BANK FOR INTERNATIONAL SETTLEMENTS

BIS Working Papers
No 913

Forecasting expected and unexpected losses

by Mikael Juselius and Nikola Tarashev

Monetary and Economic Department

December 2020

JEL classification: G17, G21, G28.

Keywords: Loss rate forecasts, cyclical turning points, expected loss provisioning, bank capital.

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Forecasting expected and unexpected losses*

Mikael Juselius (*Bank of Finland*)
Nikola Tarashev (*Bank for International Settlements*)

December 17, 2020

Abstract

Extending a standard credit-risk model illustrates that a single factor can drive both expected losses and the extent to which they may be exceeded in extreme scenarios, ie "unexpected losses." This leads us to develop a framework for forecasting these losses jointly. In an application to quarterly US data on loan charge-offs from 1985 to 2019, we find that financial-cycle indicators – notably, the debt service ratio and credit-to-GDP gap – deliver reliable real-time forecasts, signalling turning points up to three years in advance. Provisions and capital that reflect such forecasts would help reduce the procyclicality of banks' loss-absorbing resources.

JEL Codes: G17; G21; G28

Keywords: Loss rate forecasts; Cyclical turning points; Expected loss provisioning; Bank capital

*This paper reflects the views of its authors, not necessarily the views of the Bank of Finland or the Bank for International Settlements. The authors would like to thank Stijn Claessens, Marco Lombardi and seminar participants at the Bank of Finland and the BIS for valuable suggestions, and Anamaria Illes for help with the data.

1 Introduction

In seeking to align banks' loss-absorbing resources with *actual* risks, authorities have introduced forward-looking elements in accounting and prudential standards. These standards target complementary aspects of the time-varying probability distribution of future losses. Namely, IASB (2014) and FASB (2016) seek to align provisions, or the amount of assets that banks write down, with expected losses (EL). In turn, BCBS (2017) sets regulatory capital for "unexpected" losses (UL), or the extent to which EL may be exceeded in extreme scenarios.

Even though EL and UL are in principle intertwined, the literature has studied them separately. It has been mostly pessimistic about the cyclical properties of EL and has addressed time variation in UL only indirectly, through attempts to identify early warning indicators of banking crises.[1] In this paper, we propose a unified, parsimonious and head-on treatment of EL and UL, which delivers accurate real-time (that is, out-of-sample) forecasts.

Theory. To motivate our empirical analysis, we extend the theoretical model behind credit-risk requirements in Basel III (Gordy, 2003, Vasicek, 1991) by incorporating time variation in the risk distribution. In this extension, the target level of a bank's loss-absorbing resources depends on two systematic factors: an *unforecastable* factor that drives the clustering of loan defaults and a *forecastable* "phase switching" factor that drives loss-rate swings in boom-bust transitions. Our model reveals that improved forecasts of the phase-switching factor should lead to adjustments to both EL and UL estimates. Depending on the strength of the default-clustering factor, these adjustments could be in opposite directions. And to the extent that time variation in EL and UL has a common driver, forecasting either aspect of the loss distribution would carry information about the other one. This prompts us to forecast EL and UL jointly.

The **empirical framework**, our main contribution, has two components. The first is a model for forecasting jointly and in *real-time* the level of *quarterly* loss rates (first moments of the loss distribution) and the attendant squared forecast errors (second moments). The second component is a methodology for aggregating the first- and second-moment forecasts in order to derive EL and UL over the *lifetime* of a bank's

[1]EL forecasts have been discussed in Covas and Nelson (2018), Abad and Suárez (2017), Goncharenko and Rauf (2020), Chae et al (2018), Lu and Nikolaev (2020), and Loudis and Ranish (2019). On early warning indicators of banking crises, see Detken et al (2014), Tölö et al (2018), Aldasoro et al (2018) and references therein.

loan portfolio.[2] We apply this framework to data from 1985q1 to 2019q2 on the US banking sector's quarterly loss rates, as measured by net charge-off rates on loans to the private non-financial sector.

The first component of our framework comprises direct forecasting models for the first and second moments.[3] In these models, we employ forecast variables that have been used extensively in the related literature (unemployment rate, output gap, credit spread, term spread, real interest rate, and the S&P 100 volatility index) or have had a good track record as early warning indicators of system-wide banking crises (debt-service ratio, the credit-to-GDP gap, and the property-price gap). We start by estimating the model on a "training sample" from 1985q1 to 1999q4, using the estimates to forecast the two moments in each of the twelve subsequent quarters, ending with 2002q4. This is in line with regulatory texts and banking industry practice (Loudis and Ranish, 2019), which indicate three years as the longest "reasonable and supportable" forecast horizon. We then proceed to construct similar out-of-sample forecasts by sequentially adding one quarter to the sample.

Our preferred forecasts stem from the joint estimation of a linear specification for the first moments, employing only the debt service ratio (DSR), and a multiplicative heteroscedastic specification for the second moments, employing only the credit-to-GDP gap (C2Y). These forecasts are better than alternatives reflecting multiple forecast variables or (more elaborate) non-linear specifications, or separate forecasts of the first and second moments. We reach this conclusion on the basis of four assessment metrics, evaluated at each forecast horizon: root mean squared errors (RMSE), the correlation between forecasted and actual series, the number of excess/insufficient turning points in the time series of forecasts relative to the actual time series, and the time distance between forecasted and actual turning points.

Some examples of the assessment results are as follows. At the 12-quarter horizon, the correlation metric for DSR-based first-moment forecasts is at around 60%, at least 15 percentage points higher than the corresponding metrics underpinned by other forecast variables. And the former metric increases by about 30 percentage points when the first-moment forecasts are derived jointly with C2Y-based second-moment forecasts. In turn,

[2] This is a departure from the standard definition of UL, which calls for estimating a high percentile of the lifetime loss distribution (Gordy, 2003). Mindful of this, we forecast such a percentile in *Appendix D* on the basis of quantile regressions. However, we do not obtain meaningful results.

[3] Direct forecasting is discussed by eg Marcellino et al (2006), and the multiplicative heteroscedastic specification, pioneered by Harvey (1976), has recently been applied by eg Adrian et al (2019) in the context of GDP growth.

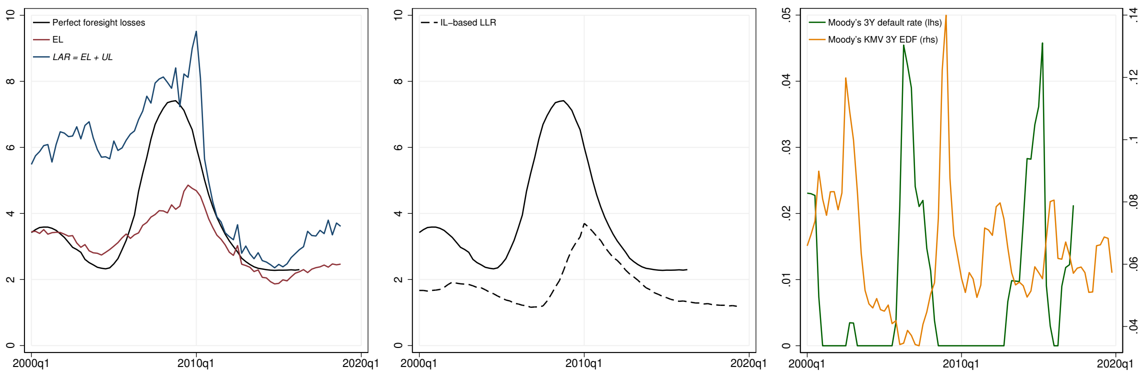

Figure 1: **Catching or missing turning points.** An observation in quarter T is equal to: a real-time forecast at T (expected lifetime losses (EL), unexpected lifetime losses (UL) set to 2 forecasted standard deviations, Moody's KMV 3-year expected default frequency (EDF)), or the corresponding materialization of risk from T onward (perfect-foresight losses, Moody's 3-year default rate), or historical loan-loss provisions (LLR) at T, based on incurred loss (IL) accounting. EL + UL = loss absorbing resources (LAR). The Moody's and Moody's KMV series refer to US high-yield non-financial corporates (NFC). A high-yield three-year default rate is defined as the share of NFCs that are rated C or CC in quarter $T + 12$ after being rated between BB+ and CCC- in quarter T.

the correlation metric for these second-moment forecasts is above 20% and consistently positive across horizons, in contrast to corresponding forecasts that condition on other variables. Last but not last, our preferred forecasts stand out by capturing exactly the number of turning points at effectively all horizons and by pinning down the timing of turning points, with errors of three or fewer quarters.

In the second part of our empirical exercise, we aggregate the forecasts of *quarterly* losses over the loan portfolio's *lifetime*. Following regulatory practice and the literature (FASB, 2016, Covas and Nelson, 2018, and Loudis and Ranish, 2019), we make standard simplifying assumptions regarding the portfolio's rundown structure, as well as the behavior of losses beyond the maximum forecast horizon. These assumptions and our first-moment forecasts deliver directly forecasts of the portfolio's EL.

The aggregation of the forecasts of quarterly second moments is more challenging. It rests on an additional set of identification restrictions (which we confirm subsequently in the data) and delivers an estimate of the standard deviation of the portfolio's lifetime losses. In turn, a scalar multiple of this standard deviation results in a UL forecast, with the magnitude of the scalar being at the bank's or its supervisor's discretion.

The EL and UL forecasts fare quite well if based on our preferred first- and second-moment specifications. While EL accounts for only 60% of the perfect-foresight lifetime loss rates at the height of the great financial crisis (GFC), it matches the time profile of these loss rates throughout the sample, pinning down each turning point (*Figure 1, left-hand panel*). By contrast, historical loan-loss reserves – based on incurred loss (IL) accounting – signal turning points with a significant delay (*centre panel*). In turn, UL tracks closely the time profile of the difference between perfect-foresight losses and EL, especially starting from 2004q3. Setting UL to two times the forecasted standard deviations of lifetime loss rates – akin to targeting the 98^{th} percentile of a normal distribution – results in the sum of EL and UL – a proxy for a bank's loss-absorbing resources – covering actual loss rates in each quarter of the sample (*left-hand panel*).

Since the success of forecasts hinges on the information content in the estimation sample, we rerun the above analysis for sub-portfolios, whose time profiles differ from that of the aggregate portfolio. Indeed the sub-portfolio comprising commercial-and-industrial (CI) loans experiences heavy losses during the recession in the early 1990's, which do not come across as sharply in the "training" sample at the aggregate level but contain useful information for subsequent forecasts. As a result, the subsequent forecasts for CI loans are more accurate in some respects than those for the aggregate portfolio, matching better the upswing during the GFC.

Literature. Our paper relates directly to the literature on loss-rate forecasts, which has been spurred recently by the introduction of EL provisioning standards for banks in 2018. A number of papers have found that forecasts become unreliable beyond a rather short horizon, of about 0-to-2 quarters (Covas and Nelson, 2018, Krüger et al, 2018, and Abad and Suárez, 2017). Likewise, Moody's KMV expected default frequencies (EDFs) reveal that this problem also permeates proprietary loss forecasts that rely on stock market information (*Figure 1, right-hand panel*). Such EL forecasts are naturally interpreted as undermining financial stability by leading to procyclical provisioning: the under-provisioning in tranquil times would underpin excessive leverage and under-pricing of risks that fuels credit booms, while the over-provisioning in recessions would amplify busts. Moreover, the above "sceptical" literature argues that EL provisioning would be more destabilizing than the unquestionably procyclical IL provisioning.[4]

By contrast, we show that it is possible to derive accurate EL forecasts up to 12

[4] For similar arguments, see Goncharenko and Rauf (2020) and Risk.net (2018a and b). For analyses of the procyclicality of IL provisioning, see Borio and Lowe (2001) and, more recently, Huizinga and Laeven (2019). For how the Spanish "dynamic provisioning" addressed this procyclicality pre-GFC, see Jiménez et al (2017). The post-GFC introduction of EL provisioning on an international scale seeks to improve the timeliness of banks' provisions (Cohen and Edwards, 2017).

quarters in advance. The essential elements that distinguish our empirical framework from that of the "sceptical" literature and deliver reliable forecasts are: first, using forecast variables with a good track record as early warning indicators of banking crises and, second, forecasting first and second moments jointly. Moreover, we also show that our forecasts are countercyclical not only in the context of the credit-risk cycle but also in that of the business cycle. While this result arises on the basis of *real-time* EL and UL forecasts – reflecting joint one-step estimates of the first and second moments of loss rates – Chae et al (2018) reach a similar conclusion for EL only when there is perfect foresight about a macroeconomic variable (a house-price index) with information content about loss rates. In turn, Lu and Nikolaev (2020) construct accurate out-of-sample EL estimates from a large set of forecasting variables.

To the best of our knowledge, we propose the first UL forecasts for a loan portfolio. Our exercise is inspired by the literature on early warning indicators of banking crises (Detken et al, 2014, Tölö et al, 2018, and Aldasoro et al, 2018). By identifying such indicators of reasonable reliability, this literature provided the intellectual impetus for the countercyclical capital buffer in bank regulation (BCBS, 2017). Our results show that some of these indicators underpin the success of *direct* UL forecasts.

Should banks be expected to develop and use accurate long-horizon forecasts of loss rates, such as the ones in this paper? Literature on the inherent procyclicality of the financial system suggests that private incentives probably get in the way (Borio and Zhu, 2008, and Gorton and Ordoñez, 2014). It is in banks' interest to ride on and contribute to credit booms for as long as these last and then pass on losses to the public purse in system-wide busts (Acharya, 2009). Thus, even if banks have accurate long-horizon forecasts at their disposal, public policy measures may be needed to ensure that such forecasts actually underpin loss-absorbing resources.

Roadmap. Section 2 fixes ideas with a simple theoretical setup that delivers time-varying expected and unexpected losses and allows for studying attendant forecasts. Section 3 develops the empirical framework. Then Section 4 applies this framework to US data, assesses the accuracy of the forecasts and relates them to business-cycle indicators. Section 5 concludes. The appendices contain proofs of theoretical results and technical aspects of the empirical strategy.

2 Theoretical model

We consider two banks with loan exposures. Each bank uses all its available information to derive expected losses (EL) and unexpected losses (UL), or the difference between the $(1-\alpha)$ quantile of the loss distribution and EL. And it sets provisions and capital equal to the perceived EL and UL, respectively. Since the sum of provisions and capital amounts to loss absorbing resources (LAR), each bank *perceives* its own probability of default (PD) as equal to α.[5] While the above elements match the model underpinning bank regulation, we extend this model by introducing time variation in the probability distribution of losses. And we assume that, in case a bank defaults, an identical bank replaces it, facing the same loss distribution that the defaulted bank would have faced. In this setting, we study the interplay of perceived EL and UL, both over time and across different information sets.

In each period t, the loan portfolio is asymptotic and composed of homogeneous loans, with each loan subject to two systematic and one idiosyncratic risk factors. It is *asymptotic* in the sense that the number of constituent loans approaches infinity, $n \to \infty$. The loans are *homogeneous* in the sense that each one is of size $1/n$ and has the same probability of default (PD), which can vary over time. A borrower j defaults on its loan in period t if the realisation of its stochastic assets, $\rho(I_t) G_t + \sqrt{1-\rho(I_t)^2} Z_{j,t}$, is below a threshold $d(I_t)$. The *systematic* and *idiosyncratic* risk factors G_t and $Z_{j,t}$ are standard normal variables that are mutually independent and i.i.d. over time. Since the realization of G_t influences the likelihood of joint loan defaults, we refer to it as the "default clustering" factor. The other *systematic* factor, I_t, influences the probability of individual default, through $d(I_t)$, and the loading on the default-clustering factor, $\rho(I_t) \in (0,1)$. It is binary, capturing turning points between boom ($I_t = 0$) and bust phases ($I_t = 1$) of the financial cycle. We will refer to I_t as the "phase switching" factor. For parsimony, we set loss-given-default to 100% for all borrowers.

Since the effect of idiosyncratic factors washes out in the aggregate, portfolio-level

[5] For our purposes, it does not matter whether α is set strategically by the bank or stems from regulatory requirements. The internationally agreed regulatory standards, Basel III (BCBS, 2017), target explicitly $\alpha = 0.001$ over a one-year horizon under the internal ratings-based approach for credit risk.

losses depend only on the systematic factors:[6]

$$\Lambda(G_t, I_t) = \Phi\left(\frac{d(I_t) - \rho(I_t) G_t}{\sqrt{1 - \rho(I_t)^2}}\right), \tag{1}$$

where Φ stands for the standard normal CDF. And since a portfolio's size is 1, portfolio losses are equal to the loan-loss rate. If a bust is to generate higher losses for any realisation of G_t, then $d(0) < d(1)$ or $\rho(0) < \rho(1)$.

2.1 Alternative information settings

In order to zoom in on the implications of forecasting the financial cycle, we assume that I_t is inherently forecastable, whereas G_t is not. We then differentiate the banks in terms of their capacity to forecast I_t. Thus, for instance, anticipating a bust would translate into higher EL and UL because of adjustments to the perceived $d(I_t)$ and $\rho(I_t)$, as opposed to adjustments related to G_t.

Concretely, the two banks differ in terms of the information they have when setting LAR at the beginning of any period t. Each information set comprises full knowledge of the model parameters and the stochastic properties of the systematic factors, $\{G_t, I_t\}$. The difference stems from what banks know about the phase in period t. One of the banks sets LAR for period t on the basis of complete knowledge (CK) of the phase $I_t = \iota_t$. This is the CK bank. The other bank's date-t information set contains only past realizations of the risk factors, of which only the latest phase, ι_{t-1}, has useful content. From the perspective of this "incomplete knowledge" – or IK – bank, the likelihood that phase ι_{t-1} prevails in t is equal to $\pi(\iota_{t-1}) \in (0, 1)$.[7]

[6]Gordy (2003) provides the theoretical underpinnings of the asymptotic single risk-factor model, which features only G_t and $Z_{j,t}$. That paper imposes weak conditions on the shape, continuity and differentiability of these factors' distribution functions and delivers an expression similar to (1).

[7]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). While a more complex perceived distribution of the loss rate would emerge from assuming that the only information about the phase stems from the history of loan losses, this would not alter the qualitative implications of different knowledge about the phase.

2.1.1 Complete knowledge

A CK bank conditions on ι_t and sets LAR to cover the realisation of losses when G_t is at its α-percentile:

$$\Lambda^{\alpha,CK}(\iota_t) = \Phi\left(\frac{d(\iota_t) - \rho(\iota_t)\,\Phi^{-1}(\alpha)}{\sqrt{1-\rho(\iota_t)^2}}\right). \tag{2}$$

In turn, the expected loss rate is equal to the probability of loan default:

$$EL^{CK}(\iota_t) = \Phi(d(\iota_t)). \tag{3}$$

Finally, the unexpected loss rate perceived by the CK bank is $UL^{\alpha,CK}(\iota_t) = \Lambda^{\alpha,CK}(\iota_t) - EL^{CK}(\iota_t)$.

2.1.2 Incomplete knowledge

An IK bank's LAR, $\Lambda^{\alpha,IK}(\iota_{t-1})$, is the implicit solution of the following equation in terms of Λ:

$$\begin{aligned} &\pi(\iota_{t-1})\,\Phi\left(\frac{d(\iota_{t-1}) - \sqrt{1-\rho(\iota_{t-1})^2}\,\Phi^{-1}(\Lambda)}{\rho(\iota_{t-1})}\right) \\ &+ (1-\pi(\iota_{t-1}))\,\Phi\left(\frac{d\left(\tilde{\iota_{t-1}}\right) - \sqrt{1-\rho\left(\tilde{\iota_{t-1}}\right)^2}\,\Phi^{-1}(\Lambda)}{\rho\left(\tilde{\iota_{t-1}}\right)}\right) = \alpha. \end{aligned} \tag{4}$$

where $\tilde{\iota_{t-1}}$ is the value of I that did *not* materialize in period $t-1$. The left-hand side of this expression is equal to the probability of losses exceeding Λ, as perceived by the IK bank (*Appendix A*).

Further, the IK bank perceives an expected loss rate equal to

$$EL^{IK}(\iota_{t-1}) = \pi(\iota_{t-1})\,\Phi(d(\iota_t)) + (1-\pi(\iota_{t-1}))\,\Phi\left(d\left(\tilde{\iota_t}\right)\right) \tag{5}$$

and an unexpected loss rate equal to $UL^{\alpha,IK}(\iota_{t-1}) = \Lambda^{\alpha,IK}(\iota_{t-1}) - EL^{IK}(\iota_{t-1})$.

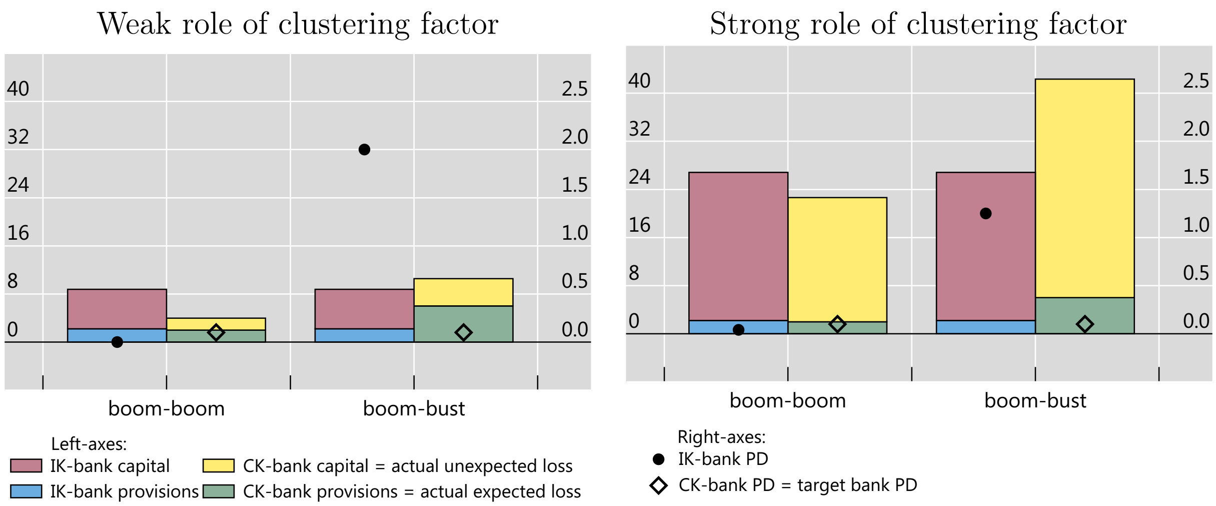

Figure 2: **Implications of missing a turning point.** The two states under each set of bars correspond to ι_{t-1} and ι_t, respectively. A bank's PD refers to ι_t. The underlying parametrization is: $pd(0) = 2\%, pd(1) = 6\%, \alpha = 0.1\%, \pi(0) = 95\%, \pi(1) = 66\%$, $\rho = \sqrt{0.01}$ (left-hand panel) and $\rho = \sqrt{0.2}$ (right-hand panel).

2.2 Failing to forecast phase switches

To illustrate consequences of failing to foresee a switch from one financial-cycle phase to another, we compare the two banks' PDs from the perspective of the CK bank, that is from the perspective of the best attainable information set. For concreteness, we focus on a scenario where the latest phase is a boom, i.e. $\iota_{t-1} = 0$. The implications are symmetric for $\iota_{t-1} = 1$. The underlying proofs are in *Appendix A*.

While the CK bank's PD is always α, the IK bank's can be lower or higher (*Figure 2*). Concretely, the latter PD is below target if the boom phase continues (boom-boom scenarios: dots below diamonds) and above target if there is a switch to a bust (boom-bust scenarios). This reflects LAR excesses, respectively shortfalls. The shortfalls are largest when the boom-to-bust switch entails both higher probabilities of loan defaults and higher loadings on the default-clustering factor, $\rho\left(\iota\right)$ (right-hand panel).

More interestingly, an IK bank with a *lower* through-the-cycle loading on the default-clustering factor has a *higher* PD. To derive this result, we set $\rho\left(0\right) = \rho\left(1\right) = \rho$ and conduct comparative statics with respect to ρ. The result is illustrated in *Figure 2*, as the boom-bust dot in the left-hand panel is above that in the right-hand panel. The underlying intuition is as follows. A lower ρ implies that phase-contingent losses are more certain, i.e. that $UL^{\alpha,CK}$ and $UL^{\alpha,IK}$ 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 the unforecastable default-clustering factor (G_t) would undo the implications of missing a phase switch (I_t). Thus, when the phase-contingent losses are more certain, it is also more certain than missing a phase switch would be detrimental.

The literature has established some stylized facts on how the importance of the clustering factor changes with the portfolio type. This importance surfaces in the data as correlations of borrowers' asset returns.[8] Düllmann et al (2007), Zhang et al (2008) and Bams et al (2012) measure such correlations and deliver the following takeaway: the importance of the clustering factor is lower in portfolios comprising exposures to smaller or riskier entities. Given this, our theoretical model implies that – all else the same – IK banks holding such portfolios would have a larger LAR shortfall in a boom-to-bust switch.

2.3 Theoretical motivation of the empirical analysis[9]

Despite its extremely stylized nature, the theoretical model captures salient features of both credit loss rates and existing real-time forecasts of these rates. Similar to *Figure 2*, real-world loss rates tend to undergo distinct boom-bust phases, with abrupt switches from low to materially higher levels and vice-versa (*Figure 1*, left-hand panel, solid lines). In turn, consistent with the message of *Figure 1* (right-hand panel), much of the literature has delivered real-time forecasts that miss the turning points of phase switches (Covas and Nelson, 2018, Krüger et al, 2018, and Abad and Suárez, 2017). Reminiscent of the IK bank, these forecasts result in LAR shortfalls when banks enter a bust phase and excesses when they exit this phase.

Seen through the prism of the theoretical model, a key objective of our empirical exercise below is to show that it is possible to enrich the information set of the IK bank in order to improve its forecasts of phase switches. Such an improvement – based on information in addition to the latest observed phase ι_{t-1} – would imply that the IK bank would *perceive* the probability of remaining in a boom as evolving over time and increasing in the run-up to a switch. Symmetrically for the probability of remaining in a bust. Effectively, the IK bank's perceptions of both EL^{IK} and $UL^{a,IK}$ – and the corresponding provisions and capital – would line up more closely with those of the CK bank, EL^{CK} and $UL^{a,CK}$.

[8] In contrast to this literature, which employs borrower-level evidence, the data we employ below is at the portfolio level and do not allow us to derive the relative role of different factors.

[9] See *Appendix A* for proofs of statements in this subsection.

The direction of such an alignment may differ between EL and UL, depending on the loan portfolio. Granted, an improvement of the IK bank's forecasts is sure to increase the perceived EL prior to a boom-to-bust switch (*Figure 2*, either panel, provisions in boom-bust scenario). However, the IK bank may adjust its perception of UL either up or down. The reason is that, from the perspective of this bank, UL is driven not only by the unforecastable default-clustering factor G_t but also by uncertainty as regards the phase-switching factor I_t. And when the latter uncertainty is sufficiently high relative to the loading on G_t, the IK bank would perceive a *higher* pre-bust UL than the CK bank (*Figure 2*, left-hand panel, capital in boom-bust scenario). From such an initial condition, aligning the IK bank's perceptions with those of the CK bank entails a decline in $UL^{a,IK}$. Conversely, if the loading on G_t is sufficiently high relative to the uncertainty about I_t, then $UL^{a,IK}$ increases in order to approach $UL^{a,CK}$ (right-hand panel, capital in boom-bust scenario). All this underscores that an improvement to the phase-switching forecasts needs to be applied consistently to EL and UL.

The ultimate reason for forecasting EL and UL jointly stems directly from equation (1), which indicates that realized losses (Λ_t) are the joint outcome of the default clustering factor (G_t) and the phase-switching factor (I_t). This suggests that an accurate forecast of I_t (and thus EL) is essential for deriving reliable estimates of the real-world distribution of G_t (and thus UL) on the basis of loss data. But it also suggests that the relationship runs in the other direction as well: UL forecasts would determine which EL forecasts are consistent with observed loss rates.

3 Framework for loss forecasts and their evaluation

The empirical framework for estimating a loan portfolio's loss distribution bears similarities and differences with the stylized theoretical model in the preceding section. Similar to that model, we will break down the estimates of LAR into EL and UL estimates. For the former, we will forecast the level of loss rates at different horizons and – using standard assumptions about the portfolio's amortization – will combine these first-moment forecasts into an estimate of the cumulative EL over the life of the portfolio. Importantly, the first-moment exercise will provide essential input into our forecasts of horizon-specific squared forecast errors (second moments). Prompted by our theoretical model, we will verify whether there are benefits in the other direction: whether the joint estimation of first and second moments improves the quality of the former. Then, relying on identification conditions, we will use the second-moment forecasts to construct the portfolio's lifetime cumulative UL. Given the pronounced boom-bust pattern of realized loss rates, we will pay close attention to the extent to which our forecasts match turning points in the data.

The empirical object of ultimate interest is the lifetime loss rate, $\mathit{\Lambda}_t$. On each date t, $\mathit{\Lambda}_t$ depends on: the length of the portfolio's (remaining) life, M, i.e. the number of periods (henceforth, quarters) it takes for the principal, L_t, to reach zero; the quarterly loss rates, μ_{t+h}, taking place h quarters into the future, where $h \in \{1, ..., M\}$; and the remaining principals at the *beginning* of the corresponding quarters, L_{t+h-1}. Concretely, $\mathit{\Lambda}_t = (1/L_t)\sum_{h=1}^{M} L_{t+h-1}\mu_{t+h}$.

3.1 First moments: forecasting expected losses

The goal of our first-moment forecasts is to generate the expected lifetime loss rate, $\mathit{\Lambda}_t^e$. To this end, we follow the spirit of FASB (2016) and implement two simplifications.[10] First, we assume a deterministic linear rundown of the portfolio, $l_{t+h} \equiv L_{t+h}/L_t = (1 - h/M)$, which implies that

$$\mathit{\Lambda}_t^e = \sum_{h=1}^{M} \left(1 - \frac{h-1}{M}\right) \mu_{t+h|t}^e \tag{6}$$

where $\mu_{t+h|t}^e$ is the date-t expectation of the quarterly loss rate materialising h quarters into the future.[11] Second, we recognise that the maximum horizon at which the data allow for meaningful forecasts – i.e. $\hat{\mu}_{t+h|t}$, where $h = 1, ..., H$ – would typically be shorter than the time to maturity: $H < M$. Beyond H, we assume that the loss rates converge linearly and deterministically to their long-term average, $\overline{\mu}$, over N quarters and remain at this level thereafter. Concretely, this implies the following structure for

[10] See in particular paragraphs 326-20-30-8 and 326-20-30-9.

[11] Equation (6) incorporates two further simplifications. First, it does not include discounting, as permitted by FASB (2016), paragraph 326-20-30-3. This simplification allows for a more transparent comparison between forecasted and realized lifetime loss rates. Second, we abstract from the fact that the run-down of the portfolio would reflect defaults. One way of modelling this would have been to set $l_{t+h} = (1 - h/M)\prod_{i=1}^{h}(1 - \mu_{t+i})$. This would have implied that Λ_t is a non-linear function of quarterly loss rates, μ_{t+h}, thus overburdening the forecasting exercise with higher moments that are of second-order in the context of equation (6).

the expected quarterly loss rates:

$$\mu^e_{t+h|t} = \begin{cases} \hat{\mu}_{t+h|t} & \text{for } h = 1, ..., H \\ \left(1 - \frac{h-H}{N}\right) \hat{\mu}_{t+H|t} + \left(\frac{h-H}{N}\right) \overline{\mu} & \text{for } h = H+1, ..., H+N \\ \overline{\mu} & \text{for } h = H+N+1, ..., M \end{cases} . \tag{7}$$

Thus, our task reduces to obtaining the forecasts $\hat{\mu}_{t+h|t}$ for each $h = 1, ..., H$, which we do by estimating the following econometric model (Jordà, 2005):

$$\mu_{t+h|t} = \sum_{k=0}^{K} \left(\beta_{h,k}\mu_{t-k} + \gamma'_{h,k} z_{t-k}\right) + \upsilon_{t+h|t}, \tag{8}$$

where K determines the forecast-variable lag, and z_t includes a constant and can include additional forecast variables either directly or as non-linear transformations.[12] When the standpoint is quarter T and the horizon is h, we use only the relevant realizations of the forecast variables available until then, $\{z_t\}_{t\in\{T_0,...,T-h\}}$, to obtain $\hat{\beta}_{h,k}$ and $\hat{\gamma}'_{h,k}$, which then lead to the following real-time forecast: $\hat{\mu}_{T+h|T} = \sum_{k=0}^{K} \left(\hat{\beta}_{h,k}\mu_{T-k} + \hat{\gamma}'_{h,k} z_{T-k}\right)$.[13] The forecast error is $\hat{\upsilon}_{T+h|T} = \mu_{T+h} - \hat{\mu}_{T+h|T}$.

3.2 Second moments: forecasting the variance of forecasting errors

The variance of the errors in the forecast of lifetime loss rates is:

$$\begin{aligned} \text{var}\left(\Lambda_t - \Lambda^e_t\right) &= E\left(l'_{t..t+M-1}\hat{\upsilon}_{t+1..t+M}\hat{\upsilon}'_{t+1..t+M} l_{t..t+M-1}\right) \\ &= l'_{t..t+M-1} E\left(\hat{\upsilon}_{t+1..t+M}\hat{\upsilon}'_{t+1..t+M}\right) l_{t..t+M-1} \\ &\equiv l'_{t..t+M-1}\Omega_t l_{t..t+M-1}, \end{aligned} \tag{9}$$

where $l_{t..t+M-1} = (l_t, ..., l_{t+M-1})'$, $\hat{\upsilon}_{t+1..t+M} = \left(\hat{\upsilon}_{t+1|t}, ..., \hat{\upsilon}_{t+M|t}\right)'$, and Ω_t is the covariance matrix of quarterly forecast errors over the entire life of the loan. We estimate first the diagonal of Ω_t, thus obtaining the diagonal matrix $diag\left(\hat{\Omega}_t\right)$, and then the

[12]Since we use direct rather than recursive forecasts based on e.g. a linear VAR system for $h > 1$, it is also straightforward to replace (8) with a non-linear forecasting function. However, an additive error term is necessary for applying the approach in Section 3.2 to second-moment forecasts. We consider one such specification in *Appendix C*, *Table C.4*.

[13]We use T to denote the timing of a forecast and t as a generic index of time.

off-diagonal terms.

The diagonal terms of Ω_t are equal to the potentially time-varying variances $E(\hat{\upsilon}^2_{t+h|t})$, for $h = 1, ..., H$, and – given the assumed deterministic reversion of loss rates in (7) – to 0 for $h = H + 1, ..., M$. For $h = 1, ..., H$, we base our forecasts on w_t, which may but need not be the same as the variable(s) underpinning first-moment forecasts, z_t. Concretely, we follow Harvey (1976) in assuming multiplicative heteroskedasticity of the following form:

$$E(\hat{\upsilon}^2_{t+h|t}) = \exp\left(\sum_{k=0}^{K}\left(\delta'_{h,k}\mu_{t-k} + \theta'_{h,k}w_{t-k}\right)\right), \tag{10}$$

where w_t contains a constant. Ultimately, from the standpoint of quarter T and for a given horizon $h = 1, ..., H$, we estimate jointly (8) and (10) using $\{z_t\}_{t\in\{T_0,...,T-h\}}$ and $\{w_t\}_{t\in\{T_0,...,T-h\}}$. This delivers, inter alia, the (h, h) element of the diagonal matrix of variances: $diag\left(\hat{\Omega}_T\right)_{h,h} = \exp\left(\sum_{k=0}^{K}\left(\hat{\delta}'_{h,k}\mu_{T-k} + \hat{\theta}'_{h,k}w_{T-k}\right)\right)$.

To estimate the off-diagonal terms of Ω_t, we need an identifying restriction. Specifically, we assume that the correlation between forecast errors for horizons h_1 and h_2 is time invariant – $\text{corr}\left(\hat{\upsilon}_{t+h_1|t}, \hat{\upsilon}_{t+h_2|t}\right) = \rho_j$, where $j = h_1 - h_2 > 0$ – and then confirm that this assumption is borne out in the data.[14] Then, denoting $\text{corr}\left(\hat{\upsilon}_{t+1..t+M}\right)$ by Θ, it follows that $\Theta = \left(diag\left(\Omega_t\right)\right)^{-\frac{1}{2}}\Omega_t\left(diag\left(\Omega_t\right)\right)^{-\frac{1}{2}}$, which delivers the second-moment forecast that we are after:

$$\hat{\Omega}_t = \left(diag\left(\hat{\Omega}_t\right)\right)^{\frac{1}{2}}\hat{\Theta}\left(diag\left(\hat{\Omega}_t\right)\right)^{\frac{1}{2}} \tag{11}$$

Thus, the only remaining estimate we need is of the correlation matrix, $\hat{\Theta}$. From the standpoint of quarter T, we obtain $\hat{\Theta}$ on the basis of the entire relevant set of in-sample forecast errors, $\left\{\hat{\upsilon}_{t+h|t}\right\}_{t\in\{T_0,...,T-H\}}$.

3.3 Forecast evaluation

Focusing on one horizon, h, at a time, we evaluate our forecasts on the basis of four metrics. The first two metrics capture the general match between forecasts and the corresponding realizations: root mean squared error (RMSE) and correlation. The other two capture the capacity of forecasts to signal turning points: the number of

[14] See *Appendix C*, *Figure C.4*.

excessive/insufficient turning points in the forecasted relative to the realized series, and the average distance between the turning points in these series.[15]

Since the latter two metrics are nonstandard, we explain them in the rest of this sub-section, focusing on first moments for concreteness. For the excessive/insufficient turning points, we first identify peaks and troughs in the realized series. By the Harding and Pagan (2002) algorithm, there is a peak at t if

$$\mu_{t-P}, ..., \mu_{t-1} < \mu_t > \mu_{t+1}, ..., \mu_{t+P},$$

where P determines the identification window. A trough is identified with the inequalities reversed. And we apply the same criteria to the forecasted series, $\left\{\hat{\mu}_{t+h|t}\right\}_t$, for each horizon $h \in \{1, ...H\}$. The absolute value of the difference between the number of identified peaks and troughs in the realized series and that number in the forecasted series delivers the metric "excessive/insufficient turning points" for each h. It would be zero for the ideal forecast.

To assess the alignment between the forecasted and realized peaks and troughs, we construct the average standardized distance between peaks and troughs. To this end, we align the time series of realized loss rates $\{\mu_t\}_t$ with the corresponding forecast, $\left\{\hat{\mu}_{t|t-h}\right\}_t$, for $h \in \{1, ..., H\}$. Given this alignment, we start from a peak (trough) in the series of forecasts and determine the minimum between (i) the time distance to the nearest peak (trough) in the realized series, and (ii) the forecast horizon, h. We then divide this minimum by h and average across all peaks/troughs in the series of forecasts. The resulting standardized metric ranges from the ideal value of 0 to 1. At the latter value, forecasts do not provide any information about turning points.

In order to assess the performance of forecasts as signals of turning points, it is necessary to consider the latter two metrics jointly. This is because the "distance" metric is depressed (respectively, inflated) artificially if there are excessive (respectively, insufficient) peaks/troughs.

[15]More standard measures of signaling capacity, such as the Area Under the Receiver Operating Characteristic, cannot be meaningfully applied in our context, as there are too few peaks and troughs in the loss rates over our prediction sample.

4 Application

We apply our empirical framework to quarterly US data from 1985q1 to 2019q2.[16] We start from a training sample consisting of the first 15 years of observations, 1985q1-1999q4, and then expand it by one quarterly observation at a time to generate our (pseudo) real time forecasts. In line with the literature (Covas and Nelson, 2018, and Loudis and Ranish, 2019), we set the maximum forecast horizon to three years ($H = 12$), the transition period to two years ($N = 8$) and the maturity of the portfolio to 7.5 years ($M = 30$).[17] With these choices, our first forecasts for the quarterly losses are $\hat{\mu}_{00q1|99q4}$,..., $\hat{\mu}_{02q4|99q4}$.

The loss-rate data are quarterly charge-off rates at different levels of portfolio aggregation. For our benchmark analysis, we consider net charge-off rates (i.e. gross charge-offs net of recoveries) on loans to the *total* private non-financial sector. Then, to shed further light on our results, we also consider charge-off rates on sub-portfolios of specific loan types. The exact definitions and sources for all data series are in Table B.1, Appendix B.

We consider eight (candidate) forecast variables. These include variables that have been extensively used to forecast credit loss rates: the unemployment rate (UR_t), an output gap measure (OG_t), a corporate credit spread (CS_t), a term spread (TS_t), an (ex-post) real short-term interest rate (RIR_t), and a volatility index (VXO_t).[18] VXO is almost identical to the better-known VIX (correlation 0.99) but is available over a longer period. We also include three financial cycle indicators that have exhibited good early warning properties for banking crises: the debt service ratio (DSR_t), the credit-to-GDP gap ($C2Y_t$) and the house price gap (HPG_t).

In applying our empirical framework to these data, we proceed as follows. When forecasting a variable, we *always* consider its two latest (auto-regressive) observations. Then, outside robustness exercises, we add the two latest observations of only *one* forecast variable at a time for first-moment forecasts and a potentially different forecast

[16]The methodology could also be applied to forecasting non-performing loan (NPL) rates. Gambera (2000) obtains first-moment NPL forecasts from both direct and iterative methods. They find that linear specifications, employing a small set of variables, deliver reliable out-of-sample forecasts over short horizons, up to two quarters. That said, much of the related literature seeks to identify NPL determinants in-sample (see Ghosh (2015) and references therein).

[17]Drehmann and Juselius (2012) find that the remaining maturity on the stock of US private non-financial sector loans is about 10 years (40 quarters) on average. With discounting, the effective remaining maturity is shorter. By setting $M = 30$ quarters, we proxy for the effect of discounting.

[18]We also considered growth rates of real credit, real GDP, real house prices and real equity prices. The performance of these variables is poorer than that of those we use.

variable for second-moment forecasts. The lag parameter in equations (8) and (10) is thus $K = 1$.

We settle on the ultimate forecasts through the following steps. We start with purely first-moment forecasts to inform ourselves on how to proceed for the joint forecasts of first and second moments. Namely, our preliminary exercise is to derive first-moment forecasts on the basis of equation (8) or a non-linear variant and a battery of forecast variables. Then, we evaluate the forecasts of the alternative models, as outlined in Section 3.3. Since the second moments are squared errors in the first-moment forecasts, it is important to avoid using poor first-moment forecasts, which would distort the forecast errors. Thus, keeping only the best first-moment model from the preliminary exercise, we move to a joint forecast of first and second moments, i.e. equations (8) and (10). At this stage, we explore alternative second-moment forecast variables and identify the best joint first- and second-moment forecasts, again on the basis of the evaluation metrics in Section 3.3. Lastly, we aggregate the chosen first-moment forecasts to derive a time series of EL over the lifetime of the bank's portfolio – per equation (6) – and use the chosen forecasts of the lifetime portfolio variance – per (11) – to construct the lifetime UL.

4.1 First moments: forecasting quarterly loss rates

We explore sequentially linear and non-linear models of the level of quarterly loss rates.

4.1.1 Linear specification

In our linear specifications, the regressors z_t in equation (8) are untransformed forecast variables. For a visual assessment of the results, we focus on a subset of the forecast variables z_t and plot the time series of the forecasts they deliver alongside the corresponding realized loss rates in *Figure 3*. We then examine more formally the performance of each z_t in *Table 1*. To avoid cluttering these and similar figures and tables below, we report results only for $h \in \{1, 4, 8, 12\}$.[19] And, in order to filter out high-frequency outliers, we identify peaks and troughs on the basis of two-year windows, i.e. setting P from Section 3.3 to 8.

As a benchmark, we start with forecasts obtained from a purely auto-regressive specification. Given the persistent boom-bust phases in the realized loss rates and

[19]For example, for $h = 4$, the series is $\hat{\mu}_{h=4} = (\hat{\mu}_{00q4|99q4}, \hat{\mu}_{01q1|00q1}, ..., \hat{\mu}_{19q2|18q2})'$.

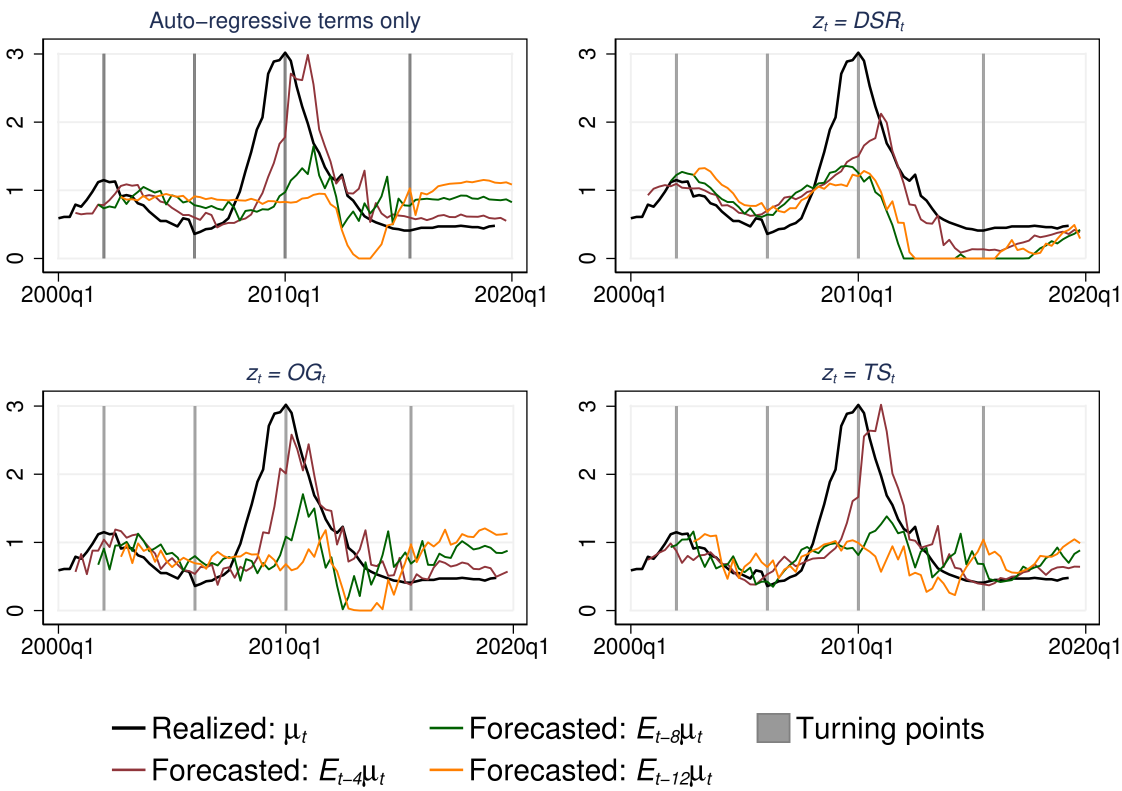

Figure 3: **First moments.** Realized and forecasted quarterly loss rates based on OLS estimates of equation (8) with lag parameter $K = 1$. The forecast variable in z_t is: DSR_t = debt service ratio, OG_t = output gap, or TS_t = term spread. An observation in quarter t is equal to: a real-time forecast formed at $T = \{t-4, t-8, t-12\}$ or the corresponding realization at t.

z_t :	–	DSR_t	$C2Y_t$	HPG_t	UR_t	OG_t	CS_t	TS_t	RIR_t	VXO_t
h					Root mean square error					
1	0.13	0.14	0.13	0.13	0.13	**0.12**	0.13	0.13	0.14	0.13
4	0.46	0.47	0.45	0.48	0.44	**0.40**	0.42	0.47	0.54	0.45
8	0.70	0.69	0.70	**0.63**	0.67	0.73	0.71	0.64	0.76	0.72
12	0.81	**0.71**	0.75	0.85	0.79	0.85	0.82	**0.71**	0.88	0.89
					Correlation (realized vs. forecasted)					
1	0.98	0.98	0.98	0.98	0.98	**0.99**	0.98	0.98	0.98	0.98
4	0.75	0.79	0.77	0.74	0.78	**0.82**	0.80	0.75	0.68	0.77
8	0.27	**0.60**	0.44	0.53	0.41	0.22	0.27	0.58	0.35	0.17
12	0.00	**0.58**	0.43	0.29	0.11	-0.06	0.06	0.40	0.13	-0.22
					Excess/insufficient turnings					
1	**0**	**0**	**0**	**0**	**0**	1	1	1	1	1
4	2	**0**	1	3	1	**0**	2	1	2	2
8	2	**0**	1	1	**0**	1	2	2	1	2
12	2	**0**	1	1	**0**	3	2	1	1	2
					Standardized average distance					
1	0.75	0.75	0.75	1.00	0.75	**0.50**	1.00	0.75	1.00	0.75
4	1.00	0.92	1.00	**0.42**	1.00	**0.42**	0.50	0.58	0.58	0.67
8	0.88	0.42	0.46	0.58	0.79	0.79	0.88	**0.38**	0.63	0.88
12	0.69	**0.17**	0.83	0.69	0.53	0.44	0.61	0.50	0.28	0.67
# best	1	**8**	1	3	3	7	0	2	0	0

Table 1: **First moments.** Evaluation of forecasted quarterly loss rates based on OLS estimates of equation (8) with lag parameter $K = 1$ and horizon h as indicated. The forecast variable in z_t is: DSR_t = debt service ratio, $C2Y_t$ = credit-to-GDP gap, HPG = house price gap, UR_t = unemployment rate, OG_t = output gap, CS_t = corporate bond spread, TS_t = term spread, RIR_t = real interest rate, or VXO_t = volatility index. The row "# best" counts the number of times an indicator is the best performer (boldface values).

the abrupt switches between these phases, these forecasts tend to miss turning points. This is visible quite clearly in the time series plots (*Figure 3*, top left-hand panel) and is then confirmed by the four metrics in *Table 1* (first column). While the "auto-regressive" forecasts are highly correlated with realized loss rates at the short horizons, the correlation deteriorates substantially for $h > 4$. At such horizons, the purely auto-regressive forecasts miss two turning points completely and flag the other ones with material lags (the distance metric is close to 1).

The debt service ratio (DSR) stands out with its forecasting performance of the level of loss rates.[20] Visually, the DSR-based forecasts help improve on the auto-regressive

[20] Relatedly, Banerjee and Kharroubi (2020) find that debt service burden at the sectoral level is an

specifications by matching better turning points in realized loss rates at all horizons (*Figure 3*, top right-hand panel). In comparison, any gains of the OG- or TS-based forecasts are difficult to detect (bottom, left-hand panel). That said, the plots reveal that no forecast accounts for the full swing of loss rates during the GFC and all underpredict these rates' post-GFC levels.

In terms of the evaluation metrics (*Table 1*), the overall top performer is DSR (second column), featuring the best score in 8 out of the 16 cases. DSR-based forecasts have low RMSEs, are highly correlated with realized loss rates, and identify the exact number of turning points. Importantly, the (standardized) distance between forecasted and realized turning points declines as the horizon increases, indicating that the DSR detects turning points well in advance. For its part, OG outperforms at the short horizons (sixth column). Since most variables fare well at such horizons, however, the gains of the OG-based forecasts are modest.

The performance of the other variables is quite patchy. That of the credit-to-GDP gap is consistent – even if mediocre – across most horizons but breaks down at the longest ones, e.g. $h = 12$. The term spread (TS) performs similarly as the DSR in terms of RMSE and the correlation metric but its quarter-to-quarter volatility results in it flagging too many peaks and troughs (*Table 1*, eighth column). We also see why Krüger et al (2018) are quite sceptical about real-time loss-rate forecasts: closely related to the VIX that they focus on, VXO is among the worst performers in our exercise (last column).

Since the best performers in *Table 1* change with the metric and the horizon, different forecast variables may carry complementary information about loss rates. We thus examine forecasts based on multi-variate versions of z_t. We report the results in *Figure C.1* and *Table C.1* in *Appendix C*.

The gains of combining several variables are modest at best, not least because such combinations do not lead to better forecasts of the magnitude of losses during the GFC and continue to underpredict post-GFC losses. Concretely, we first examine a specification that combines the best performers at specific horizons in *Table 1*, ie the debt service ratio, the output gap and the term spread (*Table C.1*, second column). This does not yield any substantive gains over forecasts based only on the debt service ratio (first column). We then consider a specification that combines the unemployment rate, the output gap and the house price gap (third column), akin to the exercise in Covas and

(in-sample) predictor of firm exits.

Nelson (2018). Given this paper's pessimism as regards real-time loss-rate forecasts, it is hardly surprising that these forecast variables deliver subpar performance, especially in the case of signaling turning points. Finally, in order to consider all forecast variables discussed in Krüger et al (2018), we add the credit spread, the term structure and the volatility index to the latter specification (fourth column). But this actually decreases forecasting accuracy and adds high frequency noise to the forecasts, highlighting the implications of in-sample overfitting.

4.1.2 Non-linear specifications

Next we allow for a non-linear relationship between forecast variables and loss rates, possibly due to amplification mechanisms that play out only in specific financial environments, corresponding to specific levels of the forecast variables. Concretely, we revert to considering a single forecast variable (in addition to auto-regressive terms), and estimate three specifications. In the first two, we let the relationship between loss rates and the forecast variable be, respectively, quadratic and asymmetric around the forecast variable's real-time mean. The third specification – which we estimate with non-linear least squares – features an exponential forecasting function and an additive error term. *Tables C.2-C.4* in *Appendix C* define each specification and report the corresponding results. The bottom line is that these richer specifications do not consistently deliver better forecasts of loss rates, even though they introduce greater complexity.

In sum, simple models, based on auto-regressive terms and a single forecasting variable forecast well the turning points of realized loss rates at long horizons. However, neither these nor more elaborate specifications – which generate limited and inconsistent improvements – succeed in matching the *amplitude* of loss-rate swings. These swings may thus be related to the inherently unexpected part of the loss process – e.g. its second moments – which we investigate next.

4.2 First and second moments: forecasting *jointly* quarterly loss rates and forecast errors' variances

The stylized theoretical model underscores benefits of forecasting second moments. For one, it implies that second and first moments rise in tandem: that is, predictability declines exactly when the expected level rises. Thus, the poor match between forecasted and realized first moments during the GFC should not be surprising and should mo-

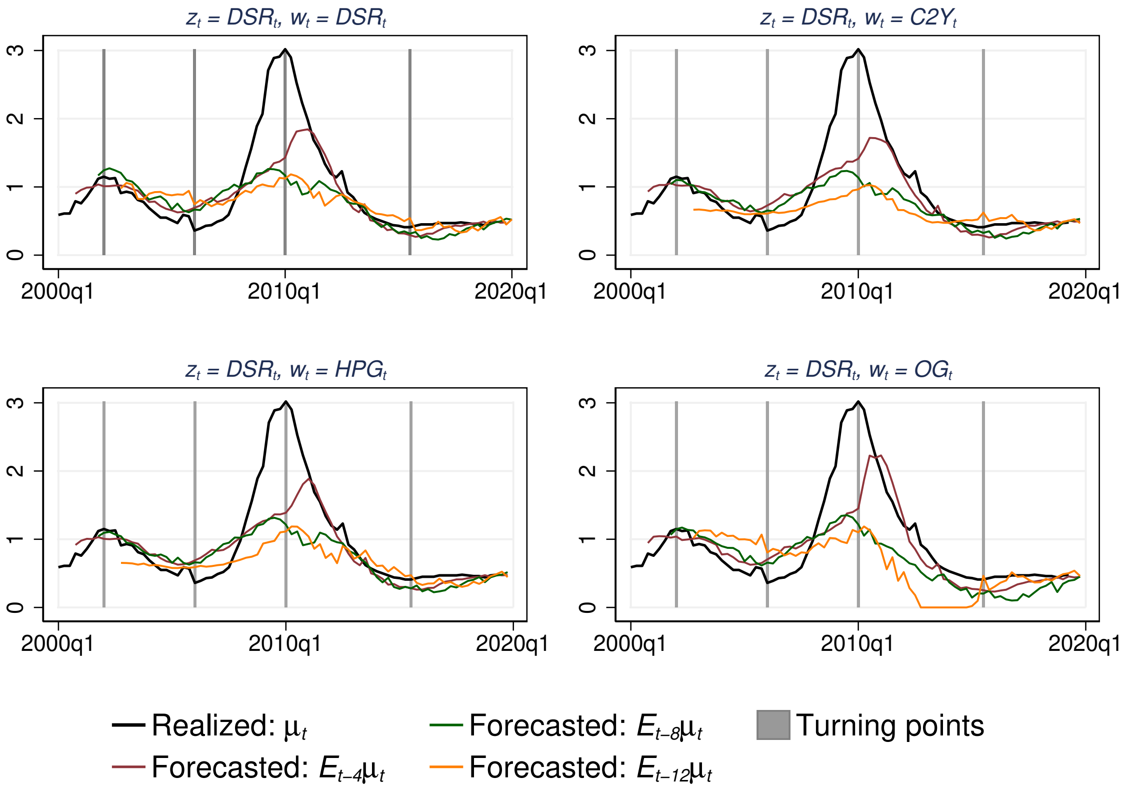

Figure 4: **First moments (estimated jointly with second moments).** Realized and forecasted quarterly loss rates based on joint maximum likelihood estimation of equations (8) and (10) with lag parameter $K = 1$. The forecast variables (in addition to auto-regressive terms) are: DSR_t = debt service ratio (in z_t for first moments) and DSR_t, $C2Y_t$ = credit-to-GDP gap, HPG_t = house price gap, or OG_t = output gap (in w_t for second moments). An observation in quarter t is equal to: a real-time forecast formed at $T = \{t-4, t-8, t-12\}$ or the corresponding realization at t.

tivate second-moment forecasts. In addition, the model suggests that studying second moments should generate useful information for first-moment forecasts. The empirical exercise to which we turn next reveals that this is indeed borne out in the data.

For our joint first- and second-moment forecasts, we apply the maximum-likelihood estimator of Harvey (1976) simultaneously to (8) and (10). Given the results in the previous section, we set $z_t = DSR_t$ in (8) and then consider sequentially each of the eight forecast variables for w_t in (10). We report the time series of forecasted quarterly loss rates and error variances, together with the corresponding actual observations, in *Figures 4* and *5*. The assessment metrics are in *Table 2*.

While the performance of second-moment forecasts varies across horizons (*Figure*

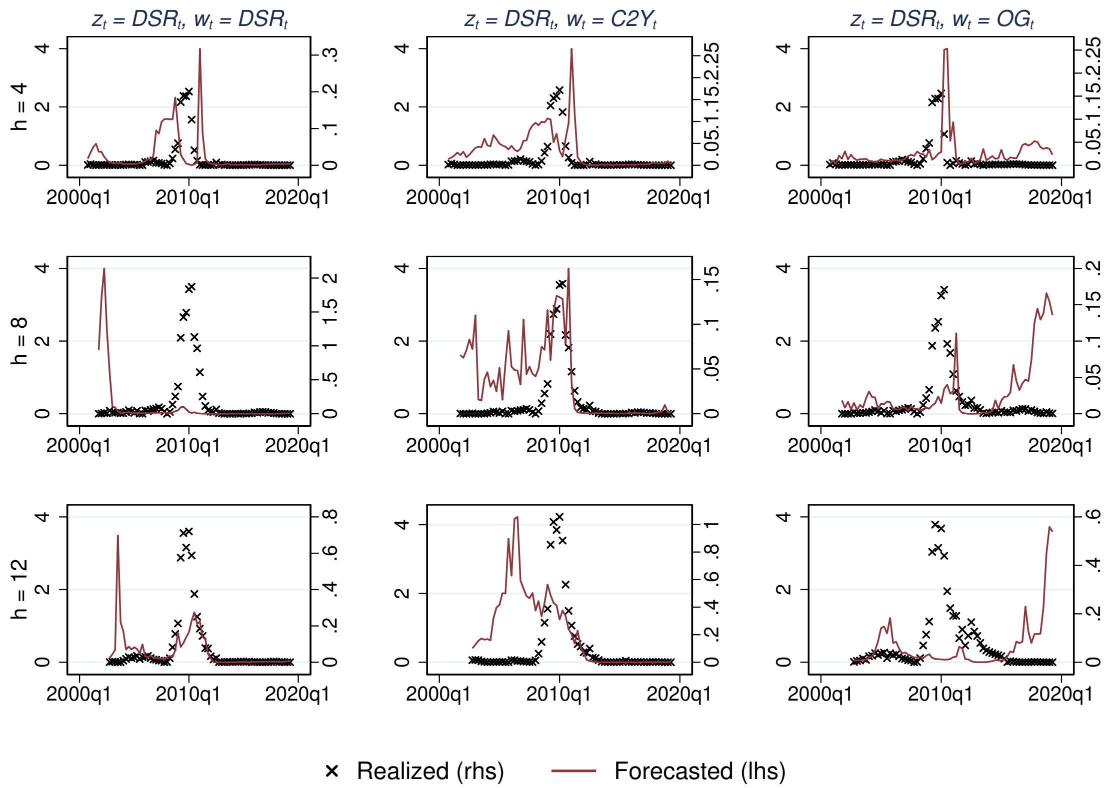

Figure 5: **Second moments (estimated jointly with first moments).** Realized and forecasted squared errors in first-moment forecasts based on joint maximum likelihood estimation of equations (8) and (10) with lag parameter $K = 1$ and horizon h as indicated. The forecast variables (in addition to auto-regressive terms) are: DSR_t = debt service ratio (in z_t for first moments) and DSR_t, $C2Y_t$ = credit-to-GDP gap, or OG_t = output gap (in w_t for second moments). An observation in quarter t is equal to: a real-time forecast formed at $T = \{t-4, t-8, t-12\}$ or the corresponding realization at t.

5), the joint estimation improves visibly the first-moment forecasts (compare *Figures 3* and *4*). The latter is in particular the case for first-moment forecasts based on the debt service ratio, the credit to GDP gap, or the property price gap. The improvement comes mostly from post-GFC quarterly loss rates being no longer under-predicted. In parallel, most – but not all – of the plots in *Figure 5* reveal a reasonable match between the high level of squared errors during the GFC and their real-time forecasts.

The evaluation metrics in *Table 2* confirm that the joint first- and second-moment estimation improves the performance of first-moment forecasts. This is seen by comparing the upper part of *Table 2* with the second column of *Table 1* (where we forecast only first moments on the basis of $z_t = DSR_t$). The improvement is strongest when the accompanying second-moment models employ one of the financial cycle measures: debt service ratio, credit-to-GDP gap or property price gap. All of these specifications lead to: a decline in RMSEs; an increase in the correlation between realized and forecasted losses, especially at the longer horizons; and a preservation or improvement of the already good turning-point identification. Using other variables for second-moment forecasts does deliver some improvements – notably, in the case of the output gap or the term structure – but only at the shorter horizons.

To assess the performance of second-moment forecasts, we need to refer to (model-dependent) squared forecast errors, as the actual variances are not observable. In this case, we use only two metrics: the RMSE and the correlation between realized and forecasted series (*Table 2*, lower part). The other two metrics are uninformative (and thus unreported), as squared forecast errors are too erratic to allow for a meaningful identification of turning points.

The best second-moment forecasts condition on the debt service ratio (DSR) or the credit-to-GDP gap (C2Y). The DSR-based forecasts are particularly impressive at the longer horizons, where they deliver the lowest RMSE and the highest correlation metric. While these forecasts are also decent at the shortest horizons, they break down at intermediate horizons, at which they deliver negative correlations between forecasted and realized series. By contrast, C2Y underpins consistently good results across all horizons, with a correlation metric above 60% for 8-quarter forecasts.

We conclude this subsection with an assessment of the statistical significance and stability of the coefficients in two forecast models, both employing DSR for the first moments, and DSR and C2Y, respectively, for the second moments (*Appendix C, Figures C.2* and *C.3*). In the first-moment part of either model, the (sum of) auto-regressive parameters and (the sum of) DSR parameters are fairly stable. In line with earlier results indicating poor predictive power of the auto-regressive coefficients, these coefficients decline substantially with the horizon and become mostly (jointly) insignificant at $h = 12$. By contrast, the DSR coefficients are significant and positive at all horizons and throughout the forecast period, indicating that a high debt service ratio systematically predicts higher losses in the future. Turning to the second-moment parts of the models, the auto-regressive coefficients are less stable than in the first-moment parts, whereas the coefficients on the DSR and, especially, C2Y remain mostly stable, positive and significant.

z_t :	DSR_t	DSR_t	DSR_t	DSR_t	DSR_t	DSR_t	DSR_t	DSR_t	DSR_t	DSR_t
w_t :	–	DSR_t	$C2Y_t$	HPG_t	UR_t	OG_t	CS_t	TS_t	RIR_t	VXO_t
h	Root mean square error									
1	**0.13**	0.14	**0.13**	0.14	**0.13**	**0.13**	**0.13**	**0.13**	**0.13**	**0.13**
4	0.45	**0.44**	0.45	0.45	**0.44**	**0.44**	0.46	**0.44**	0.46	0.47
8	0.61	0.58	0.59	**0.56**	0.59	0.59	0.60	0.61	0.62	0.61
12	0.71	**0.63**	0.68	**0.63**	0.70	0.73	0.67	0.66	0.72	0.70
	Correlation (realized vs. forecasted)									
1	**0.99**	0.98	**0.99**	0.98	0.98	**0.99**	**0.99**	0.98	**0.99**	0.98
4	0.78	**0.83**	**0.83**	0.82	0.80	0.80	0.80	0.80	0.79	0.76
8	0.67	0.69	0.74	**0.75**	0.70	0.70	0.68	0.67	0.66	0.67
12	0.46	0.70	**0.86**	**0.86**	0.52	0.46	0.58	0.54	0.47	0.50
	Excess/insufficient turnings									
1	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**
4	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**
8	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	2	**0**
12	**0**	**0**	2	**0**	**0**	**0**	2	**0**	**0**	**0**
	Standardized average distance									
1	**0.75**	**0.75**	**0.75**	**0.75**	**0.75**	**0.75**	**0.75**	**0.75**	**0.75**	**0.75**
4	0.67	0.67	**0.50**	0.67	0.67	0.75	**0.50**	0.67	0.75	0.67
8	0.33	0.38	**0.29**	**0.29**	0.33	0.33	0.33	0.33	0.33	0.33
12	0.22	0.36	0.31	0.36	0.25	0.25	0.28	0.31	**0.17**	**0.17**
	Quarterly variances, $E(v_{t+h}^2)$									
	Root mean square error									
1	0.13	**0.07**	0.10	0.11	0.10	**0.07**	0.14	0.17	0.14	0.34
4	0.56	0.58	0.58	0.65	0.55	0.55	**0.54**	0.56	0.56	0.62
8	0.80	0.93	0.86	0.82	0.79	0.81	0.81	0.82	**0.76**	0.82
12	1.01	**0.90**	1.04	1.33	0.96	1.04	1.18	0.97	1.11	1.06
	Correlation (realized vs. forecasted)									
1	0.15	0.27	0.23	0.19	0.17	**0.36**	0.11	0.13	0.18	0.09
4	0.00	-0.01	**0.21**	-0.12	0.07	0.11	0.43	0.10	-0.04	0.04
8	-0.06	-0.07	**0.64**	-0.13	-0.06	-0.07	0.03	-0.17	-0.02	-0.11
12	-0.29	**0.36**	0.23	-0.11	-0.21	-0.26	-0.13	-0.13	-0.22	-0.27
# best	7	11	**12**	10	7	10	8	7	8	8

Table 2: **First and second moments.** Evaluation of forecasts based on joint maximum likelihood estimates of equations (8) and (10) with lag parameter $K = 1$ and horizon h as indicated. The forecast variables (in addition to auto-regressive terms) are: DSR_t = debt service ratio (in z_t for first moments) and DSR_t, $C2Y_t$ = credit-to-GDP gap, HPG = house price gap, UR_t = unemployment rate, OG_t = output gap, CS_t = corporate bond spread, TS_t = term spread, RIR_t = real interest rate, or VXO_t = volatility index (in w_t for second moments). The row "# best" counts the number of times an indicator is the best performer (boldface values).

4.3 Lifetime loss rates

As Sections 4.1 and 4.2 deliver respectively quarterly loss rates $\{\hat{\mu}_{t+h|t}\}_{t,h}$ and the attendant variance-covariance matrix $\{\hat{\Omega}_t\}_t$, we revert to Sections 3.1 and 3.2 to derive loss forecasts over the loan portfolio's lifetime, Λ_t^e, and the lifetime variance of these losses. In this section, we study the properties of two versions of Λ_t^e and lifetime variance estimates. For each version, we estimate first and second moments simultaneously, employ the debt service ratio (DSR) in a linear specification of the first moments – per equation (8) – and rely on a multiplicative heteroskedasticity specification for the second moment, per equation (10). The versions differ in that one of them employs the DSR and the other one the credit-to-GDP gap (C2Y) for the second-moment specification. As indicated by *Tables 2* and *C.1-C.4*, these forecasts perform consistently the best across horizons among a battery of alternatives.

To assess first- and second-moment forecasts of lifetime loss rates, we compare them to the corresponding perfect-foresight magnitudes. Importantly, the latter are different from *actual* loss rates. We construct perfect-foresight first moments on the basis of equations (6) and (7), where we replace forecasts $\hat{\mu}_{t+h|t}$ with actual observations μ_{t+h} for horizons $h = 1, ..., H$ but maintain the subsequent deterministic reversion to the historical long-term average. Likewise, we base perfect-foresight second moments on the actual residual obtained from estimating equation (8) for $h = 1, ..., H$, and set these residuals to zero for $h > H$. In this way, we assess the forecasts over the "reasonable and supportable" horizon.

Our assessment reveals that forecasts of lifetime losses perform in general even better than the forecasts of quarterly losses, as horizon-specific errors tend to wash out. This is seen by comparing *Figures 4 and 5* with *Figure 6*. The forecasted and perfect-foresight lifetime losses exhibit synchronous ups and downs throughout the sample period, even though the former falls short of the latter when the GFC sets in (*Figure 6*, top panels). In an indication that the GFC-induced losses were genuinely higher than what one could have expected in real time, the forecasted lifetime variances of portfolio losses explain well the swings in the corresponding squared forecast errors from 2006 onward (bottom panels). At the beginning of the sample period, the second-moment forecasts appear as overly conservative.

The good forecasting performance at the level of lifetime losses is confirmed more formally in *Table 3*, where we report results for the same evaluation metrics as above. The correlations between forecasted and perfect-foresight series are above 0.8 for the

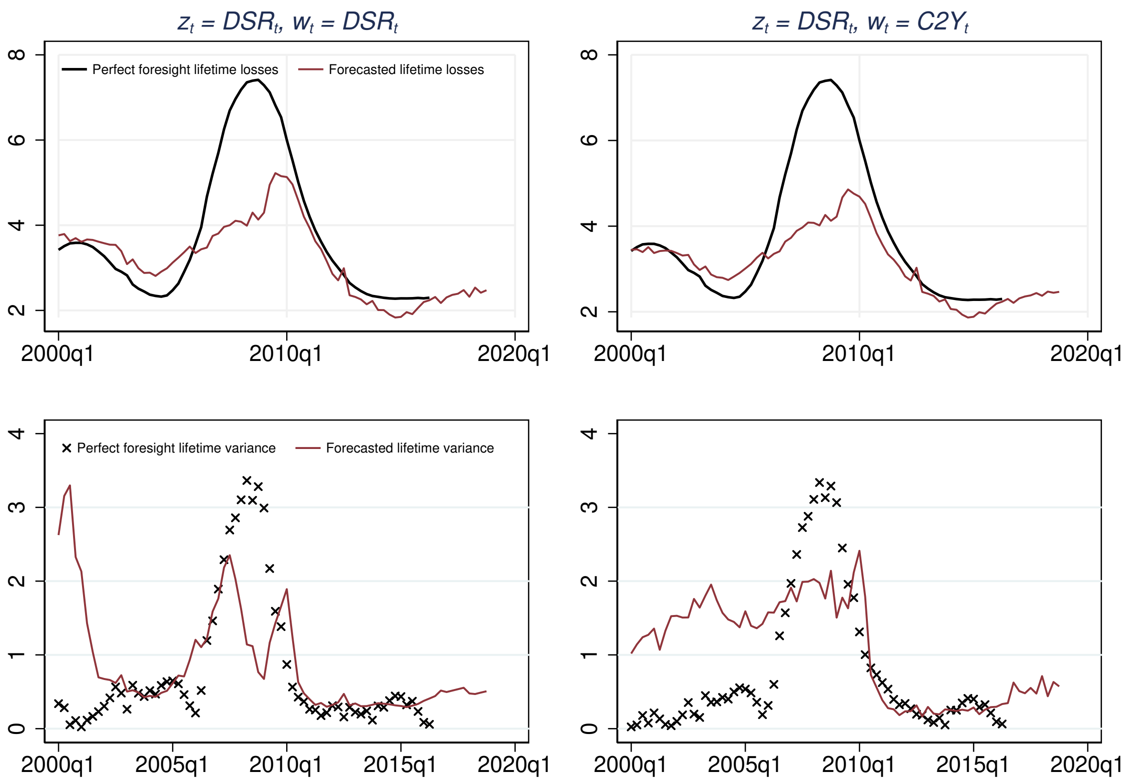

Figure 6: **First and second moments over the portfolio's lifetime.** An observation in quarter T is equal to: a real-time forecast at T or the corresponding materialization of risk from T onward. The series are based on forecasted and realized quarterly magnitudes as well as expression (7) for the deterministic evolution of losses beyond the maximum forecast horizon. The forecast variables (in addition to auto-regressive terms) are: DSR_t = debt service ratio (in z_t for first moments) and DSR_t or $C2Y_t$ = credit-to-GDP gap (in w_t for second moments).

		Lifetime EL				Variance of lifetime forecast errors		# best
z_t	w_t	RMSE	Corr.	Turn.	Dist.	RMSE	Corr.	
DSR_t	DSR_t	**1.19**	0.81	**0**	**0.33**	3.38	0.16	3
DSR_t	$C2Y_t$	1.23	**0.86**	**0**	**0.33**	**2.49**	**0.58**	**5**

Table 3: **Lifetime expected loss (EL) and error variance.** The evaluations are based on forecasted and perfect foresight lifetime magnitudes that incorporate corresponding quarterly magnitudes as well as expression (7) for the deterministic evolution of losses beyond the maximum forecast horizon. The forecast variables (in addition to auto-regressive terms) are: DSR_t = debt service ratio (in z_t for first moments) and DSR_t or $C2Y_t$ = credit-to-GDP gap (in w_t for second moments). The column "# best" counts the number of times an indicator is the best performer (boldface values).

first moments and almost 0.6 for the C2Y-based second-moment results. Each model identifies the right number of turning points and the standardized average distance between forecasted and perfect-foresight turning points is 0.33, or two quarters divided by the average six-quarter forecast horizon. Overall, the model that employs C2Y for second-moment forecasts outperforms the one employing DSR for all but one of the evaluation metrics.[21]

4.3.1 Provisions and capital

What would banks' provisions and capital have looked like if they had been based on our forecasts? To answer this, we first assume that expected loss (EL) provisions are set according to our forecasts of lifetime loss-rate levels, on the basis of either of the two models that we focus on in Section 4.3. In addition – remaining agnostic about the probability distribution of loss rates – we assume that capital is set to match unexpected losses (UL) that equal two, three or four standard deviations of lifetime losses.

Our forecasts would have resulted in provisions that account for the accumulation, as opposed to the materialization, of credit risk to the extent that they track the lifetime losses realized over the forecast horizon. Indeed, as indicated in the top panels of *Figure 6*, risk started accumulating in 2004q4 (black lines) and this is when the forecast-implied provisions would have started picking up as well. By the peak of the GFC in 2009, these provisions would have resulted in banks accumulating loan-loss reserves that amount to about 60% of the subsequently realized losses. And the reserves

[21] While these are indirect forecasts of lifetime loss rates' first and second moments – resulting from aggregating quarterly loss-rate forecasts – they fare generally better than direct forecasts of lifetime losses: compare *Table 3* to *Table D.1* in *Appendix D*.

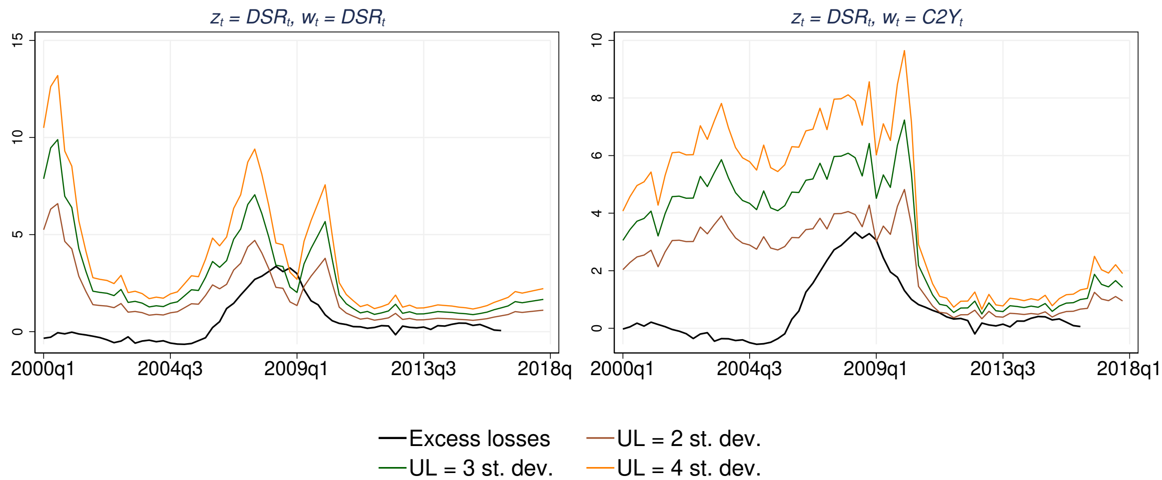

Figure 7: **Capital adequacy.** "Excess" loss is equal to the difference between perfect-foresight and expected losses (*Figure 6*, top panels). Unexpected losses (UL) are equal to two, three or four times the forecasted standard deviation of lifetime forecast errors (*Figure 6*, bottom panels).

would have started declining in 2009q3, three quarters after credit risk starts declining. By contrast, historical reserves – based on incurred loss accounting – start picking up only in 2007q3 and start declining in 2010q2 (*Figure 1*).[22]

The capital implied by our UL forecasts would have been largely sufficient to cover perfect-foresight losses. Concretely, "excess" losses (relative to forecast-implied provisions) will have been effectively covered by capital set equal to four times the DSR-based standard-deviation forecasts or two times the C2Y-based forecasts (*Figure 7*). While only the DSR-based UL forecasts match closely the time profile of excess losses throughout most of the sample time period, both DSR- and C2Y-based forecasts deliver a particularly close match from 2008q4 onward.

Importantly, the forecasts we have derived so far fare better – in the sense of delivering sufficient loss absorbing resources – than *direct* forecasts of lifetime portfolio losses (*Appendix D*). This is the case when we construct UL on the basis of second-moment forecasts (*Figure D.1*) or on the basis of quantile regressions (*Figure D.2*). The main reason is that these alternative forecasts fail to account for the magnitude of loss rates' increase during the GFC.

[22] Admittedly, historical LLR are not fully comparable with the other series in the figure. While LLR reflect the outstanding loans in banks' portfolios, perfect-foresight losses and model-based provisions reflect the assumptions in expression (7).

	Dependent variable									
	EL	IL	EL	IL	EL	IL	EL	IL	EL	IL
CBO	0.020* (0.011)	−0.013* 0.007								
LW			0.034* 0.018	−0.023** 0.011						
IMF					0.032*** 0.010	−0.002 0.009				
OECD							0.029*** 0.009	0.003 0.006		
HP									0.017 0.016	−0.005 0.009
R-squared	0.97	0.99	0.97	0.99	0.97	0.99	0.97	0.99	0.96	0.99

Table 4: **Cyclicality of provisions.** The dependent variable is either lifetime expected losses (EL), per Section 4.3.2, when the forecast variables (in addition to auto-regressive terms) are debt service ratio (for first moments) and credit-to-GDP gap (for second moments), or historical loan-loss reserves under incurred loss (IL) accounting. The regressors – contemporaneous to the dependent variable – are alternative output-gap measures: CBO = Congressional Budget Office, LW = Laubach-Williams, IMF = International Monetary Fund, OECD = Organisation for Economic Co-operation and Development, and HP = Hodrick–Prescott. Each reported coefficient comes from a regression equation including also a constant and two auto-regressive terms. The sample is from 1999q4 to 2018q4 for EL (effective sample of 75 quarters) and 1999q3 to 2019q2 for IL (effective sample of 78 quarters). Statistical significance at the 10% and 5% levels is denoted by * and **, respectively.

4.3.2 Cyclicality of forecasts

Would the provisions generated by our forecasts have helped dampen the procyclicality of the financial system? Concretely, would these provisions have been higher during macro-economic expansions and lower in recessions. We address this question by sequentially regressing the EL provisions implied by our first-moment forecasts on alternative *contemporaneous* output-gap measures (as well as two auto-regressive terms). Then, we repeat the exercise, replacing EL provisions with US banks' actual loan loss reserves (LLRs), underpinned by incurred loss (IL) provisions. We report the results in *Table 4*.

The overall conclusion is that replacing IL with EL provisions would have helped dampen procyclicality. In the EL regressions, the coefficients on alternative measures of the output gap are always positive and are significant in four out the five cases. By contrast, in the IL regressions, the coefficients are negative in four of the cases and significantly so in two of them.[23]

[23] These results are consistent with two stylised facts: (i) high GDP growth going hand-in-hand with low loss rates and short-term loss forecasts (Banerjee et al, 2020, and Fatouh and Giansante, 2020); and (ii) accurate long-term forecasts of credit losses increasing while concurrent loss rates are low.

4.4 Digging deeper: sub-portfolios

To understand better how our findings depend on properties of the underlying data, we rerun the above exercises on specific sub-portfolios that exhibit different time profiles of the loss rates. The sub-portfolios we consider comprise: real estate (RE), commercial and industrial (CI), or consumption (CO) loans. For each sub-portfolio, we plot forecasted and perfect-foresight time series for lifetime losses and forecast errors (*Figures 8* and *9*) and report the corresponding values of the assessment metrics (*Table 5*).[24]

The overarching message is that forecasts perform well to the extent that the initial – or training – sample contains information about turning points in loss rates and the boom-bust amplitude in these rates. In the case of RE and CO loans, there is little such information up until the GFC. Thus, unsurprisingly, each of the alternative models and forecast variables fares dismally when it comes to forecasting losses on such loans during the GFC. By contrast, the boom-bust cycle that CI loans undergo during the training period is comparable to that during the GFC. As a result, we obtain accurate forecasts of the magnitude of loss rates on such loans and the attendant turning points. In some respects, these forecasts are even better than those for the overall portfolio: compare the first row in *Table 3* with the second row in the middle panel of *Table 5*. Much of the improvement is due to a better match of the up-swing in loss rates during the GFC: compare *Figures 6* and *8*. As earlier, the best forecasts are those that rely on financial cycle indicators, notably the DSR.

Taken together, the sub-portfolio results bode well for future forecasts. They suggest that the errors in the forecasts of the *aggregate* loan-loss rates during the GFC reflect the lack of key information about the RE and CO sub-portfolios, as opposed to model misspecification or a wrong choice of forecast variables. As such information accumulates over time, future real-time forecasts should be better than the current ones.

5 Conclusions

This paper investigates whether it is possible to derive accurate real-time forecasts of credit loss rates. We pay particular attention to whether the forecasts correlate with realized loss rates and signal turning points in these rates even at the longest "reasonable and supportable" horizons. Our main finding is that direct forecasts – relying on straightforward specifications and using financial cycle indicators, such as the debt ser-

[24]Results at the level of quarterly losses are available upon request.

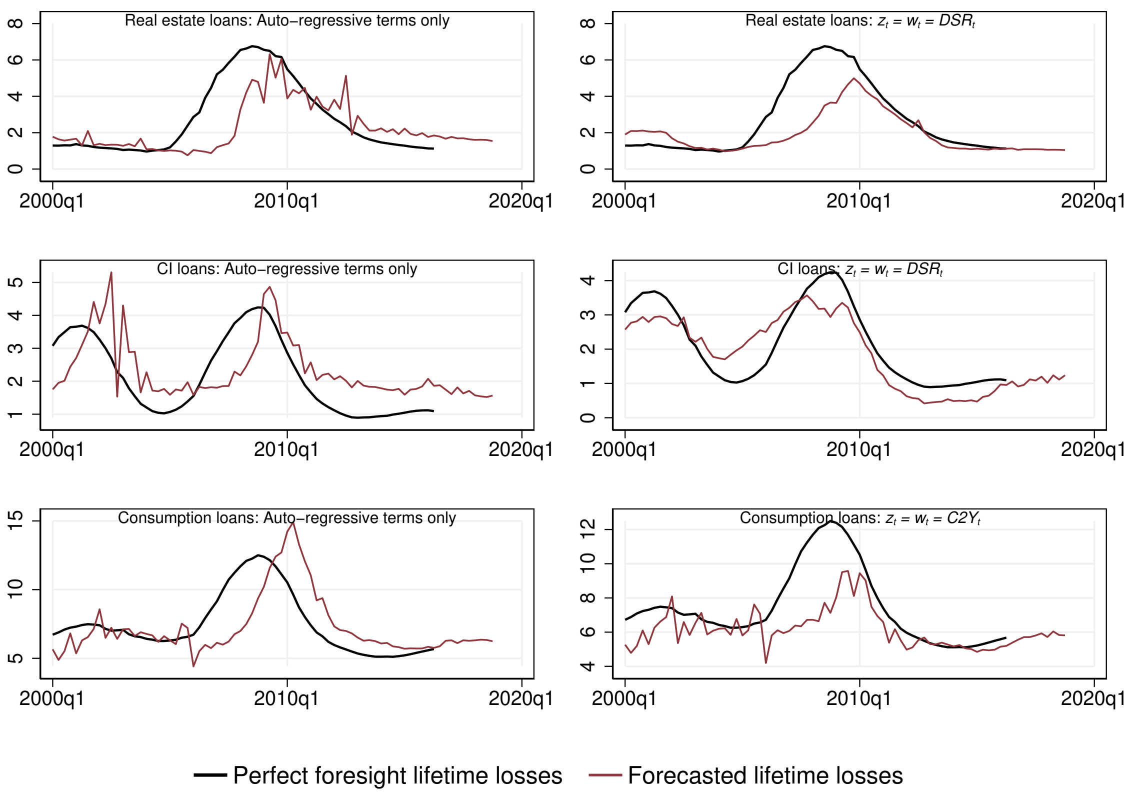

Figure 8: **First moments over the lifetime of sub-portfolios.** An observation in quarter T is equal to: a real-time forecast at T or the corresponding materialization of risk from T onward. The series are based on forecasted and realized quarterly magnitudes as well as expression (7) for the deterministic evolution of losses beyond the maximum forecast horizon. The sub-portfolios are: real estate loans, commercial and industrial (CI) loans, and consumer loans. The right hand panels show results when the forecast variable (in addition to auto-regressive terms) is the same for first- and second-moment forecasts: DSR_t = debt service ratio (real estate and CI loans) or $C2Y_t$ = credit-to-GDP gap.

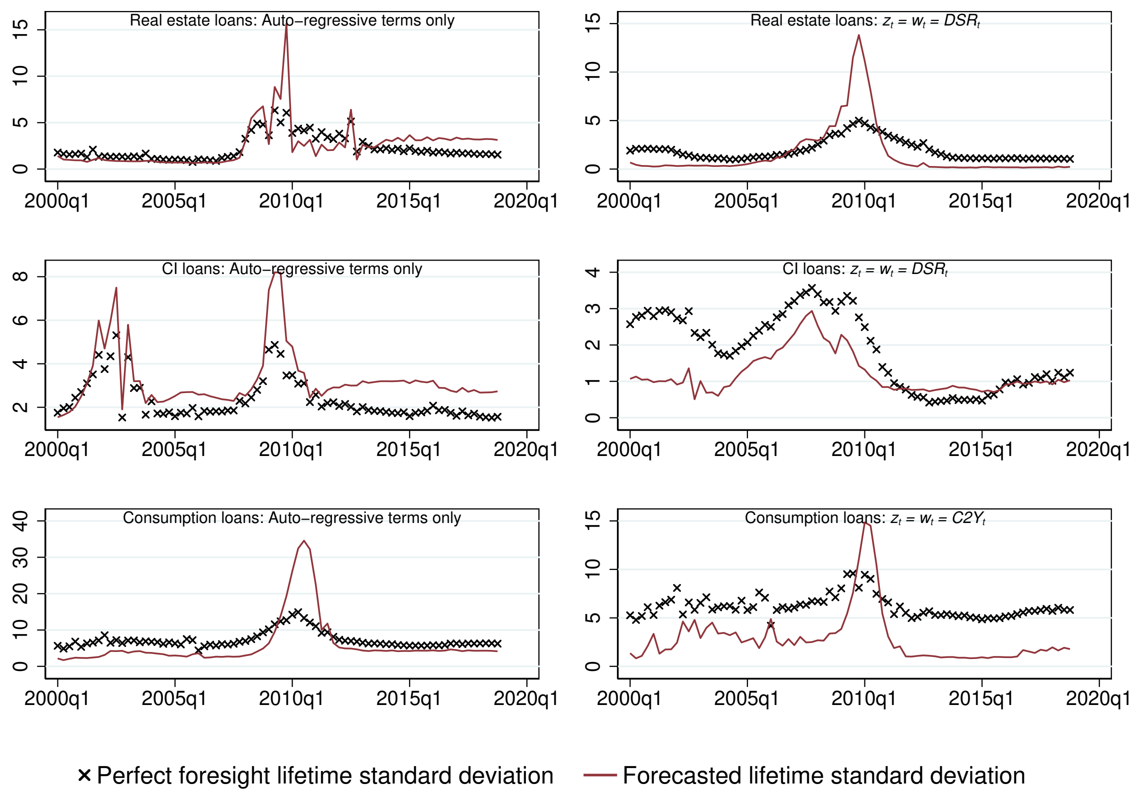

Figure 9: **Second moments over the lifetime of sub-portfolios.** An observation in quarter T is equal to: a real-time forecast at T or the corresponding materialization of risk from T onward. The series are based on forecasted and realized quarterly magnitudes as well as expression (7) for the deterministic evolution of losses beyond the maximum forecast horizon. The sub-portfolios are: real estate loans, commercial and industrial (CI) loans, and consumer loans. The right hand panels show results when the forecast variable (in addition to auto-regressive terms) is the same for first- and second-moment forecasts: DSR_t = debt service ratio (real estate and CI loans) or $C2Y_t$ = credit-to-GDP gap.

	Lifetime EL				Variance of lifetime forecast errors		# best
z_t, w_t	RMSE	Corr.	Turn.	Dist.	RMSE	Corr.	
				Real-estate loans			
–	1.56	0.65	**0.00**	0.74	33.93	-0.06	1
DSR_t	**1.51**	**0.76**	**0.00**	**0.49**	33.65	0.08	**4**
$C2Y_t$	1.59	0.73	1.00	0.91	193.32	0.03	0
HPG_t	1.72	0.54	1.00	0.91	**16.33**	**0.28**	2
UR_t	1.59	0.64	1.00	0.74	81.30	-0.02	0
OG_t	1.61	0.64	**0.00**	0.74	19.24	-0.02	1
CS_t	1.61	0.64	1.00	0.57	108.29	0.13	0
TS_t	1.59	0.66	1.00	0.57	35.12	-0.01	0
RIR_t	1.72	0.70	**0.00**	0.57	29.17	0.08	1
VXO_t	1.66	0.67	**0.00**	0.91	337.06	0.13	1
				C&I loans			
–	1.02	0.56	2.00	0.74	18.30	0.23	0
DSR_t	**0.62**	**0.86**	**0.00**	0.41	**2.40**	0.18	**4**
$C2Y_t$	1.00	0.65	1.00	0.47	13.19	0.00	0
HPG_t	1.90	0.21	1.00	0.73	18.36	**0.56**	1
UR_t	0.92	0.70	1.00	**0.29**	13.08	0.25	1
OG_t	0.91	0.63	1.00	0.65	13.48	-0.15	0
CS_t	1.15	0.54	1.00	0.41	13.27	0.27	0
TS_t	0.76	0.76	1.00	0.53	7.43	-0.05	0
RIR_t	0.96	0.75	2.00	0.57	8.44	-0.26	0
VXO_t	1.11	0.58	2.00	0.69	22.83	0.38	0
				Consumer loans			
–	2.14	0.54	2.00	0.68	255.90	0.57	0
DSR_t	**1.71**	0.67	1.00	0.81	67.16	0.13	1
$C2Y_t$	1.98	**0.72**	**0.00**	**0.56**	43.30	-0.02	**3**
HPG_t	2.18	0.57	2.00	0.79	133.94	0.45	0
UR_t	2.11	0.71	1.00	0.97	**14.44**	0.17	2
OG_t	2.09	0.57	1.00	0.84	151.94	0.59	0
CS_t	2.24	0.41	1.00	0.97	488.60	0.00	0
TS_t	1.76	0.67	1.00	0.68	48.93	0.03	0
RIR_t	2.06	0.57	1.00	0.84	181.53	0.39	0
VXO_t	2.24	0.54	1.00	0.84	299.99	**0.61**	1

Table 5: **Lifetime expected loss (EL) and error variance for sub-portfolios.** The evaluations are based on forecasted and perfect-foresight lifetime magnitudes that incorporate corresponding quarterly magnitudes as well as expression (7) for the deterministic evolution of losses beyond the maximum forecast horizon. The forecast variable (in addition to auto-regressive terms) is the same for first- and second-moment forecasts: DSR_t = debt service ratio, $C2Y_t$ = credit-to-GDP gap, HPG = house price gap, UR_t = unemployment rate, OG_t = output gap, CS_t = corporate bond spread, TS_t = term spread, RIR_t = real interest rate, or VXO_t = volatility index. The row "# best" counts the number of times an indicator is the best performer within each portfolio (boldface values).

vice ratio and credit-to-GDP gap – fare remarkably well, especially if first and second moments of the loss distribution are estimated jointly. This finding is important for policy, as it indicates that banks are in a position to set adequate provisions and capital, in line with the objectives of the expected-loss-provisioning and Basel III standards.

Future research would assess the generality of the above findings. To what extent are they relevant for portfolios weighted towards specific exposure types – e.g. SMEs vs large corporates – or specific jurisdictions? How would the result change for different measures of loss rates, such as those based on delinquencies, or for loss rates on traded credit securities?

Finally, it is key to acknowledge that even the best forecast models would be in a position to flag in advance some but not all boom-bust switches. Such models would combine economic reasoning with systematic empirical relationships to forecast busts that are *endogenous* outcomes of excessive risk taking during booms. However, any forecast model would inevitably fail to predict busts that are rooted in inherently *exogenous* shocks, such as the fallout of the Covid-19 outbreak. To be prepared for the outsize losses that stem from such shocks, the financial system should build forecast-independent – that is, precautionary – loss-absorbing resources.

References

Abad J and J Suárez (2017), "Assessing the cyclical implications of IFRS 9 – a recursive model", ESRB Occasional Paper Series No 12.

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

Adrian T, N Boyarchenko and D Giannone (2019), "Vulnerable growth", *American Economic Review*, vol. 109(4), pages 1263–89.

Aikman D, J Bridges, S Hacioglu Hoke, C O'Neill and R Akash (2019), "Credit, capital and crises: a GDP-at-risk approach", Bank of England Working Paper 824.

Aldasoro I, C Borio and M Drehmann (2018), "Early warning indicators of banking crises: expanding the family", BIS Quarterly Review, March 2018.

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

Banerjee R, G Cornelli and E Zakrajšek (2020), "The outlook for business bankruptcies," *BIS Bulletin* No. 30, October.

Banerjee R and E Kharroubi (2020), "The financial vulnerabilities driving firms to the exit," *BIS Quarterly Review*, March.

Basel Committee on Banking Supervision (BCBS) (2017), "Basel III: Finalizing post-crisis reforms", December.

Borio C and P Lowe (2001), "To provision or not to provision", *BIS Quarterly Review*, September.

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

Chae S, R Sarama, C Vojtech and J Wang (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.

Cohen B and G Edwards (2017), "The new era of expected credit loss provisioning", *BIS Quarterly Review*, March.

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

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

Drehmann M and M Juselius (2012), "Do debt service costs affect macroeconomic and financial stability?", *BIS Quarterly Review, September.*

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

Fatouh M and S Giansante (2020), "Expected loss model and the cyclicality of bank credit losses and capital ratios", mimeo.

Financial Accounting Standards Board (FASB) (2016), "Accounting standards update No. 2016–13, Financial instruments – Credit losses (Topic 326): Measurement of credit losses on financial instruments", June.

Gambera M (2000), "Simple forecasts of bank loan quality in the business cycle", *Emerging Issues*, Federal Reserve Bank of Chicago, April issue.

Ghosh A (2015), "Banking-industry specific and regional economic determinants of nonperforming loans: Evidence from US states", *Journal of Financial Stability*, vol. 20, pages 93–104.

Goncharenko R and A Rauf (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*, vol. 12(3), pages 199–232.

Gorton G and G Ordoñez (2014), "Collateral crises", *American Economic Review*, vol 104, no 2, February.

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

Harvey A (1976), "Estimating regression models with multiplicative heteroscedasticity", *Econometrica*, vol. 44(3), pages 461–465.

Huizinga H and L Laeven (2019), "The Procyclicality of banking: Evidence from the euro area", CEPR Discussion Paper 13605.

International Accounting Standards Board (IASB) (2014), "IFRS 9 financial instruments", July.

Jiménez G, S Ongena, J Peydró and J Saurina (2017), "Macroprudential policy, countercyclical bank capital buffers, and credit supply: Evidence from the Spanish dynamic provisioning experiments", *Journal of Political Economy* 125, pages 2126–77.

Jordà O (2005), "Estimation and inference of impulse responses by local projections", *American Economic Review*, vol. 95(1), pages 161–82.

Koenker R (2017), "Quantile regression: 40 years On", *Annual Review of Economics*, vol. 9(1), pages 155–176.

Krüger S, D Rösch and H Scheule (2018), "The Impact of loan loss provisioning on bank capital requirements," *Journal of Financial Stability* 36, pp 114–129.

Loudis B and B Ranish (2019), "CECL and the credit cycle", Finance and Economics

Discussion Series 2019-061, Board of Governors of the Federal Reserve System.

Lu Y and V Nikolaev (2020), "Expected loan loss provisioning: An Empirical model", Chicago Booth Research Paper No. 19-11.

Marcellino M, J Stock and M Watson (2006), "A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series", *Journal of Econometrics*, vol. 135(1-2), pages 499–526.

Peura S and E Jokivuolle (2004), "Simulation based stress tests of banks' regulatory capital adequacy", *Journal of Banking and Finance*, vol. 28(8), pages 1801–1824.

Risk.net (2018a), "Banks grapple with IFRS 9 and CECL loan loss forecasting", March 14.

Risk.net (2018b), "Banks ask Fed to delay CECL impact on stress testing", August 15.

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

Vasicek, O (1991), "Limiting Loan Loss Probability Distribution", KMV Working Paper.

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

A Proofs

We first consider $\Lambda^{\alpha,CK}(\iota_t)$, per (2). Since Φ is an increasing function, $\Lambda^{\alpha,CK}(\iota_t)$ increases in $d(\iota_t)$. In turn, $\Lambda^{\alpha,CK}(\iota_t)$ is guaranteed to be increasing in $\rho(\iota_t)$ if $\Phi^{-1}(\alpha) < 0$ or, equivalently, if $\alpha < 0.5$. This would be always the case in practice since α is the targeted probability of bank failure over a year. Thus, $d(0) < d(1)$ or $\rho(0) < \rho(1)$ implies that a bust calls for higher loss-absorbing resources than a boom.

We then derive equation (4) and show that it has exactly one solution in terms of Λ. Provided that the latest phase is ι_{t-1}, an IK bank knows the following for a given level Λ of LAR:

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

which leads to the left-hand side of (4). Since this expression is monotonic in Λ and converges to 1 (respectively, 0) as $\Lambda \to 0$ (respectively, 1), there is exactly one value of Λ that sets it equal to α.

To derive the IK bank's excess/shortage of LAR, we show that $\Lambda^{\alpha,IK}(\iota_{t-1})$ is sandwiched as follows: $\Lambda^{\alpha,IK}(\iota_{t-1}) \in \left(\Lambda^{\alpha,CK}(0), \Lambda^{\alpha,CK}(1)\right)$. Since each of the two summands on the left-hand side of (4) is strictly decreasing in Λ, the left-hand side of (4) is smaller (respectively, larger) than α if $\Lambda^{\alpha,IK}(\iota_{t-1}) \geq \Lambda^{\alpha,CK}(1)$ (respectively, if $\Lambda^{\alpha,IK}(\iota_{t-1}) \leq \Lambda^{\alpha,CK}(0)$). Finally, since (4) has exactly one solution, we obtain the desired result.

To derive the dependence of the IK bank's PD on a through-the-cycle loading on

G_t, we set $\rho(0) = \rho(1) = \rho$ while preserving $d(1) > d(0)$. When an IK bank's LAR is equal to Λ and the phase is ι_t, this bank defaults in period t with probability $PD^{IK} = \Phi\left(\frac{d(\iota_t)-\sqrt{1-\rho^2}\Phi^{-1}(\Lambda)}{\rho}\right)$, which implies $\Lambda = \Phi\left(\frac{d(\iota_t)-\rho\Phi^{-1}\left(PD^{IK}\right)}{\sqrt{1-\rho^2}}\right)$. Referring to (4), imposing $\iota_{t-1} = 0$ and $\iota_t = 1$, and using the latter equation to substitute for Λ delivers:

$$\frac{d(1) - d(0)}{\rho} = \Phi^{-1}\left(PD^{IK}\right) - \Phi^{-1}\left(\frac{\alpha - (1 - \pi(0))\, PD^{IK}}{\pi(0)}\right),$$

where the left-hand side is strictly positive and decreases in ρ, and the right-hand side increases in PD^{IK}. Thus, this expression implies $dPD^{IK}/d\rho < 0$, which confirms that an IK bank's PD is higher for a lower ρ.

To show that an improvement in the IK bank's information reduces $\left|EL^{CK}(1) - EL^{IK}(0)\right|$ and $\left|UL^{\alpha,CK}(1) - UL^{\alpha,IK}(0)\right|$, we first note that such an improvement is equivalent to decreasing $\pi(0)$ in period t when $\iota_t = 1$. It is immediate that $\partial\left|EL^{CK} - EL^{IK}\right|/\partial\pi(0) = \partial\left(EL^{CK} - EL^{IK}\right)/\partial\pi(0) > 0$ and $\lim_{\pi(0)\to 0}\left(EL^{CK} - EL^{IK}\right) = 0$. Then, we note that, by the earlier sandwiching result, (4) implies $d\Lambda^{\alpha,IK}(0)/d\pi(0) < 0$ and $\lim_{\pi(0)\to 0}\Lambda^{\alpha,IK}(0) = \Lambda^{\alpha,CK}(1)$. The desired result in terms of UL follows immediately.

Finally, we show that, for a sufficiently low $\rho(\iota_{t-1})$ and $\rho(\iota_t)$, the IK bank would enter a boom-bust switch – whereby $\iota_{t-1} = 0$ and $\iota_t = 1$ – with a higher perceived UL than the CK bank. As $\rho(\iota_t) \to 0$, $UL^{\alpha,CK}(\iota_t) \to 0$. In turn, when $\rho(0) \to 0$ and $\rho(1) \to 0$, the IK bank perceives probability of $1 - \pi(\iota_{t-1})$ that date-t losses will be higher than $EL^{IK}(\iota_{t-1})$. Thus, as long as $1 - \pi(\iota_{t-1}) \geq \alpha$ – i.e. as long as the perceived probability of switching phases is higher than the targeted PD for the bank – $\lim_{\substack{\rho(0)\to 0\\ \rho(1)\to 0}} UL^{\alpha,IK}(\iota_{t-1}) > 0$. By continuity, this implies that there exists $\underline{\rho} > 0$ such that $UL^{\alpha,IK}(0) > UL^{\alpha,CK}(1)$ – that is, the IK bank is over-capitalized – provided that $\rho(0) < \underline{\rho}$ and $\rho(1) < \underline{\rho}$.

B Data

Variables	Definition	Sample	Source
μ_t	Charge-Off on total Loans and Leases at Commercial Banks	85q1-19q2	Fed
DSR_t	Debt service ratio	79q1-18q4	BIS
$C2Y_t$	Deviation between credit-to-GDP ratio and its one-sided HP-filtered trend ($\lambda = 400000$)	79q1-18q4	BIS
HPG_t	Deviation between real residential property price and and its one-sided HP-filtered trend ($\lambda = 400000$)	84q1-18q4	BIS
UR_t	Unemployment rate	79q1-18q4	IMF
OG_t	Deviation between real GDP and its one-sided HP-filtered trend ($\lambda = 1600$)	82q1-19q1	National data
CS_t	Spread between Seasoned Baa and Aaa Corporate Bond Yields	79q1-19q3	Moody's
TS_t	Spread between 10 year and 1 one year yeild on government bonds	79q1-19q3	Fed
RIR	Difference between the 3-month federal funds rate (annualized) and yoy CPI index growth	80q1-19q1	Datastream, national data
VXO_t	CBOE S&P 100 Volatility Index	86q1-20q2	CBOE

Table B.1: **Variable definitions, data sources, and sample period.** Sources: Federal reserve board (Fed), Bank for International Settlements (BIS), and Chicago Board Options Exchange (CBOE).

C Additional results

z_{1t} :	DSR_t	DSR_t	UR_t	UR_t
z_{2t} :		OG_t	OG_t	OG_t
z_{3t} :		TS_t	HPG_t	HPG_t
z_{4t} :				CS_t
z_{5t} :				TS_t
z_{6t} :				VXO_t
h	Root mean square error			
1	0.14	0.14	**0.12**	0.13
4	0.47	0.48	**0.44**	0.45
8	0.69	0.71	**0.64**	0.69
12	**0.71**	0.73	0.90	0.84
	Correlation (realized vs. forecasted)			
1	0.98	0.98	**0.99**	0.98
4	**0.79**	0.78	0.78	0.77
8	**0.60**	0.57	0.51	0.38
12	**0.58**	0.54	0.26	0.14
	Excess/insufficient turnings			
1	**0**	1	**0**	1
4	**0**	**0**	1	4
8	**0**	**0**	1	5
12	**0**	**0**	1	4
	Standardized average distance			
1	0.75	**0.50**	1.00	0.75
4	0.92	0.67	0.58	**0.42**
8	0.42	**0.33**	0.67	0.38
12	0.17	**0.11**	0.53	0.47
# best	**8**	6	5	1

Table C.1: **First moments, with multiple forecast variables.** Evaluation of forecasted quarterly loss rates based on OLS estimates of equation (8), with lag parameter $K = 1$ and horizon h as indicated. The forecast variables in z_t are subsets of the following list: DSR_t = debt service ratio, $C2Y_t$ = credit-to-GDP gap, HPG = house price gap, UR_t = unemployment rate, OG_t = output gap, CS_t = corporate bond spread, TS_t = term spread, RIR_t = real interest rate, or VXO_t = volatility index. The row "# best" counts the number of times an indicator is the best performer (boldface values). The first column reproduces the results for $z_t = DSR_t$ from *Table 1*.

z_{1t} :	DSR_t	DSR_t	$C2Y_t$	HPG_t	UR_t	OG_t	CS_t	TS_t	RIR_t	VXO_t
z_{2t} :	–	DSR_t^2	$C2Y_t^2$	HPG_t^2	UR_t^2	OG_t^2	CS_t^2	TS_t^2	RIR_t^2	VXO_t^2
h	Root mean square error									
1	0.14	0.13	0.13	0.14	0.14	**0.12**	0.20	0.13	0.15	0.14
4	0.47	0.44	0.51	0.65	0.46	**0.41**	0.77	0.48	0.56	0.54
8	0.69	0.65	0.90	0.79	0.69	0.72	0.86	**0.64**	0.79	0.72
12	0.71	**0.70**	0.92	1.06	0.78	0.84	0.86	0.73	0.87	0.91
	Correlation (realized vs. forecasted)									
1	0.98	0.98	0.98	0.98	0.98	0.99	**0.96**	0.98	0.98	0.98
4	0.79	**0.86**	0.70	0.52	0.76	0.82	0.13	0.74	0.65	0.65
8	**0.60**	0.48	-0.06	0.41	0.38	0.26	-0.17	0.55	0.28	0.24
12	**0.58**	0.53	-0.03	-0.01	0.04	0.03	0.01	0.37	0.07	-0.23
	Excess/insufficient turnings									
1	**0**	**0**	**0**	**0**	**0**	1	1	1	1	**0**
4	**0**	1	2	2	1	1	2	1	2	1
8	**0**	1	3	1	1	3	2	**0**	3	2
12	**0**	**0**	2	1	1	3	4	2	1	2
	Standardized average distance									
1	0.75	0.75	1.00	1.00	0.75	**0.50**	0.75	0.75	1.00	0.75
4	0.92	0.92	0.83	0.75	1.00	0.92	0.75	**0.58**	0.67	0.92
8	0.42	0.46	0.79	0.50	0.88	0.54	0.54	**0.38**	0.54	0.75
12	**0.17**	0.19	0.89	0.75	0.69	0.47	0.47	0.33	0.33	0.75
# best	**7**	4	1	1	1	3	1	4	0	1

Table C.2: **First moments, with quadratic terms.** Evaluation of forecasted quarterly loss rates based on OLS estimates of equation (8) with lag parameter $K = 1$ and horizon h as indicated. The forecast variables (in addition to auto-regressive terms) are $z_t = (z_{1t}, z_{2,t})'$, where , $z_{2t} = z_{1t}^2$ and z_{1t} is: DSR_t = debt service ratio, $C2Y_t$ = credit-to-GDP gap, HPG = house price gap, UR_t = unemployment rate, OG_t = output gap, CS_t = corporate bond spread, TS_t = term spread, RIR_t = real interest rate, or VXO_t = volatility index. The row "# best" counts the number of times an indicator is the best performer (boldface values). The first column reproduces the results for $z_t = DSR_t$ from *Table 1.*

z_{1t} :	DSR_t	DSR_t	$C2Y_t$	HPG_t	UR_t	OG_t	CS_t	TS_t	RIR_t	VXO_t
z_{2t} :	–	DSR_t^+	$C2Y_t^+$	HPG_t^+	UR_t^+	OG_t^+	CS_t^+	TS_t^+	RIR_t^+	VXO_t^+
h					Root mean square error					
1	0.14	0.13	0.13	0.14	0.13	**0.12**	0.14	0.13	0.14	0.13
4	0.47	0.43	0.47	0.55	0.46	**0.40**	0.56	0.48	0.56	0.50
8	0.69	**0.63**	0.71	0.78	0.68	0.71	0.83	0.64	0.78	0.73
12	**0.71**	**0.71**	0.73	1.03	0.78	0.82	0.84	0.74	0.87	0.92
					Correlation (realized vs. forecasted)					
1	**0.98**	**0.98**	**0.98**	**0.98**	**0.98**	0.99	**0.98**	**0.98**	**0.98**	**0.98**
4	0.79	**0.87**	0.75	0.63	0.77	0.83	0.60	0.74	0.66	0.71
8	**0.60**	0.57	0.27	0.36	0.39	0.31	-0.09	0.55	0.30	0.21
12	**0.58**	0.52	0.33	-0.01	-0.04	0.12	0.07	0.34	0.09	-0.29
					Excess/insufficient turnings					
1	**0**	**0**	**0**	**0**	**0**	1	1	1	1	**0**
4	**0**	**0**	2	1	1	1	2	1	2	1
8	**0**	1	3	1	1	1	4	3	3	3
12	**0**	**0**	2	1	**0**	3	4	**0**	1	2
					Standardized average distance					
1	0.75	**0.50**	1.00	1.00	0.75	**0.50**	0.75	0.75	1.00	0.75
4	0.92	0.50	0.67	0.67	1.00	0.92	**0.42**	0.58	0.58	1.00
8	0.42	0.46	0.79	**0.33**	0.67	0.71	0.38	0.38	0.63	0.79
12	**0.17**	**0.17**	0.83	0.39	0.69	0.47	0.47	0.39	0.33	0.61
# best	**9**	**9**	2	3	3	3	2	2	1	2

Table C.3: **First moments with asymmetric effects.** Evaluation of forecasted quarterly loss rates based on OLS estimates of equation (8) with lag parameter $K = 1$ and horizon h as indicated. The forecast variables (in addition to auto-regressive terms) are $z_t = (z_{1t}, z_{2,t})'$, where $z_{2t} = z_{1t}^+ = I_{z_{1t} > \bar{z}_{1t}} \cdot z_{1t}$, $\bar{z}_{1t}$ is a real-time mean and z_{1t} is: DSR_t = debt service ratio, $C2Y_t$ = credit-to-GDP gap, HPG = house price gap, UR_t = unemployment rate, OG_t = output gap, CS_t = corporate bond spread, TS_t = term spread, RIR_t = real interest rate, or VXO_t = volatility index. The row "# best" counts the number of times an indicator is the best performer (boldface values). The first column reproduces the results for $z_t = DSR_t$ from *Table 1*.

z_t :	DSR_t	DSR_t	$C2Y_t$	HPG_t	UR_t	OG_t	CS_t	TS_t	RIR_t	VXO_t
Est.:	OLS	NLS	NLS	NLS	NLS	NLS	NLS	NLS	NLS	NLS
h					Root mean square error					
1	**0.14**	0.18	0.33	0.29	0.30	0.27	0.46	0.34	0.38	0.32
4	0.47	0.48	0.59	0.58	0.49	**0.42**	0.56	0.62	0.76	0.58
8	0.69	**0.60**	0.65	0.71	0.68	0.72	0.76	0.68	0.77	0.73
12	0.71	**0.66**	0.70	0.94	0.78	0.83	1.10	0.72	0.86	0.87
					Correlation (realized vs. forecasted)					
1	**0.98**	0.97	0.94	0.95	0.93	0.95	0.92	0.93	0.92	0.95
4	0.79	**0.82**	0.74	0.71	0.73	0.81	0.77	0.64	0.53	0.73
8	0.60	**0.67**	0.50	0.55	0.39	0.22	0.27	0.48	0.16	0.16
12	**0.58**	**0.58**	0.45	0.22	0.05	-0.07	-0.06	0.35	-0.01	-0.26
					Excess/insufficient turnings					
1	**0**	**0**	0	0	0	0	1	1	1	1
4	**0**	**0**	1	3	1	0	2	1	2	2
8	**0**	**0**	1	1	0	1	2	2	1	2
12	**0**	**0**	1	1	0	3	2	0	1	2
					Standardized average distance					
1	0.75	**0.50**	0.75	1.00	0.75	0.75	1.00	0.75	1.00	1.00
4	0.92	0.83	1.00	**0.42**	1.00	**0.42**	**0.42**	0.67	0.83	0.75
8	0.42	0.46	0.67	0.58	0.79	0.79	0.88	**0.33**	0.63	0.88
12	**0.17**	0.39	0.89	0.97	0.58	0.47	0.72	0.83	0.44	0.69
# best	8	**10**	0	1	0	2	1	1	0	0

Table C.4: **First moments with exponential specification.** Evaluation of forecasted quarterly loss rates based on NLS estimates of $\mu_{t+h|t} = \exp\left(\sum_{k=0}^{K}(\beta_{h,k}\mu_{t-k} + \gamma'_{h,k}z_{t-k})\right) + v_{t+h|t}$ with lag parameter $K = 1$ and horizon h as indicated. The forecast variable in z_t is: DSR_t = debt service ratio, $C2Y_t$ = credit-to-GDP gap, HPG = house price gap, UR_t = unemployment rate, OG_t = output gap, CS_t = corporate bond spread, TS_t = term spread, RIR_t = real interest rate, or VXO_t = volatility index. The row "# best" counts the number of times an indicator is the best performer (boldface values). The first column reproduces the results for $z_t = DSR_t$ from *Table 1*.

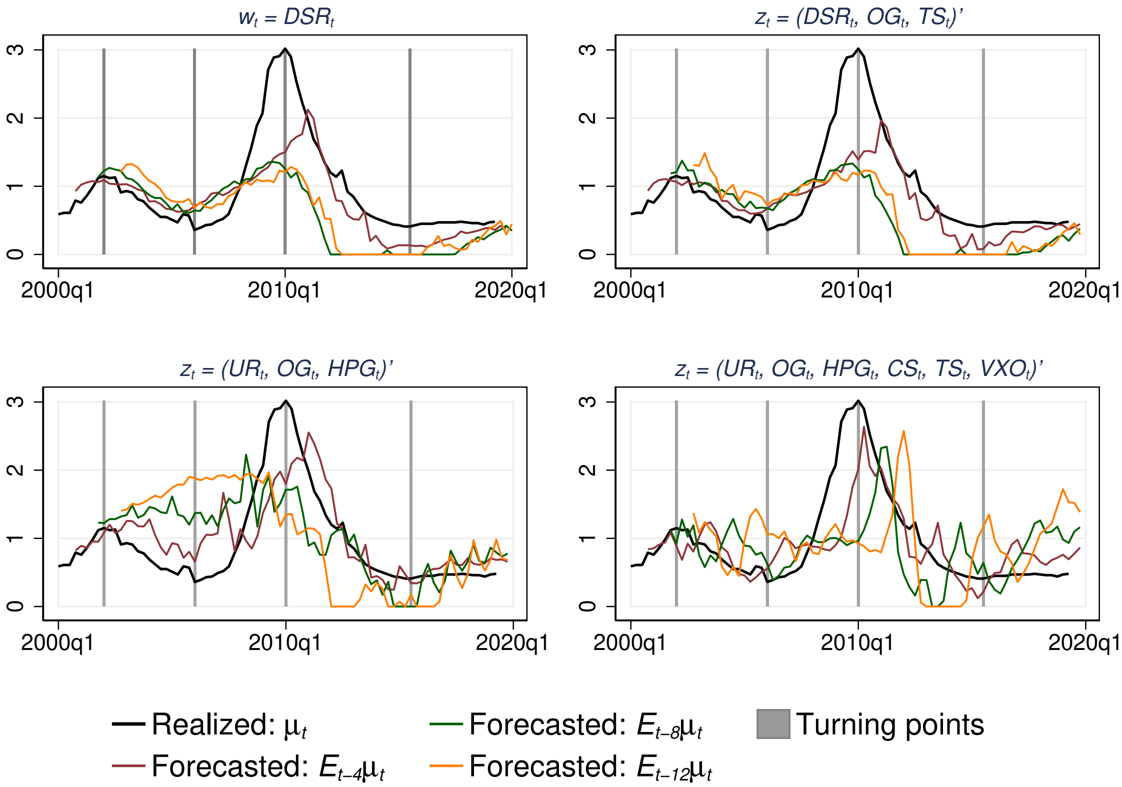

Figure C.1: **First moments, with multiple forecast variables.** Realized and forecasted quarterly loss rates based on OLS estimates of equation (8) with lag parameter $K = 1$. The forecast variables in z_t are subsets of the following list: DSR_t = debt service ratio, HPG_t = house price gap, UR_t = unemployment rate, OG_t = output gap, CS_t = corporate bond spread, TS_t = term spread, or VXO_t = volatility index. An observation in quarter t is equal to: a real-time forecast formed at $T = \{t-4; t-8; t\text{–}12\}$ or the corresponding realization at t. The top-left panel reproduces the results for $z_t = DSR_t$ from *Figure 3*.

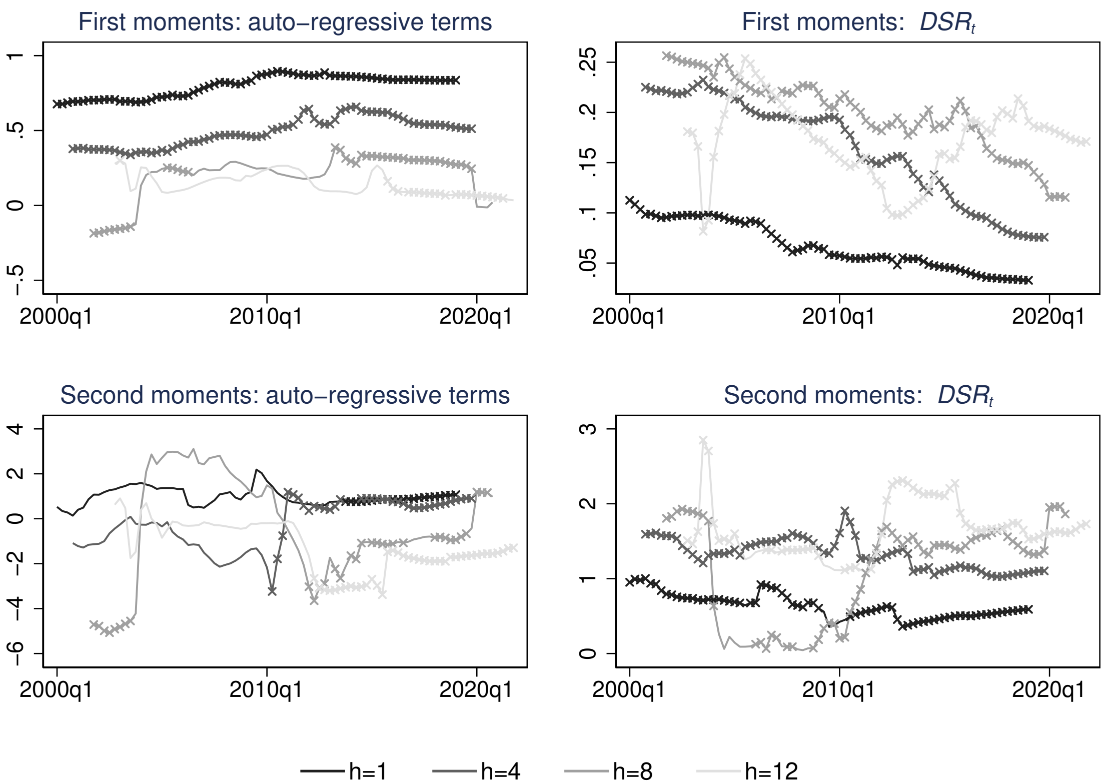

Figure C.2: **Stability of estimates at different horizons.** Sum of coefficients over $k = \{0, 1\}$ on auto-regressive terms and forecast variables per equations (8) and (10), where $z_t = w_t = DSR_t =$ debt service ratio, lag length $K = 1$ and horizon h as indicated. An observation in quarter T refers to a real-time estimate at T. Statistical significance at the 5% level is denoted with "x".

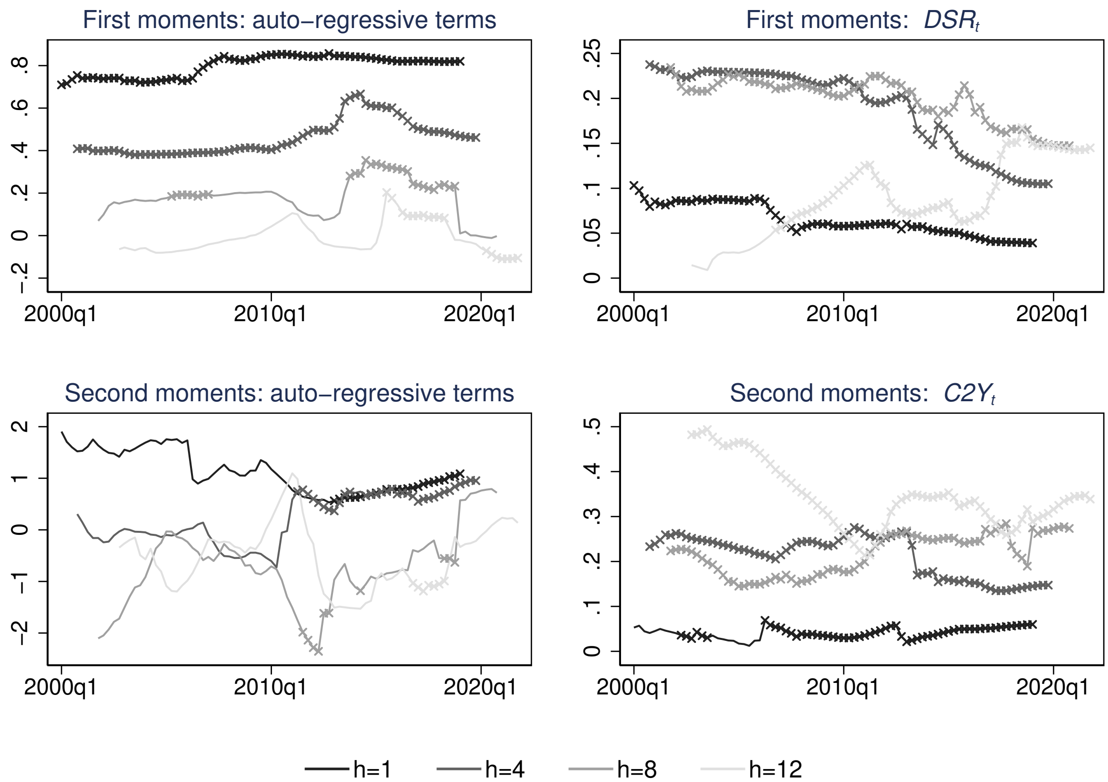

Figure C.3: **Stability of estimates at different horizons.** Sum of coefficients over $k = \{0, 1\}$ on auto-regressive terms and forecast variables per equations (8) and (10), where $z_t = DSR_t =$ debt service ratio, $w_t = C2Y_t =$ credit-to-GDP gap, lag length $K = 1$ and horizon h as indicated. An observation in quarter T refers to a real-time estimate at T. Statistical significance at the 5% level is denoted with "x".

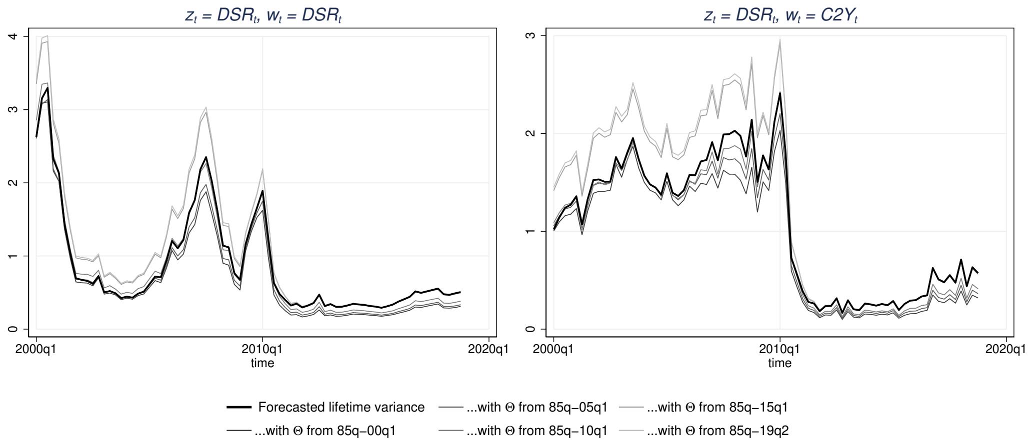

Figure C.4: **Stability of correlation matrix.** Forecasted lifetime variance based on estimating the correlation matrix, Θ, (i) in real-time (black solid line) or (ii) from fixed samples as indicated in the legend (lines in various scales of gray). The model is estimated per (8) and (10) with lag parameter $K = 1$ and forecast variables (in addition to auto-regressive terms) as follows: $z_t = w_t = DSR_t =$ debt service ratio (left panel) and with $z_t = DSR_t$, $w_t = C2Y_t =$ credit-to-GDP gap (right panel).

D Direct forecasts of lifetime loss rates

This appendix is in two parts. First, we derive *direct* forecasts of lifetime portfolio losses and the variance of forecast errors. This is in contrast to the main text, where we build first- and second-moment forecasts of lifetime losses from corresponding forecasts of quarterly loss rates. Second, we forecast percentiles of the probability distribution of lifetime losses. The forecast of a high percentile, together with a first-moment forecast, delivers unexpected losses (UL), in line with regulatory texts.

We start by forecasting lifetime perfect-foresight losses, Λ_t^{PF}, constructed as in Section 4.3. Given a sample of $t = 1, ..., T$ observations on μ_t, we construct Λ_t^{PF} for $t = 1, ..., T - H$, where the last lifetime loss, Λ_{T-H}^{PF}, uses observations up to T. For a direct real-time forecast of Λ_T^{PF}, we estimate the following versions of (8) and (10):

$$\Lambda_t^{PF} = \sum_{k=0}^{K} \left(\beta_k^D \mu_{t-k} + \gamma_k^{D\prime} z_{t-k} \right) + \upsilon_t^D \tag{12}$$

$$E\left((\upsilon_t^D)^2 \right) = \exp\left(\sum_{k=0}^{K} \left(\delta_k^D \mu_{t-k} + \theta_k^{D\prime} w_{t-k} \right) \right) \tag{13}$$

for $t = 1, ..., T - H$ and evaluate them at $t = T$. Two remarks are in order. First, while (12) and (13) may appear as contemporaneous relationships, Λ_t^{PF} embeds quarterly loss rates up to H quarters in the future. Second, since the one-quarter lag of the dependent variable in (12) is not available in real time, we replace it with the latest available quarterly loss rate, μ_t.

Targeting lifetime losses directly yields real-time forecasts that are comparable to those in the main text in terms of first moments but inferior in terms of second moments. Again, the best results in the current context stem from specifications in which the the forecast variable for first moments is the debt service ratio (DSR_t) and that for second moments is either DSR_t or the credit-to-GDP gap ($C2Y_t$). The similarity of the first-moment results, between each of these specifications and the corresponding ones in the main text are seen by comparing the left-hand panels of *Table D.1* and *Table 3*, as well as *Figure D.1* (red lines) and *Figure 6* (top panels). In turn, the deterioration of second-moment forecasts surfaces most clearly in *Figure D.1* (blue lines), where the implied loss-absorbing resources (LAR) remain consistently and materially below the GFC peaks, in contrast to the capital implied by the aggregation of quarterly loss forecasts (*Figure 7*). Perhaps unsurprisingly, we loose information about the variability

		Lifetime EL				Variance of lifetime forecast errors		# best
z_t	w_t	RMSE	Corr.	Turn.	Dist.	RMSE	Corr.	
DSR_t	DSR_t	**1.07**	**0.81**	**0.00**	**0.33**	**2.80**	-0.12	**5**
DSR_t	$C2Y_t$	1.17	0.80	**0.00**	**0.33**	2.99	**0.79**	3

Table D.1: **Lifetime expected loss (EL) and error variance, direct forecasts.** The evaluations compare perfect-foresight lifetime magnitudes with their direct forecasts, per equations (12) and (13). The underlying forecast variables (in addition to μ_t) are: DSR_t = debt service ratio and $C2Y_t$ = credit-to-GDP gap. The column "# best" counts the number of times an indicator is the best performer within each portfolio (boldface values).

of losses when we "smooth" the data via (12) and (13) prior to forecasting.

In the current context, we can also forecast specific percentiles of the lifetime loss distribution, in line with recent "at risk" assessments (eg Adrian et al, 2019, and Aikman et al, 2019). We do so on the basis of a linear specification for the τ^{th} conditional quantile of this distribution

$$\Lambda_t^{PF} = \sum_{k=0}^{K} \left(\beta_{\tau,k}^{Q} \mu_{t-k} + \gamma_{\tau,k}^{Q\prime} z_{t-k} \right) + \upsilon_{\tau,t}^{Q} \tag{14}$$

which we estimate by solving (eg Koenker, 2017):

$$\underset{\xi}{\operatorname{argmin}} \sum_{t=K+1}^{T-H} \left(\tau I_{\upsilon_{\tau,t}^{Q} \geq 0} \left| \upsilon_{\tau,t}^{Q} \right| + (1-\tau) I_{\upsilon_{\tau,t}^{Q} < 0} \left| \upsilon_{\tau,t}^{Q} \right| \right) \tag{15}$$

where $\xi = (\beta_{\tau,0}^{Q}, ..., \beta_{\tau,K}^{Q}, \gamma_{\tau,0}^{Q\prime}, ..., \gamma_{\tau,K}^{Q\prime})'$ and I is the indicator function.

We obtain poor forecasts of the 95^{th} percentile of the lifetime loss distribution (*Figure D.2*, green lines). For one, these forecasts undershoot perfect-foresight losses most of the time. This is in contrast to the unexpected loss (UL) that we set equal to two times the forecasted standard deviations in the main text (*Figure 6*). In addition, there are instances in which the forecasts of the 95^{th} percentile are lower than the corresponding median forecasts (*Figure D.2*, orange lines). While this latter result may seem counter-intuitive, it stems from the fact that the forecasts are *real-time* and, especially those of the 95^{th} percentile, rely on a small part of an already short sample.

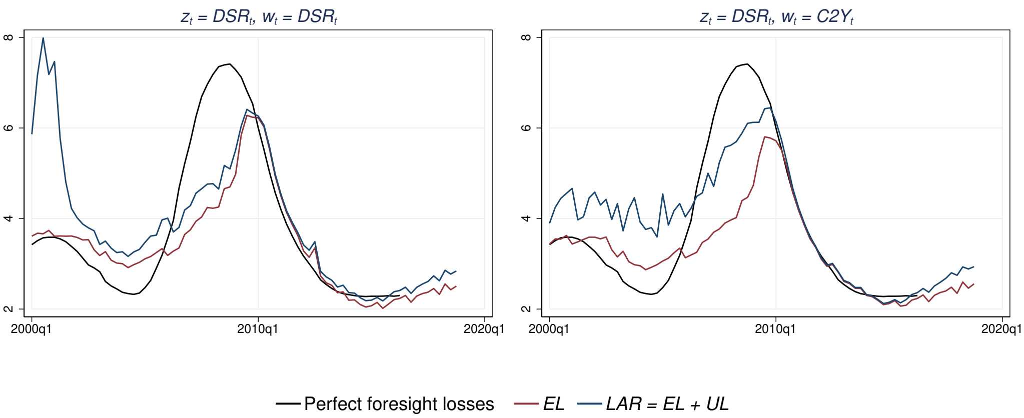

Figure D.1: **Direct forecasts of lifetime losses.** EL = expected losses, UL = unexpected losses set to two times the forecasted standard deviation, and loss-absorbing resources = LAR. An observation in quarter T is equal to: a real-time forecast at T or the corresponding materialization of risk from T onward. Based on estimating equations (12) and (13), where the underlying forecast variables (in addition to μ_t) are: DSR_t = debt service ratio (in z_t for first moments) and DSR_t or $C2Y_t$ = credit-to-GDP gap (in w_t for second moments).

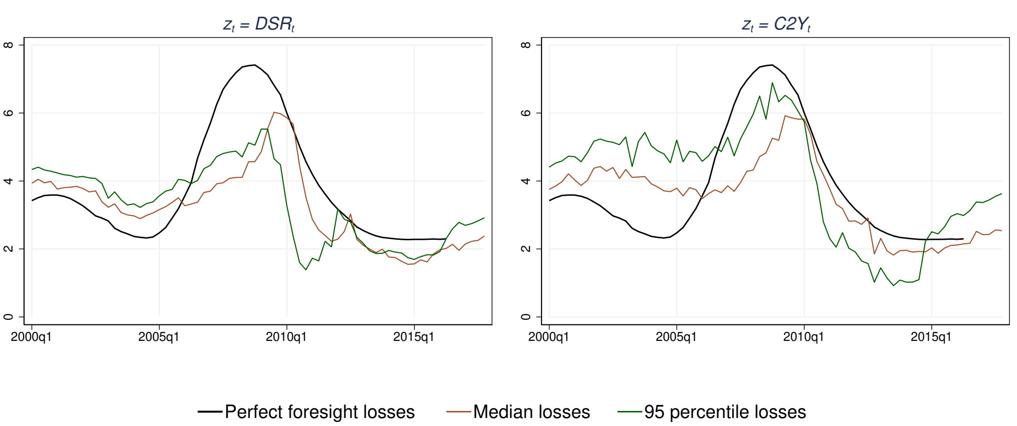

Figure D.2: **Median, and 95^{th} percentile of lifetime losses, direct estimates.** Based on quantile regressions. The underlying forecast variables (in addition to μ_t) are: DSR_t = debt service ratio and $C2Y_t$ = credit-to-GDP gap.

Previous volumes in this series

All volumes are available on our website www.bis.org.