NPLMARKETS

25 February 2022

Valuation of Non-Performing Loans: Calibration of unsecured recovery curves

Burkhard Heppe, PhD
Chief Technology Officer
NPL Markets Ltd.
burkhard.heppe@nplmarkets.com
+44 (0) 20 3984 6288

Table of Contents

Executive Summary

The statistical valuation of granular loan portfolios requires investors to determine the most suitable recovery curves. A recovery curve is a model for the amount and timing of expected collection and expense cash flows used in a discounted cash flow analysis. The calibration of recovery curves is complex for many reasons. For instance, NPL portfolios are often subject to a selection bias as only those loans are offered for sale that did not recover within a certain period after default. Investors need to determine to what extent their own historical performance data are representative for the loans on offer. Data sets made available to investors may be incomplete or unreliable resulting in missing information and censoring bias when investors do not have access to the full recovery history of similar loans. This article describes some key model decisions that investors must take when valuing granular portfolios of unsecured non-performing loans. Banks can benefit from similar model decisions and ensure the consistent use of internal and external data for the valuation of NPL and other valuation purposes like the determination of IFRS 9 loan loss provisions or LGD risk parameters. We propose a tractable competing risk Markov chain model to capture some of the observed dynamics of recovery cash flows. The model provides a simple closed formula for the investor's gross cash flow expressed as a multiple of the last 12 months collections and a closed formula for the net present value of recovery cash flows for homogeneous sub-portfolios. We offer some empirical evidence of recovery curves of European NPL. In Italy after 2020, we find that special loan servicers increased the weighted average life of the expected unsecured recovery curves by around 1.5 years. Some servicers do not adjust their recovery expectation for elapsed time after default in stark contrast to the empirical evidence and theoretical prediction.

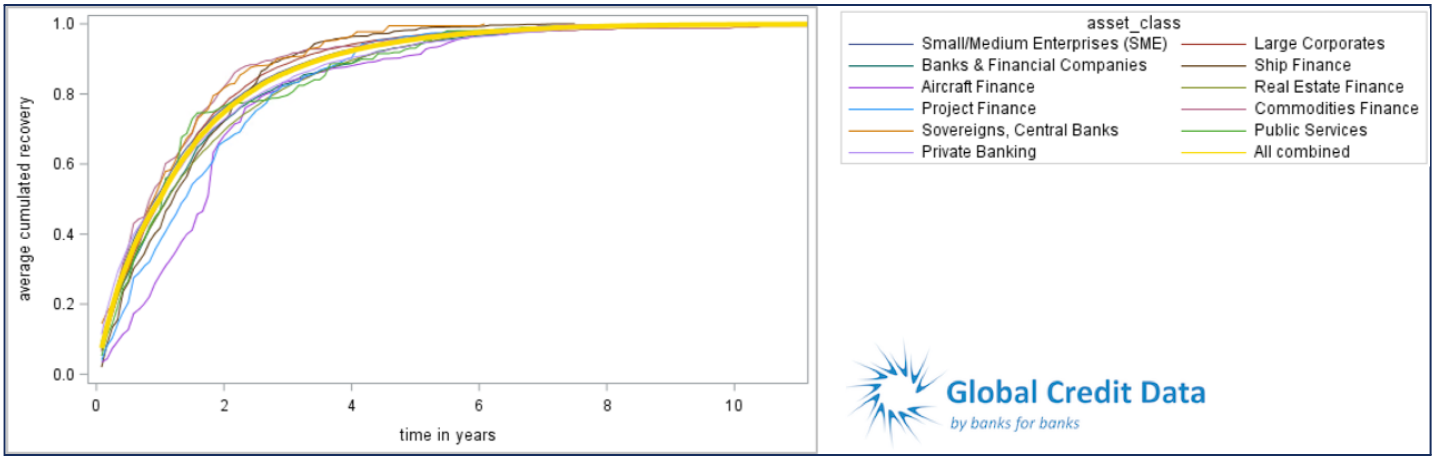

Figure 1: Cumulative recovery rates for different wholesale asset classes reported by Global Credit Data (GCD 2020 with permission). The recovery curves for different asset classes have different ultimate recovery rates. The picture shows the normalised timing of recoveries as percentage of the ultimate recovery rate.

Overview and Initial Examples

We describe some of the key steps and assumptions required to determine the value of granular unsecured non-performing loans. The model methodology discussed here excludes loans with property or non-property collateral for which different valuation methods apply based on the liquidation value of the collateral (see for example NPL Markets 2020a). Recovery curves differ by jurisdiction of the borrower and type of loan, i.e. are country and asset class specific and within one country and asset class most investors would apply different curves for different sub segments by borrower, detailed loan type, age of the loan after default, exposure size and other characteristics. For academic research on the drivers of loan recovery on wholesale loans see, for instance, Betz et. al (2021) and references therein. For drivers of consumer debt collection rates see, for instance, Kriebel and Yam (2020) who show the importance of additional information gathered by debt collectors during the workout process for accurate predictions of recoveries. They divide information used for predicting recovery rates into spatial information (e.g. local unemployment rates), external credit assessments (credit scores), customer relationship experience, and financial status information.

Figure 1 shows example recovery curves from Global Credit Data for different wholesale loan asset classes. The x-axis shows time after default in years. Figure 1 shows that more than 90% of the ultimate recovery rates are reached after around 4 years (the ultimate recovery rate is not shown). In contrast, Figure 2 shows three recovery curves for different unsecured retail portfolios in Italy in quarters after default. While the GCD example might give the impression that recoveries follow a universal curve across asset classes, this is clearly not the case for the Italian retail portfolios which differ both by the speed at which they reach their ultimate recovery rate and the ultimate recovery rate itself. Pool 1 (black line) reaches the ultimate recovery rate of 40% after 10 quarters whereas pools 2 and 3 appear not to have reached their ultimate recovery rate of around 80% after 20 quarters. Despite

these differences, the shapes of the curves in these examples are similar with the largest recoveries happening early after default and an asymptotic tail of recoveries later on. The asymptotic behaviour makes it hard to determine the ultimate recovery rate exactly and some practitioners cut-off the tail after a set number of years. The dashed lines in Figures 2 show the fit of a simple exponential curve which motivates us to consider a constant hazard competing risk model which predicts those exponential curves.

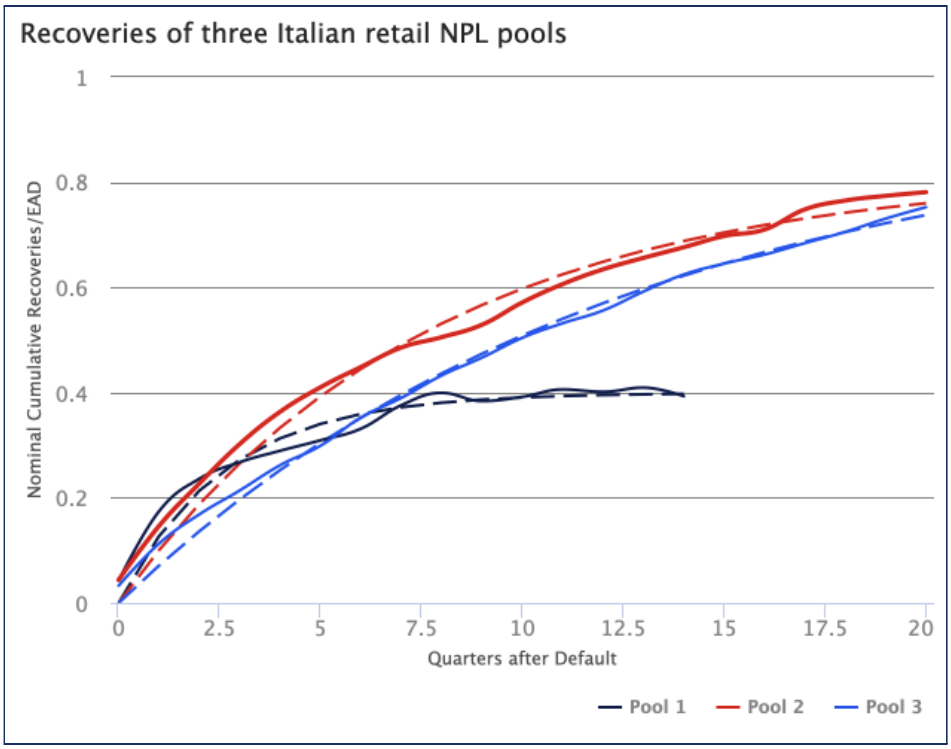

Figure 2: Cumulative recovery rates for three different Italian retail portfolios. The curves show the actual timing and ultimate recoveries. The dashed lines show the model fit of a simple competing risk model described below.

The valuation of loans based on a discounted cash flow analysis requires the projection of (undiscounted) recovery cash flows as described by the recovery curve and a suitable discount rate. How many parameters are required to describe the recovery curve? In the extreme, one could model the future recovery cash flow in each period after default until no further recoveries are expected. At the lower end of the number of parameters as Figure 2 demonstrates, at least two parameters are required, the ultimate recovery rate and a time measure describing the speed at which the ultimate recovery rate is (approximately) reached. The dashed lines in Figure 2 show the fit of a two parameter model using only the ultimate recovery rate REC and weighted average live WAL of the undiscounted cash flows. In this example, the fit is good and the resulting valuation error using the fitted model is small. We find that in practice, the valuation error introduced by a misspecification of REC and WAL much exceeds the error caused by inaccuracies in the shape of the curve. Where calibration data are scarce we recommend to focus on REC and WAL rather than on a larger number of parameters describing the curves in more detail.

The examples in Figure 1 and 2 highlight the following challenges when dealing with recovery cash flow data. First, the ultimate recovery rate for some portfolios can only be observed several years after default and data covering, say, a 5 year period may not be sufficient to observe the ultimate recovery

rates. In GCD and some of the academic literature, the timing parameter is sometimes chosen as the time to resolution, also called time to workout. A typical definition of time to workout could be the time at which 90% or 95% of ultimate recovery cash flows have been received. Given the rather regular asymptotic behaviour of the curves we prefer to use the WAL (or duration) of the cash flows as a measure describing time to recovery. The WAL is commonly used in the valuation of non-performing loans. Hence, to model recovery curves we need at least two parameters, the ultimate recovery rate (REC) and the weighted average life (WAL) in addition to information about the shape of the recovery curve.

The non-linear shape highlights another fact about unsecured recovery curves. The expected value of remaining recovery cash flows a few years after default can be very low. In Figure 2 pool 1, the remaining recoveries three years after default are practically zero as the cumulative recovery curve is completely flat at that point. On the other hand, the examples of Figures 1 and 2 also show that many loans still expect some recoveries three years after default and hence it can be expected that not all loans in the bank's workout department are either written off completely or fully covered by loan loss reserves after three years. The threshold of three years after default has recently been introduced in European bank capital regulation as the prudential backstop for unsecured loan loss provisions.

Calendar Provisions under the Prudential Backstop

The European Union in its effort to tackle the increased level of NPL after the financial crisis agreed in 2019 to establish a minimum loss coverage level for NPL depending on the years after default (Regulation (EU) 2019/630, also referred to as the prudential backstop). The aim of the prudential backstop is to ensure that banks set aside sufficient own resources for when loans become and remain non-performing and to create appropriate incentives to avoid the accumulation of NPL on the balance sheet. The new rules apply to exposures that were originated after 26 April 2019 and became effective in January 2021. Table 1 shows the minimum coverage levels for unsecured and secured loan categories. Unsecured loans that are in default for more than 3 years have to be fully covered by loan loss reserves. For loans secured by immovable collateral or other eligible collateral, a gradual increase of the annual minimum loss coverage would apply over a period of nine or seven years respectively, starting three years after the loans are classified as non-performing. The prudential backstop must be calculated at exposure level requiring banks to track their eligible collateral values (where applicable).

Minimum coverage level (in %)									
After year	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Unsecured	0	35	100						
Secured by other CRR eligible collateral	0	0	25	35	55	80	100		
Secured by immovable collateral	0	0	25	35	55	70	80	85	100

Table 1: Minimum loss coverage as per the Regulation (EU) 2019/630 Prudential Backstop.

Given the conservative design of the prudential backstop it is expected that banks will indeed find it advantageous to sell their loans after three years rather than holding them with full provisions on their balance sheet. The higher costs of holding on to fully provisions loans results in a lower break-even sale price for NPL (NPL Markets 2020b).

Accrued Interest, Legal and Servicing Costs

The valuation of NPL by investors not only requires the projection of gross cash flows but also of workout costs, legal and procedural costs and servicing fees. Banks need to deduct external workout costs when estimating LGD. Investors will often use an external loan servicer who gets paid a servicing fee for working out the loans. Banks are not required to include their internal workout costs in the calculation of LGD, but limit their cost assumptions to external expenses only. There can be many different cost and expense items and incentive fees agreed with the loan servicer can be complex. Any complete valuation model will need to be able to handle the complexity of different costs and fees. In what follows, however, we will mainly focus on gross cash flow recovery curves.

Banks express recoveries and losses as a percentage of exposure at default (EAD). For fully funded exposures the EAD will be close to the principal balance at default. The interest that has accrued and remains unpaid at the date of default is either capitalised i.e. added to principal balance or ignored in the calculation of EAD. Interest will continue to accrue after the date of default and sometimes for many years. The treatment of accrued interest differs across jurisdictions and asset classes and some countries have introduced consumer protection legislation with the aim to prevent individual borrowers from a spiralling debt burden caused by compounding penalty interest rates and collection fees. While the topic of accrued interest is complex it can have a material impact on total recoveries and the valuation of NPL.

The lender, servicer or debt collector will normally be entitled to increase the legal claim against the borrower with certain fees and expenses incurred when working out the loan. For small retail loans, these cost items can be a material percentage of total claim. For small retail loans it is not uncommon

for the maximum legal claim of principal, interest and fees to exceed the EAD by a factor of two or three. Some countries like Germany have recently introduced caps on the amounts of fees that can be charged to the borrower. For highly seasoned corporate NPL that are sold 3 or more years after default, accrued interest and fees can be a material component of legal claim (or gross book value GBV), however, the potential future increases in GBV after sale are often ignored by investors. For most retail loans and loans sold closer to the date of default, accrued interest and fees should not be ignored.

To our knowledge, the academic LGD literature is mostly silent on the impact of accrued interest, fees and expenses charged to borrowers on recoveries. Even the external workout cost that banks must take into account for their LGD calculation has received little attention. In our experience, all these effects are material and cannot be ignored when pricing NPL.

Incomplete Recovery Processes

It is a well-known fact that loans that recover later on average have lower (undiscounted) recovery rates than loans that resolve early (e.g. Betz et al 2021). If loans of a particular asset class or segment resolve, say, on average after 4 years then the observed recovery rates for completed cases (also referred to as closed or resolved cases) for loans that defaulted within the last 4 years will be biassed as there are still unresolved or open cases which are expected to have lower recoveries than the closed cases. Recovery projections based on closed cases only can overestimate actual recoveries. Bank supervisors are aware of this and require banks to include incomplete recovery processes in their data sets to estimate loss parameters (ECB 2019).

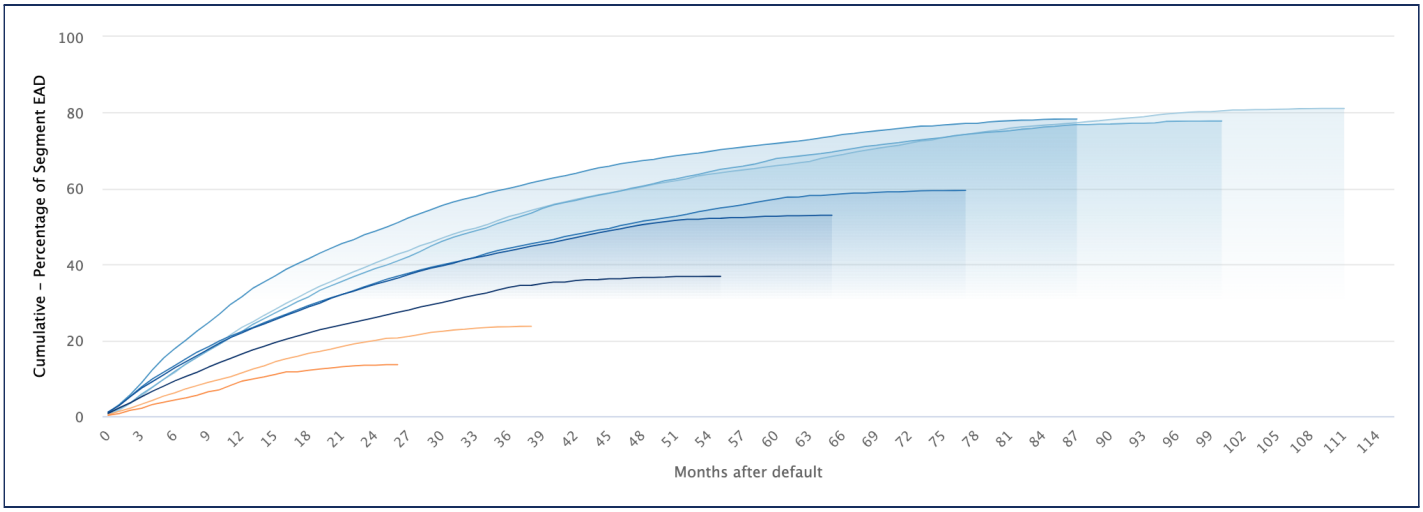

Figure 3a: Example recovery curves for unsecured consumer loans by year of default.

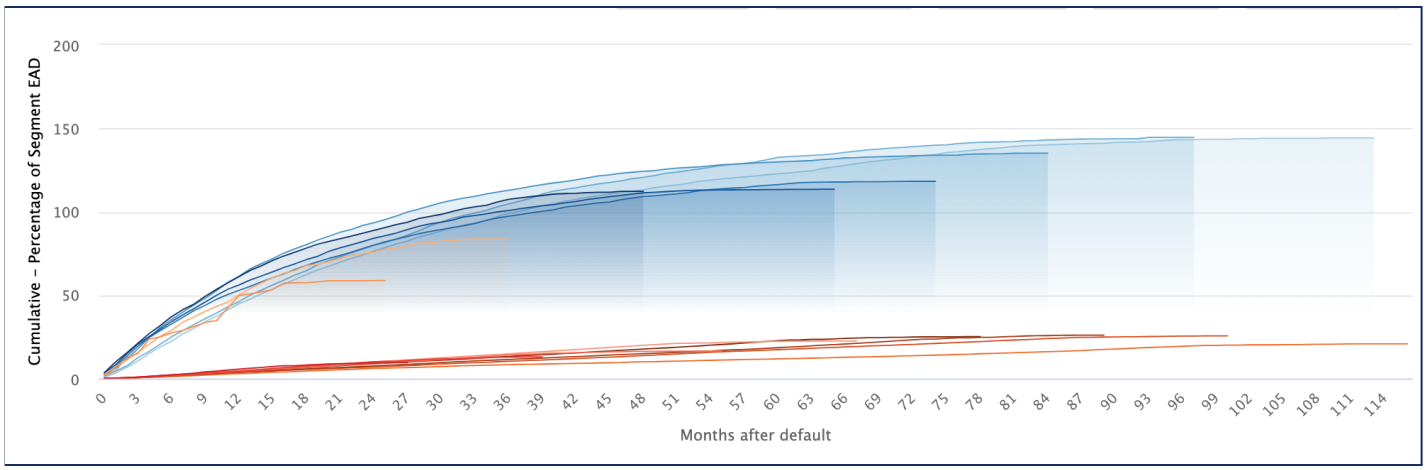

Figure 3b: The same recovery data as in Figure 3a by default year (vintage) and in addition segmenting the data into closed (blue) and open (orange) cases.

Figure 3a shows the recovery curves for unsecured consumer loans by year of default (default vintage). In this example, the more recent vintages have lower initial recoveries than those shown by older vintages at the same time after default. Here, the recent vintages were impacted by Covid whereas the older vintages had very few recoveries left by the time Covid hit in early 2020. Other factors that can contribute to pronounced vintage effects are changes in loan origination standards, changes in workout procedures, or regulatory changes.

Figure 3b shows the same example data, but now also segmented by closed and open cases. The chart shows dramatically different curves for closed and open cases. Whereas the closed cases for most vintages recovered more than EAD due to accrued interest and fees, open cases point to an ultimate recovery rate of less than 25%. This example may be extreme, but it demonstrates the relevance of the estimation bias of recovery rates if unresolved cases are excluded. For unsecured loans, it is well-known that the distribution of recoveries is bimodal at 0 and 1 i.e. most loans either recover fully or nothing at all. Cases will be closed when the full claim has been recovered. Cases that have not recovered fully will typically remain open unless there are specific reasons to write off the remaining claim like if the borrower has moved abroad to an unknown address or in the case of fraud.

Internal versus External Recovery Data

The examples of the previous sections show bank recovery data and, in the case of GCD in Figure 1, data pooled by a large number of international banks in different countries. Banks observe all recovery cash flows after default on their loans up until resolution or the date at which the loans are sold to investors. Detailed loan recovery histories can be scarce especially for so-called low default portfolios or for new loan types that were only originated recently. Changes in the definition of default in recent regulatory initiatives further complicate the creation of long consistent recovery histories. GCD was

founded by AIRB banks to pool their loan recovery data to help with benchmarking and calibrating loss given default models after Basel II (GCD 2020).

	Exposure class	AIRB	FIRB	SLSC	Number of participating institutions
LDP	LCOR	61	50	0	93
	COSP	28	17	39	67
	CGCB	25	30	0	46
	INST	31	39	0	58
HDP	CORP	61	50	0	93
	SMEC	58	47	0	89
	SMOT	70	0	0	70
	RETO	79	0	0	79
	RSMS	64	0	0	64
	MORT	87	0	0	87
	RQRR	40	0	0	40
ALL	**ALL**	**102**	**57**	**39**	**109**

Table 1 (from EBA 2021): EBA benchmark survey of regulatory risk parameters comparing low default portfolios (LDP, e.g. large corporate LCOR) and high default portfolios (HDP, e.g. SME corporates (SMEC) or Retail Other (RETO).

Of over 100 large European banks surveyed by EBA in the 2020 benchmark study, all banks applied the advanced IRB approach for high default portfolios (meaning retail pools and SME loans) whereas only 50-60% of the banks estimated their own LGD for low defaults portfolios and used the supervisory LGD parameters of the foundation IRB approach instead (Table 1). Even for high default portfolios where large banks might have a sufficient number of internal data points to calibrate LGD models, the recovery histories on the most recent defaults will be incomplete given the multi-year tail of recoveries as discussed before.

Supervisors expect banks to estimate LGD on their own recovery experience. Banks may supplement their own historical data on defaulted exposures with external data. The inclusion of external data requires careful consideration. Banks must ensure that the external data are representative of the population of the institution's actual obligors or facilities. Where internal data are scarce, proving representativeness will be difficult, but external data can still be used subject to a suitable margin of conservatism. In particular, banks should provide evidence that the model's performance does not deteriorate when including information derived from the external data, and that the parameter estimates are not biassed. To assess these issues, supervisors expect the banks to conduct quantitative and qualitative analyses specifically designed for this purpose (ECB 2019 Section 3.2).

To value illiquid NPL, i.e. NPL for which no observable market prices exist, which is generally the case except for the actively traded debt of a few large corporate insolvencies, investors require suitable recovery curves derived from historical data. Those recovery curves can be derived from internal data or derived from external data provided by loan servicers or NPL transaction platforms that help the investor to value the loans. All the complexities and biases mentioned above in the context of the estimation of bank risk parameters are equally relevant to investors. Investors must determine

whether the available internal or external calibration data are representative for the portfolio under consideration for valuation. The potential bias in the recovery rates of closed cases can be important if the seller provides recovery curves based on close cases only. Specifically, as most portfolios are sold to investors only after the bank has attempted to work out the loans internally and has been unsuccessful in doing so, the age effect of time after default and the negative selection bias must be considered carefully. In the following, we describe a statistical methodology based on survival analysis which is designed to capture the dynamics of recovery cash flows while allowing for incomplete recovery processes and unbiased estimates based on a combination of internal and external data. The proposed methodology is applicable to the valuation of NPL by investors and the estimation of regulatory risk parameters by banks.

A Competing Risk Model of Recovery

We propose a simple model for recovery cash flows after default applicable to the workout of granular portfolios of unsecured retail or SME loans. We model the undiscounted gross cash flows over time and initially ignore legal expenses and servicing fees which can be deducted from gross cash flow in a second step which is not detailed here. A valuation requires discounting the predicted net cash flows with an appropriate discount rate. To derive a fair market valuation the correct discount rate would be a risk neutral discount rate derived from traded instruments or the internal rate of return expected by the investor. For the calculation of Stage 3 loan loss provisions under IFRS 9, the appropriate discount rate would be the effective interest rate of the loan. For the determination of workout LGD, banks should use a discount rate that reflects the costs of holding the defaulted assets during the workout period including an appropriate risk premium and the discount rate is sometimes prescribed by bank supervisors.

Our model is based on a Markov chain for competing transitions which is part of statistical survival analysis. Markov chain models are well-known in credit risk for the modelling of rating migrations of corporate loans or the roll-rate analysis of retail loans. For performing loans, default and prepayment are often considered competing risks. For the workout of NPL, Markov chain competing risk models are less common. A recent exception is Chamboko Bravo (2020) who use a multi-state Markov chain to model the transitions to and from various states for US residential mortgages. The authors model four transitions from default: short sale, repurchase, prepayment and foreclosure (Figure 4). Early examples of using survival analysis for LGD include Dermine de Cavalho (2006) who considered losses on SME loans from a Portuguese Bank and Witzany et al (2012) who considered retail loans of a Czech bank. Fenech et al (2016) use survival analysis to model SME recovery outcomes after default. Chomboko Bravo (2020) model transitions applicable to the loan in its entirety, our approach has been inspired by Witzany (2012) who modelled the recovery cash flows as a counting process to model partial recoveries.

The recovery rate expressed as a percentage of exposure at default (EAD) is often modelled as a fraction between 0 and 1. However, as mentioned in practice the ultimate recovery rate often falls outside the unit interval as expenses and additional advances for some loans exceed collections and in other cases the lender can collect all of EAD plus accrued interest resulting in a recovery rate of more than one. For high-interest unsecured retail loans like credit cards, the highest recovery rates can range between 1.5 and 4 (see for instance Figure 3b).

Using survival analysis enables us to capture censored information which can arise from incomplete workout processes (right censoring) or from expenses and further advances after default (left truncation). We treat repayment and loss allocation (or write-off) as competing events. Whereas default and prepayment-in-full are binary events applicable to a loan in its entirety, a model for recovery and loss must be able to deal with partial outcomes.

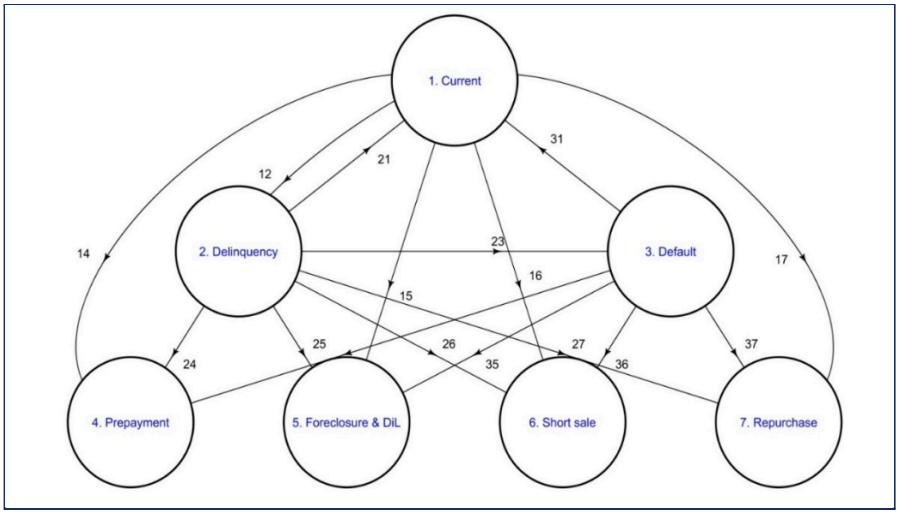

Figure 4: Example transitions for mortgage default and loss (from Chomboko and Bravo 2020). The NPL Markets valuation tool for secured and unsecured NPL also uses four transitions from default: Foreclosure/Judicial, Repayment in full, Discounted payoff (Partial repayment), and Reperforming (NPLM 2020a).

In our set up the key measure to be predicted are the ultimately realised recovery rate REC expressed as a fraction of EAD. In addition, we require a measure for time in workout for which we take the cash flow weighted average life, WAL. REC and WAL can be observed only on defaulted and resolved loan[1] while banks must estimate the expected recovery rate on performing loans based on available information. Not all banks have explicit models for WAL. For NPL to be sold, both the seller and the buyer must estimate REC and WAL based on all available information that for NPL includes post-default information such as the nature of default (e.g. past due or bankruptcy) or any partial recovery and expense cash flows that have occurred up until the valuation date.

We observe the positive or negative recovery cash flows generated during the workout of the loan. The workout process may be internal or outsourced to an external party where a debt collection company is paid a fee for collecting the payment on behalf of the loan owner. The process may also combine an ordinary collection and sale of the receivable to a third party. In any case the work-out process

[1] We use the term closed and resolved synonymously as well as open and unresolved.

involves significant costs as negative recovery cash flows in addition to further advances extended for reasons other than cost.

Data preparation and counting processes

Our setup follows Witzany (2012), i.e. we model recovery cash flows as a counting process suitable for survival analysis. Time is expressed in months after default and cash flows are expressed as a positive natural number equal to the rounded percentage point of EAD. For example, this means each loan starts with 100 observations that are all right censored (i.e. the status variable "Event" is set to 0 - censored) as neither a recovery nor a loss event have yet been observed. A positive recovery cash flow after 10 month equal to 8% of EAD will then be recorded by setting the Event variable of 8 of the 100 observations to 2 - repaid. Assume further that the resolution date is 22 months after default with no further recoveries, then Witzany proposes to keep the unrecovered 92 observation units as censored with Event 0-censored. In contrast here, we propose to treat the allocation of loss as a competing event and set the Event variable to 1-loss at the resolution time. If the workout continued beyond the last available date of observation (an unresolved default), then we would instead keep the Event variable equal to 0-censored. The counting process framework requires that we cannot model cumulative negative recoveries nor recoveries exceeding 100% of EAD plus advances so our modelled recovery rate will always fall in the unit interval. Note that our approach estimates a weighted average recovery rate for the portfolio at hand.

Competing risk analysis

We propose to treat recovery and loss allocation as competing risks that are both of standalone interest in a Markov chain model with the three states: workout, loss, and repayment. Whereas the repayment state is directly observed from cash flows, the loss event requires a subjective decision of the lender or servicer when to charge-off the loan and/or to close the account. A multi-state model describes how an observation unit moves between a series of discrete states in continuous time. The first step in a multi-state model analysis is to set up the transition matrix. The transition matrix specifies which direct transitions are possible. Here we have the transient workout state and the two absorbing states repayment and loss.

from/to	Workout	Repayment	Loss
Workout	$\lambda_{11} = -\lambda_{loss} - \lambda_{rec}$	$\lambda_{12} = \lambda_{rec}$	$\lambda_{13} = \lambda_{loss}$
Repayment	0	0	0
Loss	0	0	0

Table 2: The generator transition matrix for a competing risk model with recovery and loss as two absorbing states. If there are no further advances or cost payments, then all loan observation units start in state Workout at t=0 where t is time after default.

The movement on the discrete state space is governed by transition intensities $\lambda_{rs}(t, X(t))$ r, s = 1 (workout), 2 (repayment), 3 (loss). These may depend on time t, or, more generally, also on a set of individual-level or time-dependent explanatory variables X(t). The intensity represents the instantaneous risk of moving from state r to state s≠r. The transition intensities or hazard rates λ_{rs} form a 3*x*3 matrix whose rows sum to zero (Table 2). In a time-homogeneous Markov model, in which the λ_{rs} are independent of t, the sojourn time in each state r is exponentially-distributed with mean $1/\lambda_{rr}$. The mean sojourn time is just the weighted average life of the recovery cash flows $WAL=1/(\lambda_{rec} + \lambda_{loss})$. The probability that a unit in state r moves next to state s is $-\lambda_{rs}/\lambda_{rr}$. Each transition specific hazard λ_{loss} and λ_{rec} can be estimated from data using either parametric or non- or semi-parametric models such as the constant hazard model or the Cox regression, respectively.

Under the assumption of constant hazard rates the model is fully tractable. The probability that a loan at time after default t is still in workout is

$$P_{workout}(t) = \exp(-(\lambda_{rec} + \lambda_{loss})t).$$

The cumulative recovery rate REC(t) expressed with the lifetime parameters REC and WAL is then equal to the probability that a unit of exposure has been recovered at time after default t with

$$P_{rec}(t) = \frac{\lambda rec}{\lambda rec + \lambda loss}[1 - \exp(-(\lambda_{rec} + \lambda_{loss})t)].$$

The probability of loss or charge-off follows as

$$P_{loss}(t) = 1 - \exp(-(\lambda_{rec} + \lambda_{loss})t) - \frac{\lambda rec}{\lambda rec + \lambda loss}[1 - \exp(-(\lambda_{rec} + \lambda_{loss})t)].$$

NPL valuation with the constant hazard model

We can now express the constant hazards in terms of our lifetime parameters REC and WAL. The cumulative recovery rate at time t after default follows a simple exponential curve

$$REC(t) = REC\,[1 - \exp(-t/WAL)].$$

The marginal recovery rate MREC(t) in period [t,t+δt] is given by the derivative of REC(t)

$$MREC(t) = REC/WAL \exp(-t/WAL)\,\delta t.$$

If the portfolio is sold at the cutoff date T after default then the seller retains all cash flows until T and the buyer receives all cash flows for t>T. The cumulative remaining recovery rate $RREC_T(t)$ for the buyer is then the difference of the cumulative recovery rate REC(t) less the cumulative recovery at T

$$RREC_T(t) = REC(t) - REC(T) = REC\,[\exp(-T/WAL) - \exp(-t/WAL)].$$

In practice, investors like a simple heuristic or rule of thumb for the expected remaining undiscounted cash flow and the net present value of the cash flows. In the simple model discussed here, both undiscounted cash flow and net present value can be calculated exactly. The total cumulative undiscounted cash flow for the investor $RREC_T(\infty) = REC \exp(-T/WAL)$ can be expressed as a multiple of the collections in the last 12 months (with t and T expressed in years)

$$RREC_T(\infty)/(REC(T) - REC(T-1)) = 1/(\exp(1/WAL)-1)$$

Interestingly, the result is independent of the cutoff date T and the lifetime recovery rate REC and only depends on the WAL. As a rule of thumb, the multiple of all remaining cash flows over last year's collections is approximately WAL less 0.5 years. Specifically, with WAL=2.0, 4.0, 6.0 years, the exact multiples are 1.54, 3.52, and 5.51, respectively.

In continuous time, if the investor discounts the marginal cash flows at t>T with the internal rate of return IRR i.e. with a discount factor of exp(-IRR(t-T)), then the net present value NPV of the investor cash flows at the cutoff date T is given as

$$NPV(T) = REC \exp(-T/WAL) / (1 + WAL\ IRR).$$

The NPV expressed as a multiple of the last 12 months cash flow then depends on the WAL and the IRR as follows

$$NPV(T)/(REC(T) - REC(T-1)) = 1/[(\exp(1/WAL)-1)(1 + WAL\ IRR)]. \quad (1)$$

With WAL=2.0, 4.0, 6.0 years and IRR=10% the multiples are 1.28, 2.51 and 3.45, respectively. A market rule of thumb we encountered for unsecured retail loans is that the NPL purchase price equals 3 to 4 times last year's cash flow. We note that such a rule must be dependent on WAL and IRR as in formula (1) and cannot apply universally.

Parameter Estimation

The parameters of the transition matrix can be estimated using standard statistical procedures including time dependent covariates. To capture partial recoveries and losses we need to transform the data to the counting process format discussed above including the indicators for censoring. If we assume constant hazard rates, then the transition matrix can be estimated using the duration method. In credit risk the duration method is well known for the estimation of the generator matrix for rating transitions (Inamura 2006). The duration method estimates the transition hazard to recovery $\lambda_{rec,\ s}$ for each homogeneous sub-segment *s* as the number of transitions $N_{rec,\ s}$ from workout to repayment divided by the total years spent in workout $T_{workout,\ s}$. Note that $T_{workout,\ s}$ does not only include the time elapsed until a recovery cash flow was received, but also includes censored observations that are still continuing in workout and the observed times spent until loss/charge off.

$\lambda_{rec,\,s} = N_{rec,\,s} / T_{workout,\,s}$

Hence, the estimation of the constant hazards and the lifetime parameters REC and WAL does not require the use of statistical software and is easy to implement in Excel. Moreover, under the assumption that hazard rates are time independent the estimation does not require complete workout histories, but can be performed on an observation time window $[t_{start}, t_{end}]$ where loan workouts are not observed from the date of default $t_{start} > 0$ and the large percentage of the portfolio might still be in workout at the end of the observation period at t_{end}. As such the estimation procedure is well suited for investors who may only have partial collection histories available to them.

We cannot expect that the constant hazard assumption is accurate for all loan workout stages, scenarios and portfolio segments. In particular, banks and debt collectors are aware that the first weeks and months after default are critical for entering into a voluntary repayment agreement with the borrower. Immediately after default, we do not expect the real recovery hazard to be constant but to increase from zero to some maximum rate observed some months after default. For example, Fenech et al (2016) investigate the recovery hazard for SME loans and find a hump-shaped hazard function peaking 23 months after default. Hence, the assumption of constant transition intensities to recovery and loss is an approximation that needs to be validated for each portfolio at hand. Nevertheless, as shown in Figure 2 the constant hazard assumption can be quite accurate. The dashed line in Figure 2 shows the fit of a simple exponential curve predicted by the constant transition intensities. The following section will provide more empirical examples of recovery curves.

Examples of Empirical Recovery Curves

Figure 4 shows an example of the unsecured recovery curves for an Italian NPL securitisation. The original loan data tape is valued at the securitisation cutoff date in July 2021. The special loan servicer develops the recovery assumptions for the securitised portfolio (original business plan BP). Each month after the cutoff date, the servicer reports the actual collection cash flows. Banks and investors will compare the actual and business plan recoveries with their own modelled cash flow projections. Figure 4 also shows a baseline scenario derived from the competing risk model described above which was calibrated to general Italian market recovery data. The actual recoveries observed in the first year of the transaction until September 2021 fall between the lower values of the original BP and the higher baseline projections. In the long run, however, the BP for unsecured retail loans exceeds the baseline curve by around 2%. In this example, the projected gross collections of the original BP and baseline model scenario are fairly close and the transaction is currently performing broadly in line with expectation.

Figure 5 shows another Italian NPL securitisation where the original BP was developed in early 2019, a year before Covid-19 hit the Italian economy. In this example, the actual recoveries of the securitised pool underperformed the original business plan by a wide margin even before Covid started.

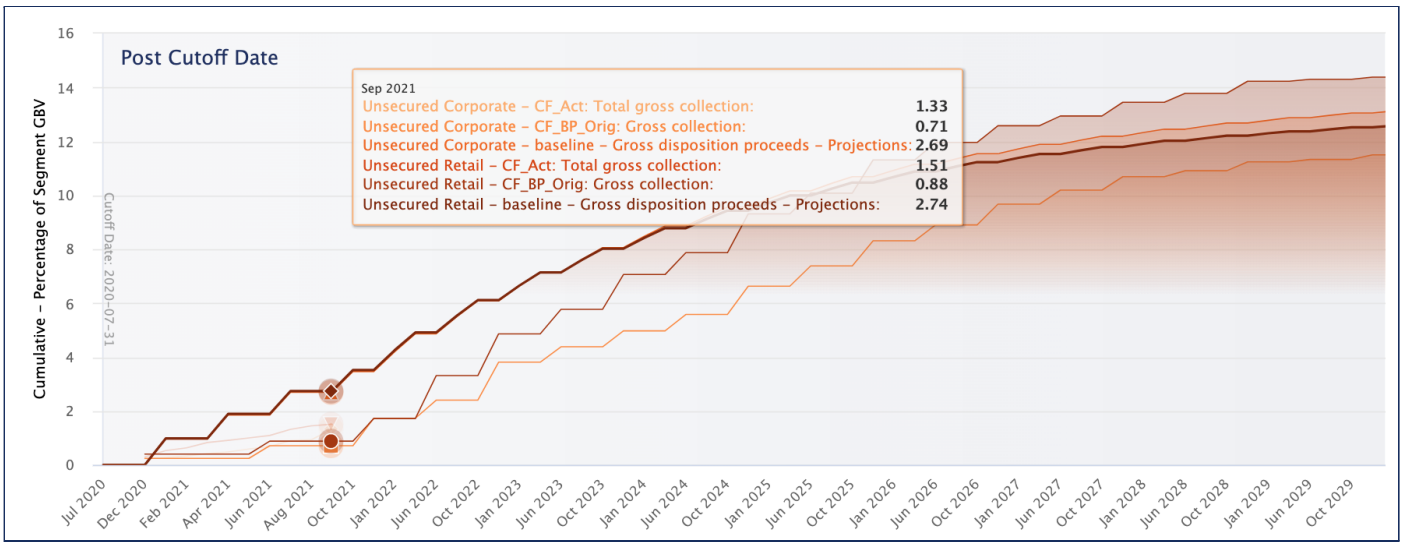

Figure 4: Unsecured recovery curves for an Italian non-performing loan securitisation. The graph shows recovery curves to corporate borrowers and private individuals, respectively. In each case, the graph shows the actual gross collection cash flows to date (CF_ACT until September 2021), the original business plan from the special servicer (CF_BP_Orig) and the baseline scenario of the competing risk model (baseline) calibrated against general Italian recovery data.

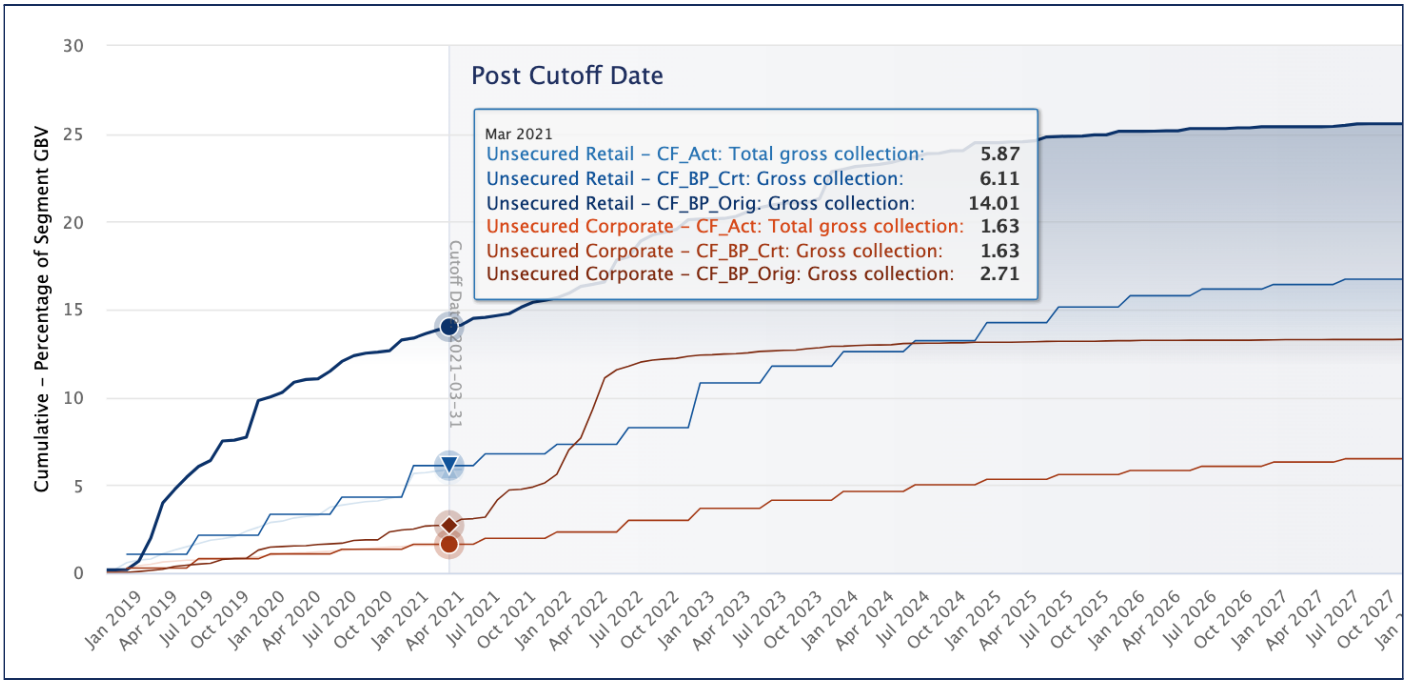

Figure 5: Unsecured recovery curves for an Italian non-performing loan securitisation that closed pre-Covid in 2019. The graph shows recovery curves to corporate borrowers and private individuals, respectively. In each case, the graph compares the original business plan from the special servicer (CF_BP_Orig) and the current business plan (CF_BP_Crt) which was revised in early 2021. The current BP was revised to match the actual recoveries until the reporting cutoff date in April 2021.

The special servicer updated the original business plan in early 2021 and materially revised the projections down. The example shows that even in a relatively data-rich market like Italy, valuing NPL is a difficult exercise and valuation errors can be large. In this example, as of mid 2021, the servicer's original BP predicted 14% gross recoveries of GBV from unsecured loans to individuals whereas the actual recoveries were only 6% of GBV. The consequences of overstating recoveries in the original BP are severe and investors in the senior tranche of this securitisation are not expected to have their principal returned in full. We argue that an accurate prediction of REC and WAL will keep valuation errors small, even if the exact shape of the curve differs from the simple parametric shape of the constant hazard competing risk model.

Table 3 shows the remaining undiscounted gross recovery rates estimated from servicer business plans done before Covid (BP Post 2020=N) and revised in late 2020 or early 2021 post Covid (BP Post 2020=Y) for unsecured Italian NPL. The table highlights that while the remaining recovery rate changed relatively little (for secured loans which are not discussed here we even observed some increases in recovery rates), the remaining weighted average life was materially extended in the revised business plans by around 1.5 years.

Legal Type	Secured Status	BP Post 2020	No. Borrowers	WAL	REC
Corporate	Unsecured	N	16183	3.7	13.2
Corporate	Unsecured	Y	34986	5.2	11.6
Private Individual	Unsecured	N	7592	3.5	19.0
Private Individual	Unsecured	Y	16016	5.1	17.1

Table 3: Unsecured recovery curve parameters pre-2020 and post-2020 calibrated from the original and updated business plans of Italian NPL securitisations. The increase in the remaining weighted average live WAL by around 1.5 years after Covid-19 has a larger impact on investor returns than the change in the remaining recovery rate REC.

Table 4 shows the variability of remaining REC and WAL for unsecured loans for different Italian loan servicers for borrowers that are not subject to an insolvency procedure. Some variability in REC and WAL can be expected due to regional heterogeneity, exposure size or default vintage. For corporate borrowers, the average remaining recovery rate for different servicers varies between 5.4% and 19% whereas the WAL varies from 2.7 years to 5.6 years.

Measure	Legal Type	Servicer A	Servicer B	Servicer C	Servicer D	Servicer E	Servicer F	Servicer G	Servicer H
REC	Corporate	17.8	11.1	11.7	18.2	19.0	16.7	5.4	14.0
REC	Private Individual	22.4	13.5	13.5	23.6	13.1	11.2	12.9	19.9
WAL	Corporate	5.5	5.3	2.7	5.4	4.4	4.2	5.6	5.6
WAL	Private Individual	5.2	5.6	2.7	5.6	3.9	6.0	5.1	5.2

Table 4: Variation in the remaining unsecured recovery curve parameters calibrated from Italian servicer business plans.

Table 5 shows the remaining undiscounted gross recovery rate by year of default (Default Vintage) for unsecured loans to corporate borrowers which are not subject to insolvency proceedings for different Italian NPL servicer based on their business plans projections post 2020. The constant hazard competing risk model predicts that marginal recoveries are continuously declining after default. As a result, the recovery expectation should decline the longer the loan stays in default i.e. earlier default vintages should have lower remaining expected recoveries. While some of the servicers show the expected decline in recovery rates the older the vintage (e.g. Servicers A, D, H) other servicers do not (e.g. Servicers B, C, E). The number of observations is large enough to rule out small sample size effects i.e. the fact that some recovery rates do not increase for younger vintages must be a feature of the statistical model used by those servicers. In our view, the vintage behaviour shown for Servicers B, C, E is unexpected and hard to justify on theoretical and empirical grounds.

Legal Type	Default Vintage	NoBorrowers	Servicer A	Servicer B	Servicer C	Servicer D	Servicer E	Servicer F	Servicer G	Servicer H
Corporate	2007 and before	5054	13.0	9.1	12.7	5.3	21.7	9.4	5.2	8.8
Corporate	2008-2009	2192	14.4	8.6	16.7	9.1	18.3	12.2	5.1	8.7
Corporate	2010-2011	2924	14.3	7.4	16.1	11.6	16.8	9.5	4.9	11.9
Corporate	2012-2013	4397	14.6	9.4	14.8	13.6	18.2	9.8	5.3	13.9
Corporate	2014-2015	4717	15.8	9.6	10.5	16.9	20.8	14.4	7.2	16.2
Corporate	2016-2017	5688	17.6	16.4	7.3	19.4	21.2	20.9	4.0	17.2
Corporate	2018-2019	5613	22.1	9.5		20.4	10.4	19.2		21.1

Table 5: Variation of the remaining unsecured recovery rate REC(t) by default vintage and different Italian servicers.

Conclusions

We discuss the problem of estimating recovery cash flows after default and propose a simple competing risk model with constant repayment and loss hazards. We discuss a few examples of empirical recovery curves observed for unsecured European NPL and explore the two important parameters describing the recoveries, the ultimate recovery rate REC and weighted average life WAL. We discuss the importance of including unresolved workout data (open cases) in the estimation for REC and WAL. We demonstrate that for unsecured loans estimating REC and WAL from resolved defaults (closed cases) only can lead to a large bias due to the bimodal distribution of recoveries. From Italian NPL securitisation we observe that the loans servicers increased the remaining WAL of their business plans by around 1.5 years between 2019 and 2021, a dramatic change with a material impact on investor returns. In comparison, we do not see a similarly large impact on the ultimate recovery rate after the onset of Covid-19. While we do not report the detailed calibration results for REC and WAL in this article, we find statistically significant differences by exposure size, borrower legal type, bankruptcy status, region and default vintage in line with academic research. The constant hazard competing risk model implies that marginal remaining recoveries continuously decline with time after default. Curiously, we do not consistently find the predicted decreases in the expected remaining recovery rates for older default vintages of unsecured loans in the business plans of some Italian NPL securitisations.

References

Betz J, Kellner R, Rösch D. 2021. Time matters: How default resolution times impact final loss rates. J R Stat Soc Series C. 2021;70:619–644. https://doi.org/10.1111/rssc.12474

Chamboko, R.; Bravo, J.M. 2020. A Multi-State Approach to Modelling Intermediate Events and Multiple Mortgage Loan Outcomes. *Risks* 2020, *8*, 64. https://doi.org/10.3390/risks8020064

Dermine, J and C. Neto de Carvalho 2006. Bank loan losses-given-default: A case study. *Journal of Banking and Finance 30*(4), 1219–1243.

ECB 2019. ECB guide to internal models.
https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.guidetointernalmodels_consolidated_201910~97fd49fb08.en.pdf

Fenech, J P et al. 2016. "Modelling the recovery outcomes for defaulted loans: A survival analysis approach." *Economics Letters* 145 (2016): 79-82.

Global Credit Data 2020. Unresolved Defaults LGD Study November 2020 and GCD European conference 2021. LGD for unresolved default. https://globalcreditdata.org/gcd_library/unresolved-defaults-lgd-study-2020/

Inamura Y 2006. Estimating Continuous Time Transition Matrices from Discretely Observed Data. *Bank of Japan Working Paper Series*, 2006. URL https://www.boj.or.jp/en/research/wps_rev/wps_2006/ data/wp06e07.pdf.

Kriebel J, Yam K 2020. Forecasting recoveries in debt collection: Debt collectors and information production. Eur Financ Manag. 2020;26:537–559. https://doi.org/10.1111/eufm.12242

NPL Markets 2020a. How to value bank loans in a crisis.
https://nplmarkets.com/en/research/article/how-to-value-bank-loans-in-a-crisis

NPL Markets 2020b. Impact of the Prudential Backstop on Bank's Balance Sheet
https://nplmarkets.com/en/research/article/impact-of-the-prudential-backstop-on-bank's-balance-sheet

Witzany, J, Rychnovský, M, & Charamza, P 2012. Survival Analysis in LGD Modeling. European Financial and Accounting Journal, 2012(1), 6-27.

Disclaimer

This paper contains confidential information about NPL Markets, current at the date hereof. This presentation is not intended to provide the sole basis for evaluating NPL Markets and should not be considered as a recommendation with respect to it or any other matter.

This document and the information contained herein are not an offer of securities for sale in the United States and are not for publication or distribution to persons in the United States (within the meaning of Regulation S under the United States Securities Act of 1933, as amended).

This presentation and the information contained herein does not constitute or form part of any (i) offer or invitation or inducement to sell or issue, or any solicitation of any offer to purchase or subscribe for, any securities or (ii) invitation or inducement to engage in investment activity within the meaning of Section 21 of the United Kingdom Financial Services and Markets Act 2000, as amended, nor shall any part of this presentation nor the fact of its distribution form part of or be relied on in connection with any contract or investment decision relating thereto.