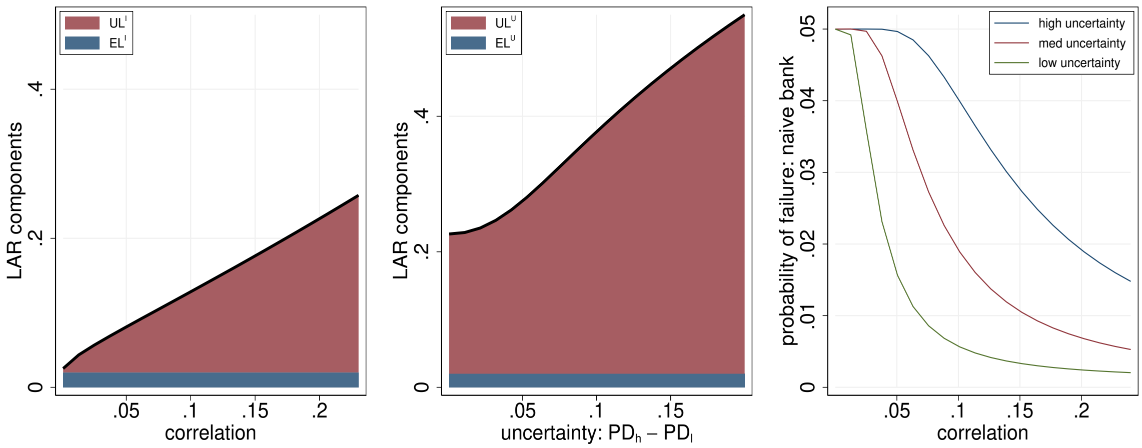

Figure 4: ***Sources of decoupling and implications for bank failure.*** Loss-absorbing resources (LAR) and their components – expected loss (EL) and unexpected loss (UL) – from the perspective of the informed bank (left-hand panel) and uninformed bank (centre panel) when the target one-year probability of bank failure is 0.1%. On the assumption that the latest phase is the low-PD one, the additional parameters are as follows: loan $PD = 2\%$ (left-hand panel, also the expected loan PD in centre panel and right-hand panels); $\rho^2 = 20\%$ (centre panel); $\pi^l = 95\%$. The right-hand panel plots the naive bank's probability of failure from the perspective of the uninformed bank, for different levels of uncertainty (PD^h–PD^l) = 5% (low), 10% (medium) and 20% (high).

Decoupling of EL and UL: loading on default-clustering factor. EL and UL may decouple if the loading on the default clustering factor changes over time. This is because such changes affect UL but not EL (see, for instance, expressions (3) and (4)). We illustrate this in Figure 4 (left-hand panel).

The empirical relevance of time variation in ρ_t is questionable. For instance, while Düllmann et al. (2007), Zhang et al. (2008) and Bams et al. (2012) do find evidence that this parameter differs across credit portfolios (see below), they do not find cyclical changes that could help explain e.g. the patterns of default clustering in Figure 1. In addition, in a setting akin to the one we use here, Zhou (2001) derives that the effect of a higher ρ_t on default clustering is similar to that of higher probabilities of default. A similar message stems from the key argument in Erlenmaier and Gersbach (2013), who study the impact of PDs on default correlations. Thus, we henceforth consider only a constant ρ, which allows us to concentrate on changes in PD_t over time.

Decoupling of EL and UL: uncertainty about the PD phase. Alternatively, the decoupling of EL and UL could stem from uncertainty about the credit-loss phase, i.e.

from uncertainty about PD_t. We consider two specific scenarios that reflect emerging uncertainty: (i) the opening of a wedge between PD_t^l and PD_t^h for a given π_t^x; (ii) π_t declining below 1 towards 1/2. In each of these scenarios, we keep EL constant.

The proposition below states that UL rises in each of these two scenarios. The proof – in Appendix B.3 – uses the fact that the uninformed bank's failure probability (expression (5)) is an increasing convex function of loans' PD.

Proposition 1 ***Effect of uncertainty on*** UL. *Suppose that each of the following two switch-to-uncertainty scenarios maintains* $EL_t^U = EL_{t-1}^U$*: (i)* $\pi_{t-1}^l = \pi_t^l$ *and* $PD_t^l < PD_{t-1}^l = PD_{t-1}^h < PD_t^h$ *or (ii)* $\pi_{t-1}^l = 1 > \pi_t^l > 0.5$, $PD_t^l < PD_{t-1}^l$ *and* $PD_t^h = PD_{t-1}^h$. *Under either scenario,* $UL_{t-1}^U < UL_t^U$.

In Figure 4 (centre panel), we illustrate the implications of scenario (i). While the blue area (EL) is flat, the red area (UL) is narrowest at the left edge, where there is no uncertainty. Moreover, the monotonic widening of the red area from left to right indicates decoupling between EL and UL even when there is an increase of *existing* uncertainty (i.e. when the wedge between PD_t^l and PD_t^h widens from a positive level). The picture is similar for scenario (ii).

4 LAR shortfalls: due to ignoring uncertainty

We now investigate how ignoring uncertainty-driven decoupling of EL and UL would affect the probability of bank failure. For this, we assume that the low-PD phase is in place in period $t-1$ and consider the naive bank from the perspective of the uninformed one. Proposition 1 implies that – from this perspective, for any given ρ and under either scenario (i) or (ii) – the naive bank has a strictly lower LAR than the one needed to attain the target failure probability of α.

Then, we consider how LAR evolves as ρ declines towards 0. In the process, the distribution of portfolio loss rates converges to a degenerate distribution with only one realisation in each phase, equal to PD_t^l and PD_t^h (Figure 5). Given the assumed phase in period $t-1$, the probability of these realisations is π_t^l and $1-\pi_t^l$, respectively. In parallel, since $\alpha < 1-\pi_t^l$ by expression (13), the level of LAR that is consistent with the target failure probability from the uninformed bank's perspective converges to PD_t^h. Ultimately, since the naive bank's LAR would then be lower than PD_t^h by Proposition 1, its failure probability would converge to $1-\pi_t^l$ as $\rho \to 0$.

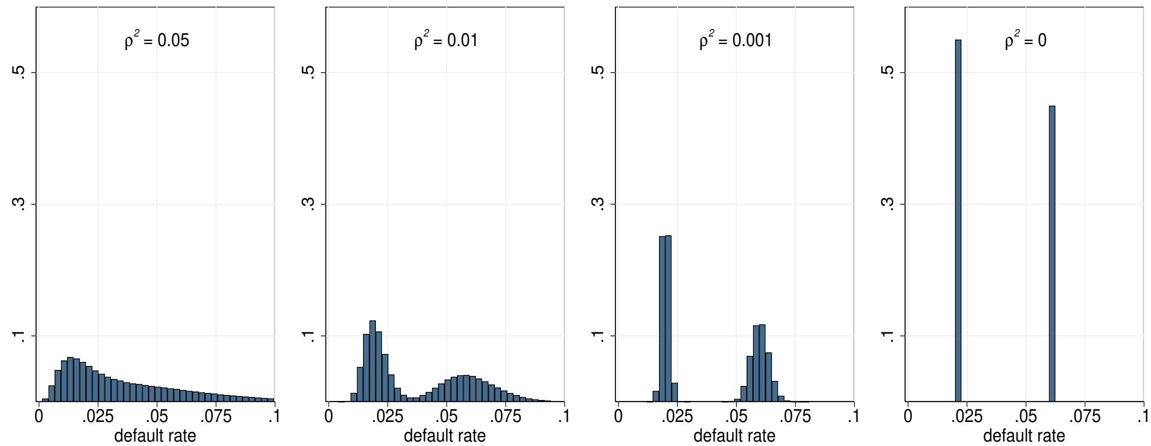

Figure 5: ***Effect of ρ on the loss distribution.*** Conditional on the latest phase being the low-PD one, with additional model parameters: $\pi^l = 55\%$; $PD^l = 2\%$ and $PD^h = 6\%$. The value of π^l is chosen to optimize readability.

Consider two banks of the naive type. Suppose that the loan portfolio of one of the banks is governed by values of ρ, π_t^l, PD_t^l and PD_t^h that imply a failure probability below $1 - \pi_t^l$. By the above reasoning, the failure probability of the other bank would be higher if it faces the same π_t^l, PD_t^l and PD_t^h but a lower ρ that is sufficiently close to 0. This is stated formally in the following proposition and proved in Appendix B.4.

Proposition 2 ***Effect of ignoring uncertainty on failure probability.*** *Suppose that the state in period $t-1$ is $x = l$. There exist $\bar{\rho} < 1$ and $\underline{\rho} < \bar{\rho}$ such that a naive bank's failure probability is higher at $\rho \in \big(0, \underline{\rho}\big)$ than at $\rho \in (\bar{\rho}, 1)$.*

For any given π_t^l, PD_t^l and PD_t^h, Appendix B.4 delivers an implicit expression for $\bar{\rho}$. For our estimates of the former three parameters within the Markov switching model (Appendix A), we derive $\bar{\rho}$ on the basis of that expression and report the results for different loan portfolios in Table A.1. All derived values of $\bar{\rho}$ are so low that $(\bar{\rho}, 1)$ encompasses any estimate of asset-return correlations that we are aware of in the literature (Düllmann et al. (2007), Zhang et al. (2008) and Bams et al. (2012)) and regulatory texts (BCBS (2017)).

In Figure 4 (right-hand panel), we employ specific parameterisations of the risk environment in Section 2. In these examples, the naive bank's failure probability increases as ρ decreases over a wide interval.

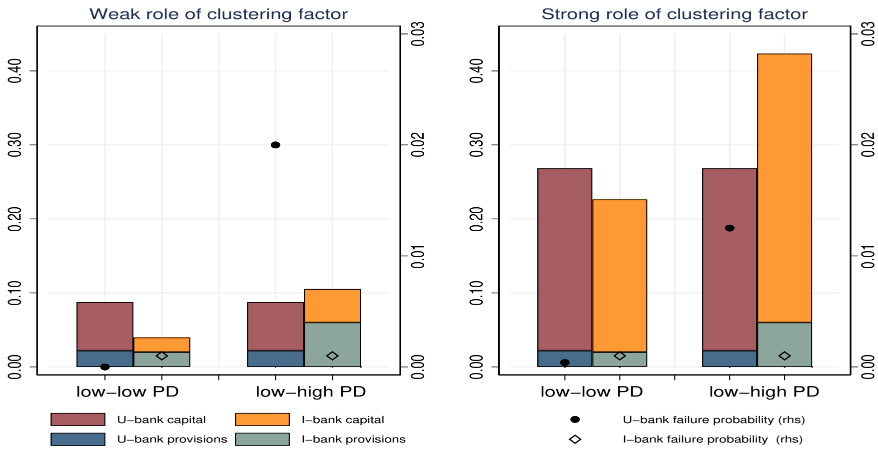

Figure 6: ***Implications of missing a turning point.*** The underlying parameterisation is: $PD^l = 2\%$, $PD^h = 6\%$, $\alpha = 0.1\%$, $\pi^l = 95\%$, $\rho^2 = 1\%$ (left-hand panel) and $\rho^2 = 20\%$ (right-hand panel).

5 LAR shortfalls: due to uncertainty

We now change perspectives to study how the uncertainty in loss forecasts – that is, the mere existence of EL-UL decoupling – affects LAR shortfall and banks' failure probability. To this end, we treat the "appropriate" LAR and the "true" failure probability to be those derived and perceived, respectively, by the informed bank and study how they differ from those of the uninformed bank. For brevity, we discus only a scenario where the latest phase has featured PD^l_{t-1}. The implications are symmetric for PD^h_{t-1}.

While the informed bank's failure probability is always at the target level, the uninformed bank's can be lower or higher (*Figure 6*, diamonds versus dots). Concretely, the latter probability is below target if the low-loss phase continues (low-to-low scenarios: dots below diamonds) and above target if there is a low-to-high switch. This reflects LAR excesses, respectively shortfalls, stemming from the sandwiching property of the uninformed bank's LAR: expression (6).

Our key insight in this context is that the overshooting of the target failure probability is more pronounced if the uninformed bank holds a portfolio that is *less* exposed to within-phase macro risk (see Appendix B.5 for a proof).

Proposition 3 ***Failure probability and exposure to default clustering.*** *Suppose that the phase sequence delivers PD^l_{t-1} and PD^h_t. When the uninformed bank sets its loss-absorbing resources according to (5), its probability of failure decreases with ρ.*

We illustrate this proposition in *Figure 6*, where the low-to-high dot in the left-hand panel is above that in the right-hand panel. A lower ρ implies that phase-contingent losses are more certain, i.e. that $UL^{\alpha,I}$ and $UL^{\alpha,U}$ are lower, which leads to lower capital levels in the left-hand panel than in the right-hand panel. The flip side of this is a lower likelihood that a benign realisation of the default-clustering factor (G_t) would undo the implications of missing a phase switch. Thus, when the phase-contingent losses are more certain, it is also more certain that missing a phase switch would be detrimental for the uninformed bank.

When the loan portfolio's exposure to within-phase macro risk is sufficiently low, under-*provisioning* would be *solely* responsible for the uninformed bank's LAR shortfall in a low-to-high scenario. Imposing a weak restriction on the failure-probability target – namely, that it is lower than the probability of switching phases – we prove the following proposition in Appendix B.6.

Proposition 4 EL ***as a driver of LAR shortfall.*** *Suppose that the phase sequence delivers PD^l_{t-1} and PD^h_t. If $1 - \pi^l_t \geq \alpha$, there exists $\underline{\rho} > 0$ such that the uninformed bank perceives a higher UL than the informed bank for any $\rho < \underline{\rho}$. In this case, the uninformed bank's LAR shortfall stems entirely from under-provisioning, due to under-estimated EL.*

We illustrate this proposition also in *Figure 6*. In the left-hand panel, despite the uninformed bank's LAR shortfall in the boom-bust scenario, the UL it perceives (and thus its capital) is actually *larger* than that of the informed bank. This is because the impact of a decline in ρ on perceptions differs between the two banks. As ρ declines to zero, losses become more and more certain from the perspective of the informed bank: the perceived UL and capital shrink to zero. From the perspective of the uninformed bank, however, a decline in ρ cannot eliminate the uncertainty about the phase-switching factor: thus, UL cannot decline to zero. With the uniformed bank's UL being larger than the informed bank's, the former bank's LAR shortfall comes from underestimating EL.

In sum, an improvement of the uninformed bank's forecasting capacity would have the following implications. It would affect the phase-switching probabilities, with e.g. π^l_{t-1} decreasing on the cusp of a low-to-high switch. All else the same, the benefit of better forecasting capacity would be greater if the portfolio of the uninformed bank is less exposed to macro risk within a phase (Proposition 3). And when such exposure is sufficiently low, the improved forecasting capacity would lead to *lower* capital and higher reserves (Proposition 4).

	Total portfolio		Business sub-portfolio		Real-estate sub-portfolio	
test	stat	p	stat	p	stat	p
CH	0.05	0.29	0.07**	0.01	0.04	0.63
HY	0.36*	0.08	0.26**	0.03	0.25	0.37
ACR	0.06*	0.09	0.06*	0.05	0.04	0.40

Table 1: ***Tests of multi-modality.*** CH: Cheng and Hall (1998); HY: Hall and York (2001); ACR: Ameijeiras-Alonso et al. (2019). The null hypothesis in each test is that the distribution of a variable is uni-modal. $*$ indicates significance at the 10% level; $**$ indicates significance at the 5% level.

6 Practical relevance

Our analysis would be relevant in practice if we can answer the question: Which portfolios load less on within-phase macro risk factors?

In addressing this question, we go through the following thought process. First, we refer to Figure 5, which illustrates that the existence of a low- and a high-loss phase gives rise to a bi-modal loss distribution. Second, we observe that the bi-modality is more distinct for a lower loading, ρ, on the within-phase macro factor. This suggests that, for a given sample size, statistical tests would be more likely to reject uni-modality when ρ is low.

Then, we turn to the two distinct sub-portfolios in our dataset, comprising respectively business and real-estate loans. We recall that each portfolio undergoes switches between two phases (see Figure 2). That said, bi-modality is more clearly visible for the business-loan portfolio than for the real-estate one. To investigate modality formally, we run a battery of bi-modality tests based on Cheng and Hall (1998), Hall and York (2001) and Ameijeiras-Alonso et al. (2019) in (Table 1). All three tests reject the null for business loans (at either the 10% or 5% significance levels). By contrast, the null cannot be rejected for real estate loans. In the context of the model in the main text, these results suggest that business loans have a lower loading factor ρ.

Of course, these results could stem from portfolio specificities other than ρ. For instance, there are fewer observations for the real-estate portfolio in the high-loss phase – reflected in a notably lower π^h estimate: 72% vs 88% for the business-loan portfolio. The lower persistence of the high-PD phase would translate into fewer observations in this phase, thus making a second mode less easy to detect. That said, while the two portfolios feature similar PD^h estimates, the PD^l estimate for the real-estate portfolio is much lower: 0.44% vs 0.81% for business loans. The bigger difference between the low- and high-loss phases for the real-estate portfolio should have made it easier to detect bi-modality all else the same.

To place these considerations on a more formal footing, we resort to Monte Carlo simulations. For a given value of the loading ρ on the within-phase macro factor and referring to equation (2), we simulate two time series of default rates: one based on the Markov-process parameter estimates for business loans and another on the corresponding estimates for real-estate loans (see Appendix A and Table A.1 in particular). Replicating each time series many times and applying the above three tests at each replication, we calculate the corresponding rates of rejecting a null hypothesis of unimodality. We repeat these steps over a range of ρ values and report the results in Table A.2. While the rejection rates are roughly similar across tests and loan portfolios for the same ρ, they decline monotonically as ρ increases. We interpret the results in Table 1 as confirming that ρ is smaller for the business-loan portfolio. Accounting for and/or mitigating phase uncertainty would make a bigger difference under this portfolio.

7 Concluding remarks

Our paper has delivered three contributions. First, a straightforward generalisation of a well-known credit-risk model allows for studying analytically the importance of (i) reducing the uncertainty about turning points in credit risk and (ii) accounting for any remaining uncertainty when setting loss-absorbing resources. Second, we find that this importance is especially high for a bank whose phase-contingent losses depend little on macro risk factors. Third, we show that testing for the number of modes in loss rates' unconditional distribution helps rank-order portfolios with respect to their within-phase exposure to macro risk.

Uncertainty about turning points in credit risk may lead expected and unexpected losses to evolve in different directions. Thus, empirical forecasts need to target different aspects of the loss distribution on the basis of multiple forecast variables. To reduce the uncertainty about turning points, these variables would need to capture risk-taking as it builds up, in the spirit of the literature on early warning indicators of banking crises (Detken et al. (2014), Tölö et al. (2018), Aldasoro et al. (2019) and references therein) and more recent advances in default-risk forecasting (Lu and Nikolaev (2022) and Juselius and Tarashev (2020)). To the extent that private incentives deviate from system-wide financial stability (Borio and Zhu (2008), Acharya (2009) and Gorton and Ordoñez (2014)) – a distortion that we have abstracted from – it will be up to a bank supervisor to ensure the proper execution of empirical forecasts and to map the forecasts into adequate loss-absorbing resources.

References

Abad, J. and Suárez, J. (2017), 'Assessing the cyclical implications of IFRS 9 – a recursive model', ESRB Occasional Paper Series 12.

Acharya, V. (2009), 'A theory of systemic risk and design of prudential bank regulation', *Journal of Financial Stability* **5**(3), 224–55.

Aldasoro, I., Borio, C. and Drehmann, M. (2019), 'Early warning indicators of banking crises: expanding the family', *BIS Quarterly Review*, March.

Ameijeiras-Alonso, J., Crujeiras, R. M. and Rodriguez-Casal, A. (2019), 'Mode testing, critical bandwidth and excess mass', *Test* **28**, 900–919.

Bams, D., Pisa, M. and Wolff, C. (2012), 'Modeling default correlation in a us retail loan portfolio', Luxembourg School of Finance WP Series 12-19.

Bank of England (2016), 'The financial policy committee's approach to setting the countercyclical capital buffer: A policy statement'.

BCBS (2005), 'An explanatory note on the Basel II IRB risk weight functions'.

BCBS (2017), 'Basel III: Finalizing post-crisis reforms'.

BCBS (2022), 'Newsletter on positive cycle-neutral countercyclical capital buffer rates', Basel Committee on Banking Supervision.

Borio, C. and Zhu, H. (2008), 'Capital regulation, risk-taking and monetary policy: a missing link in the transmission mechanism?', BIS Working Paper No. 268.

Chae, S., Sarama, R., Vojtech, C. and Wang, J. (2018), 'The impact of the current expected credit loss standard (CECL) on the timing and comparability of reserves', Finance and Economics Discussion Series 2018-020.

Cheng, M. Y. and Hall, P. (1998), 'Calibrating the excess mass and dip tests of modality', *Journal of the Royal Statistical Society* **60**, 579–589. Series B.

Covas, F. and Nelson, W. (2018), 'Current expected credit loss: Lessons from 2007-2009', BPI Staff Working Paper 2018-1.

Detken, C., Weeken, O., Alessi, L., Bonfim, D., Boucinha, M., Castro, C., Frontczak, S., Giordana, G., Giese, J., Jahn, N., Kakes, J., Klaus, B., Lang, J., Puzanova, N. and Welz, P. (2014), 'Operationalising the countercyclical capital buffer: indicator selection, threshold identification and calibration options', ESRB Occasional Paper Series 5, European Systemic Risk Board.

Düllmann, K., Scheicher, M. and Schmieder, C. (2007), 'Asset correlations and credit portfolio risk – an empirical analysis', Deutsche Bundesbank Discussion Paper No13/2007.

Erlenmaier, U. and Gersbach, H. (2013), 'Default Correlations in the Merton Model', *Review of Finance* **18**(5), 1775–1809.

European Union (2013), 'Prudential requirements for credit institutions and investment firms', Regulation (EU) No 575/2013 of the European Parliament and of the Council, European Union.

FASB (2016), 'Accounting standards update no. 2016–13, financial instruments – credit losses (topic 326): Measurement of credit losses on financial instruments'.

Federal Reserve Board (2022), '2022 Stress Test Scenarios', Board of Governors of the Federal Reserve System.

FSF (2009), 'Report of the financial stability forum on addressing procyclicality in the financial system'.

Goncharenko, R. and Rauf, A. (2020), 'Loan loss provisioning requirements in a dynamic model of banking', mimeo.

Gordy, M. (2003), 'A risk-factor model foundation for ratings-based bank capital rules', *Journal of Financial Intermediation* **12**(3), 199–232.

Gorton, G. and Ordoñez, G. (2014), 'Collateral crises', *American Economic Review* **104**(2).

Hall, P. and York, M. (2001), 'On the calibration of silverman's test for multimodality', *Test* **11**, 515–36.

Harding, D. and Pagan, A. (2002), 'Dissecting the cycle: a methodological investigation', *Journal of Monetary Economics* **49**(2), 365–381.

Harris, T. S., Khan, U. and Nissim, D. (2018), 'The Expected Rate of Credit Losses on Banks' Loan Portfolios', *The Accounting Review* **93**(5), 245–271.

Hartigan, J. A. and Hartigan, P. M. (1985), 'The DIP test of unimodality', *Annals of Statistics* **13**, 70–84.

IASB (2014), 'IFRS 9 financial instruments'.

Juselius, M. and Tarashev, N. (2020), 'Forecasting expected and unexpected losses', BIS Working Papers No 913.

Krüger, S., Rösch, D. and Scheule, H. (2018), 'The impact of loan loss provisioning on bank capital requirements', *Journal of Financial Stability* **36**, 114–29.

Loudis, B. and Ranish, B. (2019), 'CECL and the credit cycle', Finance and Economics Discussion Series 2019-061, Board of Governors of the Federal Reserve System.

Lu, Y. and Nikolaev, V. V. (2022), 'Expected Loan Loss Provisioning: An Empirical Model', *The Accounting Review* **97**(7), 319–346.

Peura, S. and Jokivuolle, E. (2004), 'Simulation based stress tests of banks' regulatory capital adequacy', *Journal of Banking and Finance* **28**(8), 1801–24.

Silverman, B. (1981), 'Using kernel density estimates to investigate multimodality', *Journal of the Royal Statistical Society* **43**, 97–99. Series B.

Tölö, E., Laakkonen, H. and Kalatie, S. (2018), 'Evaluating indicators for use in setting the countercyclical capital buffer', *International Journal of Central Banking* **14**(2), 51–112.

Vasicek, O. (1991), 'Limiting loan loss probability distribution', KMV Working Paper.

Zhang, J., Zhu, F. and Lee, J. (2008), 'Asset correlation, realised default correlation, and portfolio credit risk', Moody's KMV.

Zhou, C. (2001), 'An analysis of default correlations and multiple defaults', *The Review of Financial Studies* **14**(2).

Appendices

A Regime-switching in the data

In this appendix, we estimate a two-state Markov process for loss rates in the US banking sector: for the total loan portfolio and two sub-portfolios, comprised of real estate and business loans. We then use the model estimates for two purposes. We first derive values for $\bar{\rho}$ in Proposition 2. Second, we conduct a Monte Carlo simulation to show that testing for bi-modality of the unconditional loss-rate distribution can reveal different degrees of within-phase exposure to macro credit-risk factors.

Our loan loss data consist of quarterly net charge-off rates on loans from the US banking sector to the non-financial private sector. These series are obtained from the Federal Reserve Board of Governors. On the assumption that the underlying loss-given-default is 50%, we obtain default-rate estimates by multiplying the raw series by 2. We estimate probabilities of default (PDs) for the total portfolio, indexed by T, and separately for the two sub-portfolios: commercial and industrial or "business" loans, B, and real-estate loans, R. The sample begins in 1985q1 and ends on 2021q1.

We fit a two-state Markov-switching model to charge-off rates (CoR):

$$CoR^x_{k,t} = PD^x_k + \varepsilon_t \tag{14}$$

where $k \in \{T, B, R\}$, $x \in \{l, h\}$ stands for the unobserved state, and ε_t is a white noise error with variance σ^2. Without loss of generality, we assume that $PD^l_k < PD^h_k$. In addition, we model x as an irreducible, aperiodic Markov chain, with phase-continuation probability π^x. The model is estimated by maximizing the implied likelihood function using numerical techniques (Table A.1). Figure 2 in the main text shows the model's fit (red vs blue lines). The estimates capture fairly well the jump dynamics in the default rates. More elaborate specifications, for instance with auto-regressive terms, deliver similar point estimates of the parameters but different standard errors of these estimates. Thus, statistical inference on the basis of the simple model requires caution.

The upper part of Table A.1 delivers two messages. First, loan PDs differ across regimes, eg 1.41% vs 2.84% for the total portfolio. Second, as assumed in the main text, the PDs in tranquil times (the PD^l_k estimates) are consistently above banks' target probability of failure, $\alpha = 0.1\%$ and those in times of stress are way below 50%.

The estimates in the upper part of Table A.1 allow us to derive values for a parameter

	Total portfolio		Business sub-portfolio		Real-estate sub-portfolio	
Param.	Estimated	Annualized	Estimated	Annualized	Estimated	Annualized
PD^l	0.35 (0.02)	*1.41* [1]	0.20 (0.02)	*0.81* [1]	0.11 (0.01)	*0.44* [1]
PD^h	0.95 (0.04)	*2.84* [1]	0.73 (0.02)	*2.11* [1]	0.74 (0.04)	*1.95* [1]
σ	0.17 (0.01)	*0.68*	0.14 (0.01)	*0.56*	0.15 (0.01)	*0.60*
π^l	0.98 (0.01)	*0.94* [1]	0.97 (0.02)	*0.89* [1]	0.98 (0.01)	*0.94* [1]
π^h	0.91 (0.06)	*0.70* [2]	0.96 (0.03)	*0.83* [2]	0.92 (0.05)	*0.72* [2]
Memo item: parameter in Proposition 2						
$\bar{\rho}^2$	*0.0020* [1]		*0.0019* [1]		*0.0061* [1]	
	0.0004 [2]		*0.0003* [2]		*0.0007* [2]	

Table A.1: ***Markov switching parameter estimates.*** Based on fitting the model in (14) to US quarterly loan loss rates from 1985q1 to 2020q4 (144 quarters). Standard errors in parenthesis. The annualized estimates (in italics) abstract from paths of quarterly states along which there is a switch reversal (e.g. from l to h and then back to l). [1] Conditional on $x = l$ in the latest quarter. [2] Conditional on $x = h$ in the latest quarter.

in Proposition 2, $\bar{\rho}$. For each set of estimates for π^l, PD_k^l and PD_k^h, we report these values squared (memo item in Table A.1), so that they can be directly compared with estimates of asset-return correlations in the literature.

Finally, we study how different values of ρ affect our ability to detect the different modes in the loss distribution associated with low and high loss regimes. For this, we run a Monte Carlo simulation, using the estimated parameters under the real-estate and business-loan portfolios, respectively. In particular, we draw 1000 time series of length 150 (a rounding of the length of our sample: 37 years or 144 quarters) of the portfolio loss rate according to equation (2) for each model. In calculating the frequency of rejecting the null that the loss rate's distribution is uni-modal, we consider three formal bi-modality tests. The first is a test by Cheng and Hall (1998), which seeks to reduce the conservatism of the dip test by Hartigan and Hartigan (1985). The second test is by Hall and York (2001), which improve on the kernel density-based test by Silverman (1981). The third is a recent test by Ameijeiras-Alonso et al. (2019), which mixes elements from the other two tests. The rejection frequencies are similar across tests and portfolios for a given ρ and decline monotonically as ρ declines (Table A.2).

ρ	0.05	0.08	0.1	0.13	0.15	0.18	0.2	0.23	0.25
				Based on business loan parameters					
CH	0.92	0.87	0.73	0.51	0.27	0.11	0.06	0.05	0.03
HY	0.99	0.97	0.89	0.54	0.15	0.04	0.03	0.02	0.01
ACR	0.94	0.90	0.79	0.47	0.22	0.10	0.06	0.05	0.05
				Based on real estate loan parameters					
CH	0.66	0.59	0.47	0.31	0.17	0.07	0.03	0.02	0.01
HY	0.95	0.93	0.82	0.61	0.34	0.15	0.09	0.05	0.03
ACR	0.73	0.65	0.58	0.41	0.24	0.11	0.06	0.04	0.03

Table A.2: ***Rejection frequencies of uni-modality, simulated data.*** Based on 1000 random draws of a time series of length 150 of the portfolio loss rate according to Equation (2) and the two-phase Markov switching process for PD_t^x. In parameterising the latter process, we refer sequentially to estimates for business loans and real-estate loans in Table A.1. Each number corresponds to the rejection frequency of the null: that the loss rate's distribution is uni-modal. The underlying tests are from: CH: Cheng and Hall (1998); HY: Hall and York (2001); ACR: Ameijeiras-Alonso et al. (2019).

B Proofs

The proofs in this appendix make use of the assumptions stated in Section 3.4.1.

B.1 Proof of sandwiching claim in Section 3.2

We derive equation (5) and show that it has exactly one solution in terms of Λ.

Given PD_{t-1}, the uninformed bank knows the following when its LAR equals Λ:

$$\begin{aligned}
&\Pr\left(\textit{loan losses} > \Lambda\right) \\
&= \pi_t^x \Pr_G\left(\textit{loan losses} > \Lambda | G, x\right) + \left(1-\pi_t^x\right) \Pr_G\left(\textit{loan losses} > \Lambda | G, \widetilde{x}\right) \\
&= \pi_t^x \Pr_G\left(\Phi\left(\frac{\Phi^{-1}\left(PD_t^x\right) - \rho G_t}{\sqrt{1-\rho^2}}\right) > \Lambda\right) + \left(1-\pi_t^x\right) \Pr_G\left(\Phi\left(\frac{\Phi^{-1}\left(PD_t^{\tilde{x}}\right) - \rho G_t}{\sqrt{1-\rho^2}}\right) > \Lambda\right) \\
&= \pi_t^x \Phi\left(\frac{\Phi^{-1}\left(PD_t^x\right) - \sqrt{1-\rho^2}\Phi^{-1}\left(\Lambda\right)}{\rho}\right) + \left(1-\pi_t^x\right) \Phi\left(\frac{\Phi^{-1}\left(PD_t^{\tilde{x}}\right) - \sqrt{1-\rho^2}\Phi^{-1}\left(\Lambda\right)}{\rho}\right)
\end{aligned}$$

which is the left-hand side of (5).

There is exactly one value of Λ that sets this expression equal to α. This is because the expression is monotonically decreasing in Λ and converges to 1 (respectively, 0) as $\Lambda \to 0$ (respectively, 1).

Next, we show that Λ_t^U is sandwiched between its complete-knowledge counterparts:

$\Lambda_t^U \in \left(\Lambda^I\left(PD_t^l\right), \Lambda^I\left(PD_t^h\right)\right)$. Since each of the two summands on the left-hand side of (5) is strictly decreasing in Λ, equation (3) implies that the left-hand side of (5) is smaller (respectively, larger) than α if $\Lambda_t^U \geq \Lambda^I\left(PD_t^h\right)$ (respectively, if $\Lambda_t^U \leq \Lambda^I\left(PD_t^l\right)$). Ultimately, since (5) has exactly one solution, we obtain the desired result.

B.2 Proof of statements in Section 3.4.2

The results $\partial UL^I/\partial\rho > 0$ and $\partial\Lambda^I/\partial\rho > 0$ follow from equations (3)-(4) and the assumption that $\Phi^{-1}(\alpha) < \Phi^{-1}(PD_t) < 0$. To see this, note that
$\partial\left(\left(\Phi^{-1}(PD_t) - \rho\Phi^{-1}(\alpha)\right)/\sqrt{1-\rho^2}\right)/\partial\rho = \left(\rho\Phi^{-1}(PD_t) - \Phi^{-1}(\alpha)\right)/\left(1-\rho^2\right)^{3/2} > 0$.

We similarly prove that $\partial UL^U/d\rho > 0$ and $\partial\Lambda^U/d\rho > 0$. From equation (3), it follows that $\Phi^{-1}(PD_t^x) = \rho\Phi^{-1}(\alpha) + \sqrt{1-\rho^2}\Phi^{-1}\left(\Lambda^I\right)$. Since $\alpha < PD_t^x$ for $x \in \{l,h\}$ by (13), and $\Lambda^U \in \left(\Lambda^I\left(PD_t^l\right), \Lambda^I\left(PD_t^h\right)\right)$ by (6), it follows that $\Lambda^U > PD_t^x$ for $x \in \{l,h\}$. Turning to (5), $\Lambda^U > PD_t^x$ implies that each summand increases in ρ, which in turn implies that $\partial\Lambda^U/d\rho > 0$. Since $\partial EL^U/\partial\rho = 0$, it then follows that $\partial UL^U/\partial\rho = 0$.

Next, we prove that UL^I is non-monotonic in PD_t. In particular, noting that $sgn\left(\partial UL^I/\partial\Phi^{-1}(PD_t)\right) = sgn\left(\partial UL^I/\partial PD_t\right)$, we show that $\partial UL^I/\partial\Phi^{-1}(PD_t)$ is positive at low PD_t and negative at high PD_t.

First, we record that $\frac{\Phi^{-1}(PD_t)-\rho\Phi^{-1}(\alpha)}{\sqrt{1-\rho^2}} < \Phi^{-1}(1-PD_t) = -\Phi^{-1}(PD_t)$,which can be rewritten as $\Phi^{-1}(PD_t)\frac{1+\sqrt{1-\rho^2}}{\rho} < \Phi^{-1}(\alpha)$. The latter inequality holds for PD_t sufficiently close to α (from above) because $\Phi^{-1}(\alpha) < \Phi^{-1}(PD_t) < 0$ by (13) and $\frac{1+\sqrt{1-\rho^2}}{\rho} > 1$.

In addition, we record that $\frac{\Phi^{-1}(PD_t)-\rho\Phi^{-1}(\alpha)}{\sqrt{1-\rho^2}} > \Phi^{-1}(PD_t)$ because $\Phi^{-1}(\alpha) < 0$ by (13) and $\rho \in (0,1)$.

Combining the two observations, $\frac{\Phi^{-1}(PD_t)-\rho\Phi^{-1}(\alpha)}{\sqrt{1-\rho^2}} \in \left(\Phi^{-1}(PD_t), -\Phi^{-1}(PD_t)\right)$.

By the bell-shape and symmetry properties of the standard normal PDF, ϕ, it then follows that $\phi\left(\frac{\Phi^{-1}(PD_t)-\rho\Phi^{-1}(\alpha)}{\sqrt{1-\rho^2}}\right) > \phi\left(\Phi^{-1}(PD_t)\right)$. Thus, $\frac{\partial UL^I}{\partial\Phi^{-1}(PD_t)} =$
$\phi\left(\frac{\Phi^{-1}(PD_t)-\rho\Phi^{-1}(\alpha)}{\sqrt{1-\rho^2}}\right)\frac{1}{\sqrt{1-\rho^2}} - \phi\left(\Phi^{-1}(PD_t)\right) > 0$ for PD_t that is sufficiently close to α (from above).

As PD_t increases, $\Phi^{-1}(PD_t)\frac{1+\sqrt{1-\rho^2}}{\rho}$ eventually rises above $\Phi^{-1}(\alpha)$ or, equivalently, $\frac{\Phi^{-1}(PD_t)-\rho\Phi^{-1}(\alpha)}{\sqrt{1-\rho^2}}$ rises above $-\Phi^{-1}(PD_t)$. Moreover, due to the exponential function embedded in ϕ, $\phi\left(\frac{\Phi^{-1}(PD_t)-\rho\Phi^{-1}(\alpha)}{\sqrt{1-\rho^2}}\right)\frac{1}{\sqrt{1-\rho^2}} - \phi\left(\Phi^{-1}(PD_t)\right)$ turns negative for a sufficiently high PD_t, implying $\partial UL^I / \partial \Phi^{-1}(PD_t) < 0$.

B.3 Proof of Proposition 1

We discuss only scenario (i) in the proposition. The proof is similar for scenario (ii).

We first prove that $\Phi\left(\frac{\Phi^{-1}(\cdot)-\sqrt{1-\rho^2}\Phi^{-1}(\Lambda)}{\rho}\right)$ is an increasing convex function in the neighbourhood of EL that solves $\Lambda = \Phi\left(\frac{\Phi^{-1}(EL)-\rho\Phi^{-1}(\alpha)}{\sqrt{1-\rho^2}}\right)$. It is immediate that the function is increasing in EL. Its second derivative is equal to $\left(\frac{\sqrt{1-\rho^2}\Phi^{-1}(\Lambda)-\left(1-\rho^2\right)\Phi^{-1}(EL)}{\rho^3}\right)\phi\left(\frac{\Phi^{-1}(EL)-\sqrt{1-\rho^2}\Phi^{-1}(\Lambda)}{\rho}\right)/\phi^2\left(\Phi^{-1}(EL)\right)$, where ϕ is the standard normal PDF and thus $d\phi(x)/dx = -x\phi(x)$. This expression is positive if $\Lambda > EL$, which is the case since $\Lambda = \Phi\left(\frac{\Phi^{-1}(EL)-\rho\Phi^{-1}(\alpha)}{\sqrt{1-\rho^2}}\right)$ is equal to EL at $\rho = 0$ and increases in ρ (see Appendix B.2).

Given that $\Phi\left(\frac{\Phi^{-1}(\cdot)-\sqrt{1-\rho^2}\Phi^{-1}(\Lambda)}{\rho}\right)$ is convex, we know that, for $EL = \pi_t^x PD_t^x + (1-\pi_t^x) PD_t^{\tilde{x}}$, $\Phi\left(\frac{\Phi^{-1}\left(\pi_t^x PD_t^x+(1-\pi_t^x)PD_t^{\tilde{x}}\right)-\sqrt{1-\rho^2}\Phi^{-1}(\Lambda)}{\rho}\right) = \alpha$ $< \pi_t^x\Phi\left(\frac{\Phi^{-1}(PD_t^x)-\sqrt{1-\rho^2}\Phi^{-1}(\Lambda)}{\rho}\right) + (1-\pi_t^x)\Phi\left(\frac{\Phi^{-1}\left(PD_t^{\tilde{x}}\right)-\sqrt{1-\rho^2}\Phi^{-1}(\Lambda)}{\rho}\right)$. Thus, there is a LAR shortfall if the bank maintains its initial LAR in the face of uncertainty. Given that the right-hand side of the latter inequality decreases in Λ, the bank eliminates this shortfall by increasing its LAR. Since the appropriate LAR increases while EL stays constant, the switch to uncertainty raises UL.

B.4 Proof of Proposition 2

To obtain the probability of the naive bank's failure, we substitute $\Phi\left(\frac{\Phi^{-1}\left(\pi_t^l PD_t^l+\left(1-\pi_t^l\right)PD_t^h\right)-\rho\Phi^{-1}(\alpha)}{\sqrt{1-\rho^2}}\right)$ for Λ in (9) and obtain:

$$\pi_t^l\Phi\left(\begin{array}{c}\frac{\Phi^{-1}\left(PD_t^l\right)-\Phi^{-1}\left(\pi_t^l PD_t^l+\left(1-\pi_t^l\right)PD_t^h\right)}{\rho}\\ +\Phi^{-1}(\alpha)\end{array}\right)+\left(1-\pi_t^l\right)\Phi\left(\begin{array}{c}\frac{\Phi^{-1}\left(PD_t^h\right)-\Phi^{-1}\left(\pi_t^l PD_t^l+\left(1-\pi_t^l\right)PD_t^h\right)}{\rho}\\ +\Phi^{-1}(\alpha)\end{array}\right) \tag{15}$$

Next, we show that expression (15) is smaller than $\left(1-\pi_t^l\right)$ for a sufficiently high $\rho \in (0,1)$.

Since $PD_t^l < PD_t^h$, $\Phi^{-1}\left(PD_t^l\right) < \Phi^{-1}\left(\pi_t^l PD_t^l + \left(1-\pi_t^l\right) PD_t^h\right)$ and thus the first summand of expression (15) is smaller than $\pi_t^l \alpha$. Then, since $\alpha < \left(1-\pi_t^l\right)$ by (13), this summand is smaller than $\pi_t^l\left(1-\pi_t^l\right)$.

Turning to the second summand, we derive a condition under which it is smaller than $\left(1-\pi_t^l\right)^2$. For this, we need that $\Phi\left(\frac{\Phi^{-1}\left(PD_t^h\right)-\Phi^{-1}\left(\pi_t^l PD_t^l+\left(1-\pi_t^l\right)PD_t^h\right)}{\rho}+\Phi^{-1}(\alpha)\right) < 1-\pi_t^l$, which is equivalent to $\Phi^{-1}\left(1-\pi_t^l\right)-\Phi^{-1}(\alpha) > \frac{\Phi^{-1}\left(PD_t^h\right)-\Phi^{-1}\left(\pi_t^l PD_t^l+\left(1-\pi_t^l\right)PD_t^h\right)}{\rho}$. This inequality holds if $\Phi^{-1}\left(1-\pi_t^l\right)-\Phi^{-1}(\alpha) > \frac{\Phi^{-1}\left(PD_t^h\right)-\pi_t^l\Phi^{-1}\left(PD_t^l\right)-\left(1-\pi_t^l\right)\Phi^{-1}\left(PD_t^h\right)}{\rho} = \frac{\pi_t^l}{\rho}\left(\Phi^{-1}\left(PD_t^h\right)-\Phi^{-1}\left(PD_t^l\right)\right)$, where the latter inequality stems from $\pi_t^l PD_t^l + \left(1-\pi_t^l\right)PD_t^h < 0.5$, by (13), which implies that $\Phi^{-1}\left(\pi_t^l PD_t^l + \left(1-\pi_t^l\right) PD_t^h\right)$ corresponds to the concave portion of Φ^{-1}.

Putting the arguments on the two summands together, we conclude that $\frac{\Phi^{-1}\left(1-\pi_t^l\right)-\Phi^{-1}(\alpha)}{\Phi^{-1}\left(PD_t^h\right)-\Phi^{-1}\left(PD_t^l\right)} > \frac{\pi_t^l}{\rho}$ is a sufficient condition for expression (15) to be smaller than $\left(1-\pi_t^l\right)$.

Referring again to (13), we obtain $\frac{\Phi^{-1}\left(1-\pi_t^l\right)-\Phi^{-1}(\alpha)}{\Phi^{-1}\left(PD_t^h\right)-\Phi^{-1}\left(PD_t^l\right)} > 1$, which implies that the above sufficient condition is satisfied at a sufficiently high ρ.

Thus, there exists $\bar{\rho} < 1$ such that expression (15) is smaller than $\left(1-\pi_t^l\right)$ for $\rho \in (\bar{\rho}, 1)$.

Then, we observe that – because $\Phi^{-1}\left(PD_t^l\right) < \Phi^{-1}\left(\pi_t^l PD_t^l + \left(1-\pi_t^l\right) PD_t^h\right) < \Phi^{-1}\left(PD_t^h\right)$ – expression (15) converges to $\left(1-\pi_t^l\right)$ as $\rho \to 0$.

Finally, by continuity, we conclude that there exists $\underline{\rho} > 0$ such that the bank's failure probability is smaller under $\rho \in (\bar{\rho}, 1)$ than under $\rho \in \left(0, \underline{\rho}\right)$.

B.5 Proof of Proposition 3

When the uninformed bank's LAR is equal to Λ^U and the phase implies PD_t, the informed bank perceives it as failing in year t with probability $FP_t^U = \Phi\left(\frac{\Phi^{-1}(PD_t)-\sqrt{1-\rho^2}\Phi^{-1}\left(\Lambda^U\right)}{\rho}\right)$, which implies $\Lambda^U = \Phi\left(\frac{\Phi^{-1}(PD_t)-\rho\Phi^{-1}\left(FP_t^U\right)}{\sqrt{1-\rho^2}}\right)$. Using this to substitute for Λ in equation (5) while imposing $x = l$ and $\tilde{x} = h$ and rearranging, we obtain:

$$\frac{\Phi^{-1}\left(PD_t^h\right)-\Phi^{-1}\left(PD_{t-1}^l\right)}{\rho} = \Phi^{-1}\left(FP_t^U\right) - \Phi^{-1}\left(\frac{\alpha-\left(1-\pi_t^l\right)FP_t^U}{\pi_t^l}\right),$$

where the left-hand side is strictly positive and decreases in ρ, and the right-hand side increases in FP_t^U. This implies $dFP_t^U/d\rho < 0$, as stated in the proposition.

B.6 Proof of Proposition 4

Suppose that the economy undergoes a low-to-high switch – implying PD_{t-1}^l and PD_t^h. As $\rho \to 0$, the informed bank perceives $UL^I \to 0$. By contrast, when $\rho \to 0$, the uninformed bank perceives that date-t losses will be higher than the expected ones, EL^U, with probability $1 - \pi_t^l$. Thus, as long as $1 - \pi_t^l \geq \alpha$ – i.e. as long as the perceived probability of switching phases is higher than the bank's targeted probability of failure – $\lim_{\rho \to 0} UL_t^U > 0$. By continuity, this implies that there exists $\underline{\rho} > 0$ such that $UL_{t-1}^U > UL_t^I$ when $\rho \in \left(0, \underline{\rho}\right)$. That is, that the uninformed bank is over-capitalized from the informed bank's perspective. This proves the proposition.

C "Regulatory" bank

The naive bank in the main text shares important similarities with a bank that sets its LAR by relying only on the model underpinning international regulatory requirements for credit risk (BCBS (2017)). The latter bank would also assume (wrongly) that it has overcome the uncertainty about phase switches and has thus pinned down the relevant loan PD at the beginning of period t. Moreover, this "regulatory" bank would map its PD estimate, PD_t^R, into the loading on the default-clustering factor:

$$\rho_t^2 \equiv \rho^2(PD_t^R; \tilde{\rho}, \tilde{\tilde{\rho}}, S) = \tilde{\rho}^2 \frac{1 - e^{-S*PD_t^R}}{1 - e^{-S}} + \tilde{\tilde{\rho}}^2 \left(1 - \frac{1 - e^{-S*PD_t^R}}{1 - e^{-S}}\right), \tag{16}$$

where $\tilde{\rho} < \tilde{\tilde{\rho}}$ denote the lower and upper limits of the loading and the parameter S determines the speed at which this loading declines from the latter to the former as PD_t^R rises. BCBS (2005) states on page 12 that the assumption of a negative relationship between PD_t^R and ρ_t reflects "empirical analysis and intuition" that higher credit risk stems largely from idiosyncratic risk factors. Given this assumption, regulation requires that banks estimate only one parameter per credit-risk exposure: the probability of default (see Section 2.3).

The other, regulatory parameters in (16) differ across exposure types (BCBS (2017)). For corporate exposures, $\tilde{\rho}^2 = 12\%$, $\tilde{\tilde{\rho}}^2 = 24\%$ and $S = 50$; for residential mortgages,

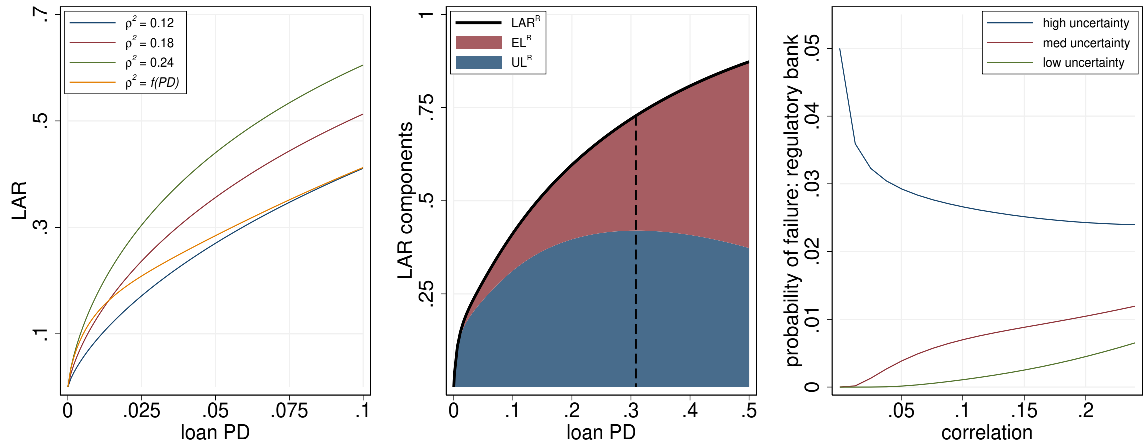

Figure C.1: ***Regulatory assumptions and implications of uncertainty.*** Loss-absorbing resources (LAR) and their components – expected loss (EL) and unexpected loss (UL) – from the perspective of the regulatory bank when the target one-year probability of bank failure is 0.1%. The dashed line is plotted at the maximum UL. In the left-hand panel the regulatory bank's LAR is compared to that of the informed bank, assuming that the latter perceives three different levels of ρ^2. For the right hand panel, it is assumed that the true level ρ^2 is as reported on the horizontal axis; $\pi_t^l = 0.95$; and the naive bank's probability of failure is from the perspective of the uninformed bank, for different levels of uncertainty, $(PD^h – PD^l) = 5\%$ (low), 10% (medium) and 20% (high), while the expected loan PD is always 2%.

$\tilde{\rho}^2 = \tilde{\tilde{\rho}}^2 = 15\%$; and, for "other retail exposures", $\tilde{\rho}^2 = 3\%$, $\tilde{\tilde{\rho}}^2 = 16\%$ and $S = 35$.

Acting like the naive bank, the regulatory bank sets LAR to:

$$\Lambda^R \left(PD_t^R; \alpha\right) \equiv \Lambda_t^R = \Phi \left(\frac{\Phi^{-1}\left(PD_t^R\right) - \rho\left(PD_t^R\right)\Phi^{-1}(\alpha)}{\sqrt{1-\rho^2\left(PD_t^R\right)}} \right). \tag{17}$$

Assuming that the regulatory bank has the same information set as the naive one, the underlying EL and UL are

$$EL_t^R = PD_t^R = \pi_t^x PD_t^x + (1-\pi_t^x) PD_t^{\tilde{x}}, \tag{18}$$

$$UL^U\left(PD_t^R; \alpha\right) \equiv UL_t^R = \Lambda_t^R - EL_t^R. \tag{19}$$

Figure C.1 reports properties of the regulatory bank's LAR using the parameterisation for corporate exposures. For one, the mapping in (16) implies a LAR that increases less strongly with loan PD than the LAR of the informed bank (left-hand panel). In addition, the regulatory bank perceives EL and UL that are joined at the hip over a realistic range of loan PDs (centre panel, recall Section 3.4.1).

Similar to the naive bank, the regulatory bank's LAR will not attain the target failure probability of α. To illustrate that, from the perspective of the uninformed bank, the regulatory bank could be safer or riskier than targeted, Figure C.1 plots its failure probability for a given PD_t^R, different levels of uncertainty, and different values of the actual ρ. To see the underlying mechanism, note first that, given PD_t^R, the LAR of the regulatory bank is constant. Then refer to Figure 5, which shows that, as $\rho \to 0$, the largest possible loss rate converges to PD_t^h. When the uncertainty stems from the difference $\left(PD_t^h - PD_t^l\right)$ – as assumed in Figure C.1 – and is sufficiently low, the level of PD_t^h would be below the bank's constant LAR. In this case, the failure probability is 0 in the limit $\rho \to 0$ (red and green lines). Conversely, when the uncertainty is sufficiently high, PD_t^h would be above the regulatory bank's LAR and this bank's failure probability will converge to one minus the continuation probability of the low-PD phase, $1 - \pi_t^l$, as $\rho \to 0$ (blue line).