Bank of Finland

BoF Economics Review

3 • 2021

Could corporate credit losses turn out higher than expected?

Mikael Juselius, Bank of Finland
Nikola Tarashev, Bank of International Settlements

Abstract

While corporate credit losses have been low since the start of the Covid-19 pandemic, their future evolution is quite uncertain. Using a forecasting model with a solid track record, we find that the baseline scenario ("expected losses") is benign up to 2024. This is due to policy support measures that have kept debt service costs low. However, high indebtedness, built up when the pandemic impaired real activity, suggests increased tail risks: plausible deviations from the baseline scenario ("unexpected losses") feature ballooning corporate insolvencies. Taken at face value, the low expected loss forecasts are consistent with low bank provisions, whereas the high unexpected loss forecasts call for substantial capital.

JEL classification: E44, E47, E65, G17, G21

The authors are grateful to Iñaki Aldasoro, Claudio Borio, Catherine Casanova, Stijn Claessens, Neil Esho, Jon Frost, Pat McGuire, Hyun Song Shin, Daniel Rees and two anonymous Bank of Finland reviewers for their comments on earlier drafts, to Anamaria Illes and Thomas Shen for excellent research assistance, and to Louisa Wagner for administrative support.

Introduction

While the Covid-19 pandemic inflicted a profound shock on the world economy, it has been a non-event in terms of corporate credit losses. The breakdown of the historical relationship between economic activity and the concurrent incidence of bankruptcies – the "bankruptcy gap" – is largely driven by the policy response to the pandemic. Tax deferrals, cash transfers, furlough schemes, debt moratoriums and equity-like injections backstopped the loss of private cash flows after March 2020. In addition, public guarantees and unprecedented monetary policy measures – aided by a strong financial sector and a temporary loosening of prudential regulation – supported the supply of credit at low interest rates. In sum, authorities provided a financial lifeline to corporates while the real economy was in a coma.

However, the unprecedented policy lifeline has resulted in high corporate indebtedness, which is a potential roadblock on the path to recovery. A key concern is that the bankruptcy gap could be followed by a delayed insolvency wave (Banerjee, Noss and Vidal Pastor (2021)). Hence there is value in estimating future credit losses on the basis of information about both liquidity conditions and solvency risks.

To forecast losses on corporate loans, we use two financial cycle indicators and assume that they have preserved their traditional relationship with credit risk. The first indicator is the private non-financial sector's debt service ratio (DSR), which captures cash flow strains. The second is the corresponding gap between the credit-to-GDP ratio and its long-term trend, which signals whether excessive leverage is building up in the household and corporate sectors. Over the past 20 years, these indicators have had a solid track record as forecasters of baseline (or "expected") losses and extreme but plausible deviations from the baseline (or "unexpected" losses) up to three years in advance (Juselius and Tarashev (2020)).

Our findings underscore the value of forecasting both expected losses and unexpected losses. On the one hand, the support of cash flows and policy measures to reduce interest payments kept the DSR low through Q1 2021. In a baseline scenario with sustained easy financial conditions, this augurs well for banks' recent loan portfolios, ie expected losses seem to be low. On the other hand, the credit-to-GDP gap spiked as indebtedness rose on the back of easy financial conditions while economic activity dropped after the start of the pandemic. Large credit-to-GDP gaps suggest strong risk-taking and have traditionally gone hand in hand with an increased likelihood of marked departures from baseline scenarios, stemming for instance from creditor retrenchment. This decoupling of unexpected losses from expected losses indicates that, while the bankruptcy gap may persist for some time, there is a high risk of its closing over the next two to three years. The decoupling is unusual. Historically, expected and unexpected losses have tended to move in the same direction, thus providing some justification to the common practice of inferring changes in unexpected losses from changes in expected losses. At the current juncture, such inference could leave banks with inadequate loss-absorbing resources.

Expected and unexpected losses: two key aspects of the loss distribution

The risk associated with future credit loss rates – concretely, net charge-offs divided by the underlying credit volume – is embodied in a probability distribution, such as the hypothetical ones in Graph 1. A loss distribution conveys the likelihood of loss rates being in a given range over some time horizon. Two aspects of a loss distribution are of particular importance: its mean and its dispersion.

The mean of a distribution (or its centre of gravity) is a number that summarises the severity of a baseline scenario, or expected losses. A forecast of the expected losses is what drives banks' provisions, which reduce reported profits and hence capital. Even if this forecast adequately reflects all relevant information, the loss rate that eventually materialises could still differ substantially from it.

The greater the distribution's dispersion, the higher the likelihood of a large wedge between realised and expected losses and, thus, of losses that banks have not provisioned for. Capital is there to absorb such unexpected losses. For concreteness, we define them as belonging to the loss range that starts at the expected level and extends two standard deviations rightwards (Graph 1, dotted whiskers). Thus, a forecast of unexpected losses stands for a forecast of the loss distribution's dispersion.

Graph 1. Hypothetical loss distributions

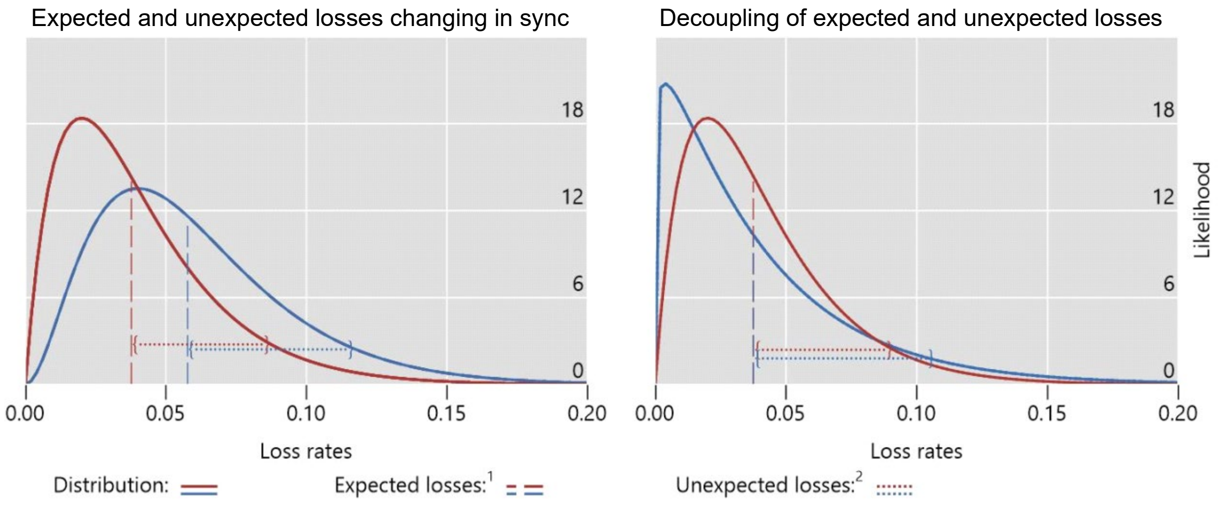

[1] Mean of the loss distribution. [2] Length of the line segment ("whisker") is equal to two standard deviations of the respective distribution.
Source: Authors' calculations.

To assess banks' loss-absorbing resources, prudential supervisors, risk managers and market participants need to monitor changes in the loss distribution (Graph 1, red vs blue lines) and adjust their risk estimates accordingly. A change may drive an increase in both expected and unexpected losses (left-hand panel). This was the case in the run-up to the Great Financial Crisis (GFC), when high risk-taking both increased the severity of the baseline scenario and raised the spectre of market disruptions that underpin extreme outcomes. When the link between expected and unexpected losses is tight, there might be a temptation to infer the latter from the former, which is typically easier to forecast.[1] However, it is also possible that, say, higher uncertainty raises unexpected losses while expected losses are stable (right-hand panel). We will argue that such a decoupling could be a fair description of the current environment.

Expected and unexpected losses may reflect different features of the economy that need to be captured by distinct indicators. After explaining our choice of two indicators of financial vulnerabilities, we use them in an econometric model to obtain forecasts of the possible evolutions of loss rates. Since we need to work with long time series, we derive such forecasts based only on US data.

Debt service ratio for forecasting expected losses

The first indicator is the debt service ratio (DSR), defined as the sum of interest payments and amortisations divided by disposable income (Drehmann et al (2015)). We use BIS data on the DSR of the US private non-financial (PNF) sector, which comprises households and corporates (Graph 2, first panel). Since these data do not employ direct observations of amortisations, they are imputed from estimates of the average interest rate and the outstanding volume of loans, with the assumption that repayments are constant over the time to maturity (held fixed for data quality reasons).

The DSR relates directly to expected credit losses. Debtors with a high DSR are in a precarious situation, with little capacity to absorb an increase in the loan rate or a drop in their income stream and still service their obligations. The DSR is thus in the spirit of micro data indicators that are based on cash balances and financial expenses and have been used recently to study private sector vulnerabilities to cash flow shortages (Gourinchas et al (2020) and Banerjee, Noss and Vidal Pastor (2021)).

Evidence from the past 26 years reveals that the DSR of the overall PNF sector provides, on average, useful information about future quarterly loss rates on commercial and industrial

[1] The global framework for banks' minimum regulatory requirements is predicated on a tight link between expected and unexpected losses on corporate debt. This stems from the assumption that obligors' probabilities of default are the key driver of both expected losses and the shape of the underlying loss distribution – that is, unexpected losses (BCBS (2017), pp 62–3).

loans – henceforth “corporate loans” – extended by banks in the United States (Graph 2, first and second panels). This property indicates that corporates’ creditworthiness depends not only on their own debt service burden but also on that of their ultimate clients, households.[2] It also signals that the DSR differs from indicators of real economic activity – such as the unemployment rate – that historically tended to co-move contemporaneously with loss rates. It is the breakdown of this co-movement during the pandemic that has led to the term “bankruptcy gap” (first panel, shaded area). At low levels before the pandemic, the DSR does not point to such a gap (second panel, darker dots).

Credit-to-GDP gap for forecasting unexpected losses

For unexpected losses, we consider a second indicator: the credit-to-GDP gap (Graph 2, third panel). The rationale for turning to the credit-to-GDP gap is that the DSR may not work well in flagging large deviations from expected losses. In particular, it rests on amortisations predicated on borrowers’ sustained access to credit markets, a condition that is likely to hold only in tranquil times. By contrast, the credit-to-GDP gap does not hinge on such a condition. It measures whether indebtedness in the PNF sector – reflecting both bank- and market-based credit – builds up too fast relative to economic activity. The credit-to-GDP gap thus captures outright solvency issues as well as the potential for creditor retrenchment when a deterioration in underwriting standards comes to the fore. It would also flag that – when leverage is high – the fallout from creditor retrenchment would be particularly severe, even if it happens for external reasons. This explains why the credit-to-GDP gap helps forecast extreme events, such as banking crises (Borio and Lowe (2002), Drehmann and Juselius (2014)).

Indeed, the evidence indicates that the US credit-to-GDP gap has tended to be higher than usual ahead of extreme losses on corporate loans (Graph 2, fourth panel). To identify such outcomes, we compute the squared deviations between realised losses and the expected loss estimates provided by the fitted line in the second panel. Forecasting such deviations delivers measures of the dispersion of the loss distribution and hence of unexpected losses.

[2] *Loss rates on US corporate loans correlate more strongly with the PNF sector’s DSR than with the non-financial corporate sector’s DSR: 80% vs 68%. The PNF sector’s DSR helps predict corporate loan losses in other countries as well: the corresponding correlation is 27% for the United Kingdom (Bank of England data, Q1 1993–Q4 2020) and 83% for Italy (Bank of Italy data, Q1 2006–Q4 2020).*

Graph 2. Loan losses, financial vulnerabilities and economic activity[1]

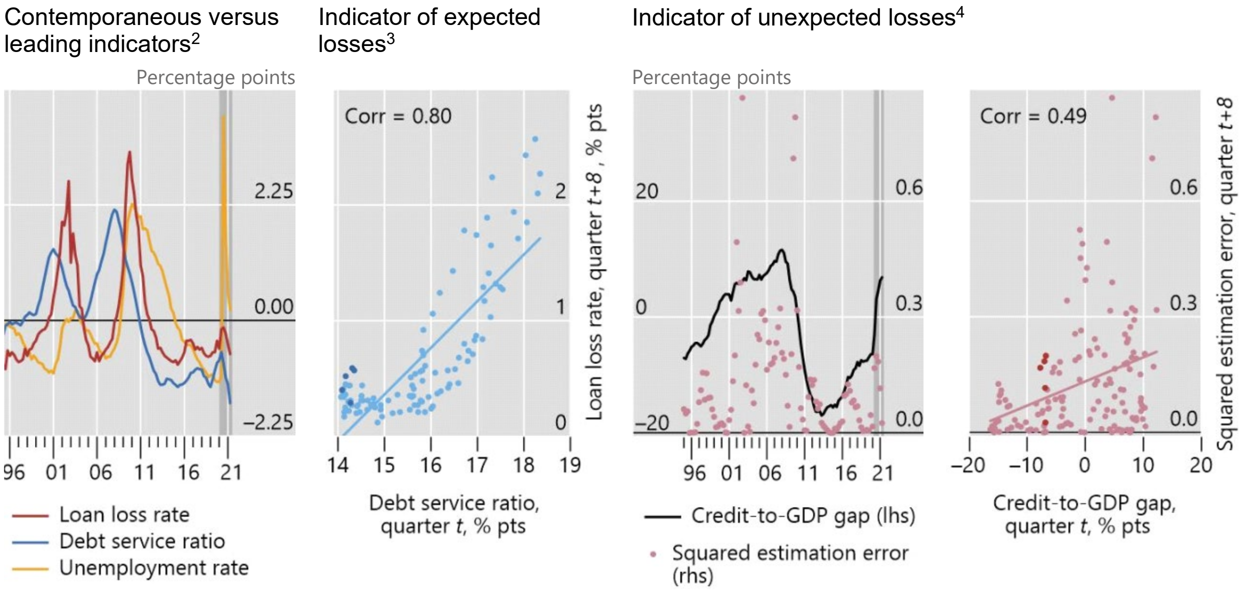

Shaded areas (first and third panels) and darker circles (second and fourth panels) indicate observations during the Covid-19 pandemic.

[1] Based on US data. Loss rates stand for net charge-off rates and refer to banks' commercial and industrial loans. The debt service ratio and the credit series underpinning the credit-to-GDP gap are for the total private nonfinancial sector. The credit-to-GDP gap in a given quarter reflects GDP over the year ending in this quarter. [2] Each series is standardised by subtracting its mean and dividing by its standard deviation. [3] From Q1 1985 to Q1 2021. [4] An estimation error is equal to the vertical distance between a dot in the second panel and the fitted line in that panel.

Sources: Board of Governors of the Federal Reserve System; Federal Reserve Bank of St Louis, FRED; BIS; authors' calculations.

Real-time loss forecasts

In principle, a number of indicators and empirical specifications could be used to forecast loss rates in *real time*, ie using information available only up to the quarter of the forecast. In Juselius and Tarashev (2020), we systematically investigate alternatives by letting the data speak. We find that elaborate specifications employing many indicators explain past losses well, but to the detriment of forecasts of future losses. We also find that the most successful forecasts of expected losses on US corporate loans rely only on the overall PNF sector's DSR. Correspondingly, the credit-to-GDP gap delivers the best forecasts of unexpected losses. This confirms our earlier arguments, prompting us to employ the two indicators in an econometric model that forecasts loss rates over multiple horizons (see appendix).

Graph 3 (first three panels) illustrates the success of expected and unexpected loss forecasts conducted in three historical quarters. For instance, using DSR and credit-to-GDP gap data up to Q1 2006, we construct forecasts of quarterly loss rates up to three years ahead (first panel, red squares and blue whiskers). These real-time forecasts of expected and unexpected losses predict the rise in loss rates during the GFC, even if not its exact timing. Likewise, forecasts formed with data up to Q1 2009 also perform well (second panel), predicting the peak of corporate loan losses at mid-2009.

The above historical forecasts exemplify a more general regularity: expected and unexpected losses tend to move in the same direction (as in Graph 1, left-hand panel). Both increase in the run-up to the GFC and decline after the crisis peaks (Graph 3, first and second panels). And both are low when the forecast is conducted at the beginning of a tranquil period, such as Q1 2013 (third panel). If forecasters perceive this regularity to be robust, they could infer changes in unexpected losses only to the extent that there are changes in the forecasted expected losses. However, we next discuss evidence that the regularity has recently broken down.

Graph 3. Real-time forecasts of expected and unexpected losses[1]

Quarterly loss rates, in per cent

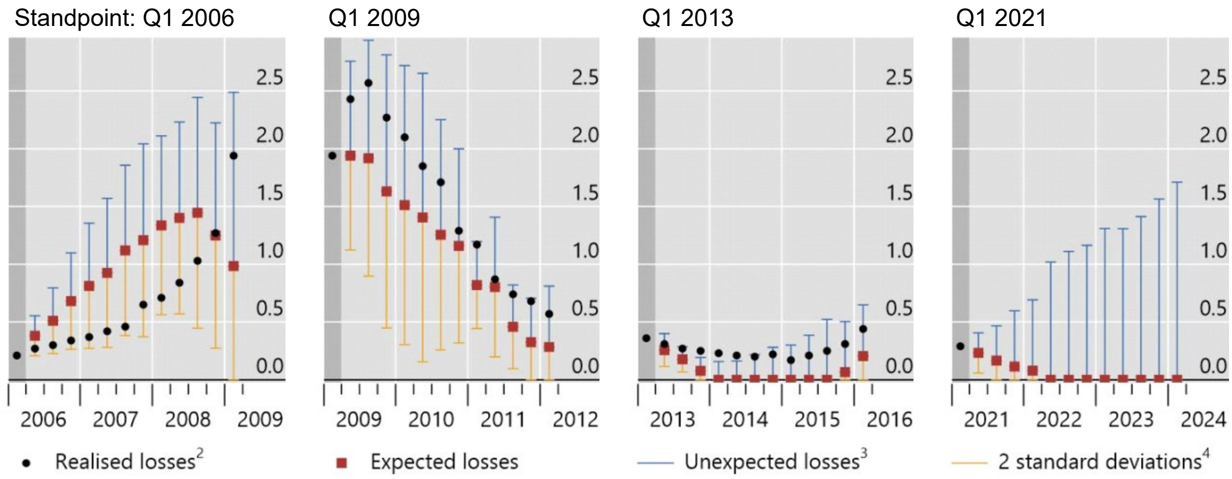

Shaded areas indicate the quarter of the forecast.

[1] The forecasts are based on a joint estimate of (i) commercial and industrial (C&I) loan loss rates as linear functions of lagged loss rates and the private nonfinancial sector's debt-service ratio; and (ii) squared loss-rate estimation errors as exponential functions of lagged loss rates and the credit-to-GDP gap, where credit is to the private nonfinancial sector. See appendix for further detail. [2] Net charge-off rates on C&I loans extended by banks in the United States. [3] Loss rates that exceed the expected level by up to two standard deviations. [4] The line's length could be shorter, to ensure that the lower bound is non-negative.

Sources: Juselius and Tarashev (2020); Board of Governors of the Federal Reserve System; authors' calculations.

Forecasting losses on portfolios held in the pandemic

The special nature of the pandemic-induced crisis and the forceful policy response to it are key drivers of the most recent observations of the two indicator variables that we use in our loss rate forecasts. The PNF sector's DSR remained low through the pandemic up to Q1 2021 (Graph 2, first panel). To an extent, this was the result of tax deferrals, cash transfers and other policy measures that supported income directly, especially that of households. Other measures – such as those to reduce interest rates and keep them low, as well as debt moratoriums and government guarantees on credit – put a lid on interest payments and backstopped cash flows, thus also keeping DSRs low. More generally, the monetary policy stance ensured the liquidity of credit markets. At the same time, enacted amidst strong precautionary motives to borrow, the policy measures incentivised the provision of credit, thus spurring the build-up of debt in the PNF sector. Since this happened when the crisis plunged the economy into a deep recession, it suggests potential excesses in risk-taking and surfaced as a spike in the credit-to-GDP gap (third panel).

Assuming that past relationships between financial imbalances and credit risk will hold in the future, the divergence of the two indicators translates into a decoupling of our forecasts of US corporate loans' expected and unexpected losses (Graph 3, fourth panel; similar to Graph 1, right-hand panel). Expected losses are forecast to decline from a low initial level (red squares), thus painting a picture of a benign baseline scenario. By contrast, unexpected losses shoot up, reaching levels comparable with but below the highest realised losses during the GFC (blue whiskers in fourth panel vs dots in first and second panels).[3] This is consistent with an increased risk of creditor retrenchment, which could lead to a wave of corporate bankruptcies down the road (Gourinchas et al (2021)). Overall, these results indicate that, while necessary and effective in containing losses so far, the policy response to the crisis has also generated considerable uncertainty: there is a wide range of possible loss rate paths up to 2024.

Our results also underscore the importance of basing loss-absorbing resources on separate forecasts of both expected and unexpected losses, with implications for banks. Taken at face value and in contrast to historical experience, expected loss forecasts on corporate loans, which drive bank provisions, point to a much more benign future than unexpected loss forecasts, which drive capital decisions. To the extent that future losses fall outside government guarantees, the high unexpected loss forecasts indicate that the adequate amount of capital for exposures to US corporate loans could be substantial.

[3] *Mojon et al (2021) predict a spike in baseline losses on the basis of projections of industry-level output.*

References

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

Banerjee, R, J Noss and J M Vidal Pastor (2021): "Liquidity to solvency: transition cancelled or postponed?", *BIS Bulletin*, no 40.

Borio, C and P Lowe (2002): "Assessing the risk of banking crises", *BIS Quarterly Review*, December, pp 43–54.

Drehmann, M, A Illes, M Juselius and M Santos (2015): "How much income is used for debt payments? A new database for debt service ratios", *BIS Quarterly Review*, September, pp 89–103.

Drehmann, M and M Juselius (2014): "Evaluating early warning indicators of banking crises: Satisfying policy requirements", *International Journal of Forecasting*, vol 30, issue 3.

Gourinchas, P-O, Ṣ Kalemli-Özcan, V Penciakova and N Sander (2020): "Covid-19 and SME failures", *NBER Working Papers*, no 27877.

——— (2021): "Covid-19 and SMEs: a 2021 'time bomb'?", *NBER Working Papers*, no 28418.

Juselius, M and N Tarashev (2020): "Forecasting expected and unexpected losses", *BIS Working Papers*, no 913.

Mojon, B, D Rees and C Schmieder (2021): "How much stress could Covid put on corporate credit? Evidence using sectoral data", *BIS Quarterly Review*, March, pp 55–70.

Appendix: Forecasting model for expected and unexpected losses

This appendix outlines the econometric model that underpins the expected and unexpected loss forecasts in Graph 3 of the Bulletin. The loss rates that the model forecasts are quarterly net charge-off rates on commercial and industrial loans extended by banks in the United States. For further details and an assessment of the model's performance relative to alternatives, see Juselius and Tarashev (2020).

The model comprises input data and an econometric specification. The two input variables are lagged loss rate observations, LR, as well as the quarterly debt service ratio, DSR, and credit-to-GDP gap, $C2Y$, which are part of the BIS credit statistics. From the standpoint of quarter t, we model the loss rate in quarter $t+h$ (with $h = 1,\ldots,12$) and the squared error in estimating this rate as follows:

$$LR_{t+h} = \alpha_h + \sum_{k=0}^{1}\left(\beta_{h,k} LR_{t-k} + \gamma_{h,k} DSR_{t-k}\right) + \vartheta_{t+h} \tag{1}$$

$$\hat{\vartheta}_{t+h}^2 = exp\left(\sigma_h + \sum_{k=0}^{1}\left(\delta_{h,k} LR_{t-k} + \theta_{h,k} C2Y_{t-k}\right)\right) + \varepsilon_{t+h}, \tag{2}$$

where the dependent variable in equation (2) equals the square of the fitted residual in equation (1).

To derive forecasts for quarter $T+h$ from the standpoint of quarter T we proceed in two steps. First, we estimate *simultaneously* the parameters of the above two equations on the basis of all available data up to quarter T. Second, on the basis of the first equation's parameter estimates and observations of LR and DSR from quarters T and T–1, we construct the *expected losses* $E_T(LR_{T+h})$. Likewise, to derive $E_T\left(\hat{\vartheta}_{T+h}^2\right)$, we use observations of LR and $C2Y$ from quarters T and T–1 in the second equation. Finally, we calculate unexpected losses as $2\sqrt{E_T\left(\hat{\vartheta}_{T+h}^2\right)}$.

Equations (1) and (2) underpin out-of-sample forecasts. In practice, these forecasts are "quasi real-time", because the data are typically revised between the quarter to which they apply and the time of the analysis.