30 April 2026

Burkhard Heppe, PhD
Chief Technology Officer
Accuria Ltd.
burkhard.heppe@
accuria.com

Valuing Non-Performing Consumer Claims with External Reference Data and Curves

Table of Contents

Overview

The European market for non-performing consumer loans (NPLs) is a well-defined institutional asset class, but pricing remains an exercise in inference under uncertainty. Large internationally-integrated debt collection platforms like Intrum in Sweden, Lowell in the UK or EOS in Germany have data warehouses with case numbers exceeding 100 million. In Germany alone 28 million claims are processed every year. However, there are several thousand debt collection agencies and institutional investors in Europe many of them small or recently founded that do not have the large data warehouses of the incumbent platforms. The EU Commission's NPL Advisory Panel's 2024 survey found most market participants believe the secondary market is functioning well but that the absence of a fully-fledged central hub still leaves smaller institutions disadvantaged (EC 2025).

This report sets out a practical, statistically grounded framework for using **external credit-collection data and reference recovery curves** to value consumer NPLs in situations where internal data is sparse, unrepresentative, or absent, most importantly for new sellers, new loan products, or recent vintages with no observed workout, and special legal-status segments such as UK Individual Voluntary Arrangements (IVAs).

We integrate the Global Credit Data representativeness guidelines (GCD 2025) with academic research on censored recovery-curve estimation, and the Bayesian-credibility blending approach for combining seller-provided collection data with external curves.

The central decision rule is simple and an important message to sellers of NPL: **the more relevant internal data the seller can supply, the more weight should fall on observed cashflows; the less relevant data is available, the more weight sits on representative external reference data with an increased margin of conservativeness**. Bayesian credibility theory makes that intuition operational and auditable. This report complements our recent analysis of the use of external data for large wholesale NPL valuation using external data (Accuria 2026).

© Accuria Ltd. 2026

1. Introduction and Market Context

1.1 Why external data matters for consumer NPL pricing

The valuation of a consumer NPL portfolio reduces to estimating the timing and magnitude of cash flows from a heterogeneous set of small-balance, statistically determined claims. Two parameters dominate: the ultimate remaining recovery rate (REC) and the weighted-average life of recoveries (WAL), which together feed a discounted-cashflow valuation. Both must be estimated from data that is, by definition, partial: most claims in any current transaction portfolio have not yet been resolved, and seller-supplied histories rarely span enough cycles to be self-sufficient. External data fills four distinct gaps:

1. **Anchoring estimates when seller data is short or absent** — under the European NPL Directive (EU) 2021/2167, bank sellers must provide loan-by-loan information using EBA templates, but historical collections are required only for a 36-month look-back via Template 5 (EBA 2022). For typical unsecured consumer claims with workouts running 7–10 years or more, this means data in the seller tape will often be left truncated (incomplete history since default) and right-censored (open cases dominate for which the ultimate recovery rate is not known).
2. **Validating seller-claimed performance** — Scope's monitoring of European NPL securitisations shows realised collections frequently diverge from servicer business plans; on a sample of Italian, Spanish and Portuguese mixed-NPL transactions, average realised unsecured recovery was 9.4% as of June 2024 versus a B-case expectation of 7.9% (Scope 2025). Many NPL securitisations in Italy and Greece are underperforming their original business plan by more than 10% (Accuria 2024).
3. **Calibrating downturn or stress scenarios** —The 2025 update to the ECB Guide to Internal Models has elevated the LGD reference value to a near-binding backstop, asking institutions to thoroughly analyse cases where it materially exceeds final estimates.
4. **Pricing segments where the investor has no internal experience** — for example a UK debt purchaser entering the IVA segment for the first time.

1.2 Specific use cases for external data

New sellers without a track record. When a credit institution sells its first NPL portfolio, the buyer cannot observe how that institution's underwriting, servicing transfer and prior collections compare to peers. Benchmarks from other bank originators or servicers and pooled reference curves become the primary anchor.

New loan products. UK BNPL grew from £0.06 billion in 2017 to over £13 billion in 2024 — 20% of UK consumers (10.9 million adults) used it in the 12 months to May 2024. FCA regulation begins on 15 July 2026, meaning the asset class will only generate properly regulated default and recovery histories from the second half of the decade. Until then, BNPL NPL pricing depends on analogue data from short-term unsecured consumer credit and point-of-sale finance.

New vintages. Cohorts that defaulted within the last 12–24 months will have most cases unresolved. Benchmark curves become critical to extrapolate the long tail.
New legal statuses. UK insolvency volumes have shifted dramatically: in the 12 months to October 2025, 57% of individual insolvencies were IVAs, 37% DROs and only 6% bankruptcies, against a 2015 split of 50/30/20. For an investor pricing a DRO-heavy tail, historical data drawn from the pre-2024 regime is materially unrepresentative.

1.3 Market structure: UK and EU
The European NPL servicing sector has consolidated around a handful of cross-border platforms — Intrum, Lowell, Cabot/Encore, Arrow, B2 Impact, Hoist, Kruk, doValue, iQera and EOS — most of which had been active forward-flow buyers of UK and EU consumer claims. Cyclical pressure on funding costs has reshaped the buyer universe with Intrum debt restructuring leading the headlines in late 2024. This thinning has tightened underwriting standards and made representativeness analysis a tangibly commercial exercise.
On the regulatory side, the NPL Directive 2021/2167 had been transposed by 26 of 27 Member States as of December 2025, with infringement proceedings still pending against six (including Spain, the Netherlands and Portugal). The Directive requires non-bank credit purchasers of EU-bank-originated consumer NPLs to appoint authorised credit servicers and obliges sellers to use EBA templates for pre-sale disclosure. UK firms are out of scope for the EBA templates, but UK sellers may adopt them voluntarily as a market standard for cross-border investors.

2. The Representativeness Framework

Representativeness is the foundational concept that determines how much weight an external dataset can carry in a given valuation. The GCD Representativeness Guidelines v1.0 provide the most recent and detailed industry standard, designed primarily for IRB credit-risk modelling but directly transferable to NPL pricing (GCD 2025).

2.1 Definition and the two faces of representativeness
GCD distinguishes representativeness for **risk differentiation** (i.e. ranking — does the model discriminate between high-risk and low-risk cases) from representativeness for **risk quantification** (i.e. calibration — does the level of the estimated parameter match the reference portfolio). A dataset can be representative for one and not the other. GCD gives the canonical example: an external pool with a different mix of secured and unsecured loans relative to the application portfolio may still rank cases correctly (differentiation) while producing a biased mean LGD (quantification).

For NPL valuation, which is essentially a calibration exercise on cashflow timing and level, **representativeness for quantification is the binding constraint**. Differentiation matters secondarily, when external data is used to build segmentation/scoring within the portfolio.

2.2 Direct vs indirect measurement

GCD describes two routes:

- **Direct measurement** compares observed risk parameters between the dataset and the reference portfolio. This is feasible only when realised parameters are observable in both — possible in Case 1 below (long seller history including closed cases), generally infeasible in Case 2.
- **Indirect measurement** compares characteristics: distributions of risk drivers, qualitative dimensions of policy and portfolio, and the cyclicality of underlying volumes. When direct comparison is impossible, indirect representativeness is the only route — and the analysis must be more thorough.

2.3 Qualitative dimensions

Following GCD (2025), four qualitative dimensions should be documented:

Dimension	Practical content
Scope of application	Geographies, customer segments, product types, channels
Definition of default	90-DPD vs unlikely-to-pay; restructuring; insolvency triggers
Lending standards / recovery policies	Underwriting, servicing transfer, judicial vs amicable mix
Economic conditions / portfolio history	Cycle coverage, vintage mix, regime shifts (e.g. COVID, cost-of-living)

For NPL secondary market work, the **definition of default** is often overlooked: an EBA-template NPL has been classified non-performing for at least 90 days past due *or* unlikely-to-pay, while reference data from a debt purchaser's older portfolios may include accounts charged-off at 180 DPD. This single difference materially shifts both REC and WAL.

2.4 Quantitative dimensions and statistical tools

GCD recommends comparing the distribution of risk drivers, the level of risk parameters, and cyclicality. The statistical toolkit has matured into a practical set:

Test	Use case	Intuition / threshold
Population Stability Index (PSI)	Compare distribution of a categorical/binned variable between two samples	Rule-of-thumb: <0.10 = stable, 0.10–0.25 = moderate shift, >0.25 = significant. Must be calibrated to sample size — Yurdakul (2018) shows the standard rule is too conservative at large *n*
Hellinger distance	Continuous distribution comparison; symmetric; bounded [0,1]	Good for shape comparison when binning is unattractive

Kolmogorov–Smirnov	Equality of two empirical CDFs	Sensitive to centre of distribution; less powerful in tails
Kruskal–Wallis	Multi-group comparison of medians (non-parametric)	Useful when comparing parameter levels across vintages or sellers
t-test	Equality of means for approximately normal data	Restrictive assumptions; use with caution on bounded recovery rates

Kruger, Schutte and Verster (2021) extend this toolkit by proposing a **model-performance-based representativeness test**: split the base (external) dataset into build and test portions; build an LGD model on the candidate dataset and test it on the base test set; if performance degrades materially, the candidate is not representative. Their method is particularly valuable when distributional tests pass but the underlying functional relationship between drivers and recoveries differs — a common situation when comparing across countries or product types.

2.5 The 12-step practical guidance from GCD
Distilled from the GCD Guidelines, the workflow is:

1. Identify the reference portfolio (the application — the portfolio for sale).
2. Identify the candidate external dataset(s).
3. Define the use case (differentiation, quantification, benchmarking).
4. Build the qualitative comparison along the four dimensions.
5. Identify the relevant risk drivers in both datasets.
6. Map common variables and harmonise definitions.
7. Compare risk-driver distributions using PSI/Hellinger/KS.
8. If feasible, compare risk-parameter levels directly.
9. Where indirect, run model-performance tests (Kruger et al.).
10. Decide whether the dataset is representative; for which use case; with which segments are excluded.
11. Implement modifications: filtering, re-weighting, hierarchical pooling, or rejection.
12. Document conclusions for audit and regulatory review.

This is the essence of any defensible external-data deployment in NPL valuation.

3. External Data Sources for UK and EU Consumer NPL Valuation

3.1 Pooled industry data

Global Credit Data (GCD) is the dominant pooled source for wholesale credit, with 55+ member banks, more than 350,000 non-retail defaulted loans and 18 years of quarterly rating-migration history. Its consumer coverage is much thinner; for consumer claims, investors typically rely on:

- Servicer-internal benchmarks held by major debt purchasers. These are commercial datasets, not generally available externally.
- Specialist data vendors (e.g. Accuria, European DataWarehouse) who provide loan-level data from completed transactions and securitisations.
- Rating-agency studies — Scope, Moody's, DBRS Morningstar and S&P publish recovery and timing benchmarks for rated NPL ABS, giving investors a meaningful comparison set, though biased toward larger Italian and Greek deals.

3.2 EBA NPL data templates

The five-template structure mandatory for EU credit-institution sales of NPLs originated post-1 July 2018 and classified non-performing after 28 December 2021 is now supposed to be the standard for NPL due diligence, but adoption is still low across Europe:

- **Template 1 — Counterparty**: borrower demographics, identifiers, credit rating/scoring source.
- **Template 2 — Relationship**: links between counterparties, loans and collateral.
- **Template 3 — Loan**: contractual data, status, interest, balances, default date.
- **Template 4 — Collateral, guarantee and enforcement**: security details, valuations, legal proceedings.
- **Template 5 — Historical collections**: 36 months of monthly collection cashflows.

The template package contains 129 fields of which 69 are mandatory; proportionality reduces requirements for loans below a €25,000 carrying-amount threshold, and for transactions classified as "non-granular" (single loans, syndicated loans). The 36-month Template 5 horizon is the structural source of the right-censoring problem addressed in Section 4.

3.3 Court and insolvency statistics

The UK Insolvency Service publishes monthly individual-insolvency statistics (IVA/DRO/bankruptcy registrations) and an annual *IVA Outcomes and Providers* report with one-year and two-year termination rates. For Continental Europe, equivalent data exist but are less harmonised. Germany's *Verbraucherinsolvenz* now allows discharge after three years for consumers (transitional provisions extended to June 2025). France's *rétablissement personnel* and the broader *surendettement* commission process are tracked by Banque de France. Belgium operates a *règlement collectif de dettes* and the Netherlands a WSNP; in each case national statistics offices and central banks are the primary data sources, but the granularity is typically below UK Insolvency Service standards.

3.4 Rating agency reference curves

Scope, S&P, Moody's and DBRS apply loan-level recovery datasets obtained from ABS performance monitoring. Scope, for example, uses "historical line-by-line market-wide recovery data on unsecured defaulted loans" between 2000 and 2019 in its Italian transactions, calibrating servicer-specific haircuts on top (Scope 2025). Their published rating reports therefore contain implicit reference curves: lifetime gross recovery rates of around 8–14% for unsecured Italian consumer/SME pools over WALs of 3–4 years, with secured recoveries typically 50–60% over 5–6 years.

3.5 Macroeconomic anchors

Loan recoveries decrease in an economic downturn and banks are required to estimate the Downturn LGD. For NPL pricing purposes similar to LGD estimates, this implies that any blue-sky recovery assumptions materially above historical realisations should carry a margin of conservatism.

4. Recovery Curve Methodology

The conceptual centrepiece of NPL valuation is the **recovery curve** — the marginal or cumulative recovery rate as a function of time since default (or since acquisition). Carleo, Rocci and Staffa (2023) provide the most rigorous recent treatment for portfolios of small-balance NPLs based on anonymised data from DoValue.

4.1 Joint measurement: rate and time

Two scalars summarise an NPL portfolio's economic performance: the recovery rate (RR) and the time to liquidate (TTL) or weighted average life (WAL). Carleo et al. note that these are not independent — pricing a long-tail unsecured pool requires a *curve*, not point estimates. The recovery curve can be defined by:

- The cumulative recovery rate at time t: R_t = (sum of recoveries up to time t) / EAD
- The conditional recovery rate at time t: c_t = (recovery in period t) / (exposure at start of period t)
- The recursion: $R_t = 1 - \prod_{i \le t} (1 - c_i)$

This formulation has a deep analogy with survival analysis. The conditional rate c_t behaves like a discrete hazard, and $(1 - R_t)$ like a survival probability of unrecovered exposure.

4.2 The censoring problem — and its importance for EBA-template data

When the seller provides only 36 months of monthly collections, most loans in the tape are right-censored: their full workout is not observed. Naively averaging cumulative recoveries at t = 36 across the pool will systematically *under-estimate* lifetime recovery, because contracts that resolved early (and contributed their full cashflow) are mixed with contracts that are only partway through their workout.

The Kaplan-Meier estimator (called the product-limit estimator in survival analysis) handles this by computing each conditional rate c_t only over the cases still at risk at time t — i.e. excluding cases censored before t. Applied to NPL recoveries (sometimes called a "default-weighted survival" approach in the credit-risk literature; see Joubert et al. 2021), the estimator gives unbiased conditional recovery rates that can be compounded into a curve extending beyond the censoring horizon, *under the assumption that censoring is non-informative* (i.e. unrelated to ultimate recovery prospects).

4.3 Smoothing and extrapolation

Empirical conditional recovery curves are noisy, especially as the at-risk set thins. Carleo et al. apply non-parametric statistical learning (regression splines, generalised additive models) to smooth the conditional curve before recompounding into the cumulative curve. Confidence bands can be obtained by non-parametric bootstrap (Carleo 2023).

4.4 The exponential recovery model

For long-tail unsecured consumer claims a parsimonious functional form is widely used:

$REC(t) = REC^* \times (1 - \exp(-t / WAL))$

where REC* is the asymptotic ultimate recovery rate and WAL is the weighted-average life — interpretable as the time at which approximately 63% of the ultimate recoveries have been collected. Two parameters and a closed form make this model attractive for blending with external benchmarks (Section 5). Extensions allow time-varying hazards (Weibull) or hump-shaped curves (gamma kernels) for legal-process-driven recoveries.

4.5 Specifics for legal-status segments

For UK IVAs, for instance, the recovery curve should follow a near-deterministic pattern: 60 monthly payments for tenants, 72 for homeowners with material equity, with a typical creditor return often quoted in the 25–35p in the £ range. Industry data show that approximately 67% of debt is typically written off, with the balance recovered over 5–6 years, but with one-year termination rates of around 6% and two-year termination rates of around 14%, requiring a layered model: a completed-IVA sub-curve plus a terminated-IVA sub-curve where recoveries depend on subsequent enforcement or bankruptcy (Insolvency Service 2024).

5. Blending External and Internal Data: Bayesian-Credibility Framework

The blending of external reference curves with seller-supplied collections is the analytical heart of the framework. We use Bayesian-credibility theory, drawing on the actuarial tradition (Bühlmann; Jewell's theorem) and the Bayesian conjugate-prior toolkit.

5.1 The credibility intuition

Given a reference curve (at a minimum REC_ref, WAL_ref for the exponential decay) and observed cash flows from n loans, we want a posterior estimate that:

- equals the reference when $n = 0$,
- equals the data mean when $n \to \infty$,
- transitions smoothly between the two.

Bühlmann's credibility weight is $Z = n / (n + \kappa_0)$ where κ_0 is a "concentration" or "prior strength" parameter encoding how informative the reference is. The credibility-weighted estimate is
posterior mean = $Z \times \text{REC\_data} + (1 - Z) \times \text{REC\_ref}$

Jewell's theorem (1974) shows that for exponential-family likelihoods with conjugate priors, the Bühlmann linear estimator coincides exactly with the Bayesian posterior mean — making the credibility weight not just a heuristic but the exact Bayes solution under standard models.

5.2 Beta-Binomial conjugate model for ultimate recovery rate REC*
Treat each loan's lifetime recovery as a Bernoulli/binomial outcome (or a fractional draw on [0,1] for which the Beta is a natural prior). The Beta(α, β) prior with mean $\text{REC\_ref} = \alpha / (\alpha + \beta)$ and concentration $\kappa_0 = \alpha + \beta$ updates with n observed loans summing to S recoveries to give:
posterior Beta($\alpha + S$, $\beta + n - S$), with mean $= (\alpha + S) / (\alpha + \beta + n) = Z \cdot (S/n) + (1 - Z) \cdot \text{REC\_ref}$
with $Z = n / (n + \kappa_0)$.
The choice of κ_0 is the central modelling decision: a high κ_0 (say, 200) means the reference is treated as worth 200 "synthetic" observations and dominates until the seller tape is large; a low κ_0 (say, 20) lets even modest seller data take over.

5.3 Gamma-Exponential conjugate model for WAL/timing
For recovery timing, an exponential likelihood with a Gamma prior on the rate $\lambda = 1/\text{WAL}$ gives a closed-form posterior. With prior Gamma(a_0, b_0) (interpreted as a_0 prior recovery events totalling b_0 time units) and observed n recoveries totalling time T:
posterior Gamma($a_0 + n$, $b_0 + T$), with posterior mean $\text{WAL} = (b_0 + T) / (a_0 + n)$.
This is again a credibility-weighted blend of the reference WAL_ref and the observed mean recovery time.

5.4 Hierarchical Bayesian structure for partial pooling
Real portfolios contain multiple sub-segments (geography × product × legal status × seasoning). Hierarchical Bayesian models allow each segment to borrow strength from segment-level priors, which themselves are informed by the portfolio-wide reference. The result is **partial pooling**: small segments are pulled toward the portfolio mean, while large segments retain segment-specific estimates. This is precisely the mechanism needed for new-vintage and new-legal-status sub-pools where direct data is sparse.

5.5 Calibrating κ_0 — the most consequential analyst judgment
GCD's representativeness assessment effectively pins down κ_0:

- High representativeness on both qualitative and quantitative dimensions → high κ_0, reference dominates until substantial data accumulates.
- Mixed representativeness (e.g. similar product but different geography) → moderate κ_0.
- Weak representativeness (e.g. proxy reference for new legal status) → low κ_0, but acknowledge that posterior uncertainty remains wide.

A practical anchor: choose κ_0 such that the reference would have the same effective sample size as a peer portfolio that passed all GCD distributional tests at standard thresholds.

5.6 An illustrative table: how *Z* evolves with sample size
The table below shows credibility weights for three illustrative κ_0 values:

Observed *n*	*Z* (κ_0 = 50)	*Z* (κ_0 = 200)	*Z* (κ_0 = 1000)
10	0.17	0.05	0.01
50	0.50	0.20	0.05
200	0.80	0.50	0.17
1,000	0.95	0.83	0.50
10,000	0.99	0.98	0.91

For typical UK consumer NPL portfolios (often 5,000–50,000 accounts), *Z* will be close to 1 unless κ_0 is set very high. The crucial regime is the new-vintage or new-legal-status sub-segment where *n* is small.

6. The Decision Framework for Two Extreme Cases

The end-to-end approach depends on whether the seller has long, rich internal data or only a minimum provided by the seller and limited internal experience. We summarise the two regimes side-by-side:

Dimension	Case 1 — Long, rich seller data	Case 2 — Limited data or new vintage
Internal data depth	5+ years incl. many closed cases	≤3 years; many cases censored
Representativeness	Internal data vs current portfolio for sale	External reference vs portfolio for sale
Direct measurement	Feasible	Generally infeasible
Indirect measurement	Light	Critical
Censoring handling	Less critical	Kaplan-Meier essential
Credibility weight Z	High (close to 1)	Low
κ_0 for external prior	Low	High

Use of external data	Benchmarking, validation, downturn calibration, gap-filling	Primary anchor for level and shape
Stress testing	Standard	Wide posterior bands, scenario layers

6.1 Case 1 — Seller provides long history including closed cases not for sale
This is the scenario where the seller has been an active originator and servicer for many years and has retained data on closed (resolved) cases. The investor's analysis emphasises:

1. **Internal representativeness**: are the seller's historical closed cases representative of the portfolio currently for sale? Vintage mix, product mix, and macro conditions are the key axes. PSI on key risk drivers (DPD bucket, balance band, customer-segment proxies, geography) between historical and current pool flags drift. The Kruger-Schutte-Verster test helps when distributions look similar but functional relationships differ.
2. **Cycle representativeness**: did the historical period span a downturn comparable to the worst plausible scenario? Post-2008 UK consumer data does include the global financial crisis but largely excludes a sustained cost-of-living/high-rate environment of the 2022–2024 type, which has affected real disposable incomes and pay-down behaviour.
3. **Direct measurement**: with closed cases, the investor can compute realised recovery curves, segment them, and overlay external benchmarks for sense-checking.
4. **Use of external data**: relegated to validation and downturn overlays. κ_0 set low; Z close to 1. External curves act as "sanity priors" rather than primary anchors.

6.2 Case 2 — Limited historical data: EBA template 36 months only, or new vintage with no workout history
This is the harder, increasingly common scenario for portfolios coming through the EU NPL Directive-2021/2167-compliant processes. Key challenges and responses:

1. **Severe right-censoring in Template 5**: only 36 months of monthly cash flows are guaranteed; for a typical unsecured consumer NPL with 7–10 year economic tail, perhaps 60–80% of the lifetime cashflow is still ahead. Without survival-style methods, naive averages dramatically understate REC*.
2. **Heavy reliance on external reference curves**: the reference must come from a source representative of the seller's geography, product, default definition and vintage profile.
3. **Critical importance of the segmentation hierarchy with fallback levels**:
    - Level 1: most specific match (e.g. UK / unsecured consumer credit card / IVA / 2023 vintage).
    - Level 2: broaden one dimension (e.g. UK / unsecured / IVA / all vintages).
    - Level 3: broaden further (e.g. UK / unsecured / all legal statuses / all vintages).
    - Level 4: cross-jurisdictional analogue (e.g. Northern European / unsecured / all). At each level, κ_0 is sized to reflect the strength of representativeness; lower-level matches deserve higher κ_0.

4. **Indirect measurement throughout**: PSI/Hellinger/KS on every drivable variable. Where direct comparison is impossible (e.g. recoveries from a brand-new BNPL product), proxy-variable comparison (analogous product types) and qualitative sign-off must be documented.
5. **Bayesian *Z* close to zero** for the youngest cases; the portfolio valuation will essentially be the reference curve modulated by characteristic-level multipliers.

6.3 Specific challenges for new vintages
A 2024-vintage cohort with 12–18 months of collection data faces three issues:

- **Timing front-loading**: early recoveries are typically the easiest and largest in unsecured consumer; extrapolating linearly will overstate the tail.
- **Macro sensitivity**: the cohort has experienced a specific macro path (in 2024–2025: easing inflation, gradual rate cuts, persistent cost-of-living pressure) that may differ from the reference's history.
- **Selection effects from cleansing**: many sellers remove "easy-to-resolve" accounts via early-stage in-house collections; the remainder is biased toward harder cases. External reference data must be drawn from comparable post-cleansing pools.

7. Practical Workflow and Worked Examples

7.1 End-to-end workflow
A coherent valuation process combines the previous sections into a sequenced workflow:

1. **Receive and clean the seller loan data tape (EBA format or proprietary)**: Templates 1–5; reconcile counterparty IDs across templates; flag missing fields.
2. **Define the segmentation**: e.g. country × product × seasoning bucket × legal status × balance band.
3. **Source candidate external reference curves** at each segmentation level.
4. **Run the GCD-style representativeness assessment** (Section 2.5) for each candidate.
5. **Apply Kaplan-Meier / product-limit estimation** to seller cashflows in Template 5, segmented as above.
6. **Fit the parametric form** (Carleo et al. exponential, Weibull or spline) at each level.
7. **Set κ_0** based on representativeness — typically a small set of values (e.g. 25, 100, 500) chosen ex ante for transparency.
8. Compute Bayesian-credibility blended REC* and WAL for each segment.
9. **Aggregate to portfolio-level cashflows**, apply servicing costs, fees and taxes.
10. **Discount at the target hurdle rate** to obtain a price.
11. **Run sensitivity and stress tests**: ±10% on REC*, ±20% on WAL, downturn overlays, κ_0 sensitivity, segmentation sensitivity.
12. **Document** the entire chain for the IC paper and for FCA/regulatory file.

© Accuria Ltd. 2026

7.2 Worked example — Bayesian blending as data accumulates

Consider a UK unsecured-credit-card NPL sub-pool of 5,000 accounts, EAD £25 million. The reference curve (drawn from a peer purchaser pool with high representativeness) gives REC_ref = 18%, WAL_ref = 6 years, with assumed κ_0 = 200 in REC* and a_0 = 200, b_0 = 1,200 in WAL.

After 12 months post-acquisition, the seller-handed-over cashflows show 1,500 closed accounts with cumulative recovery of £540,000 against £6 million EAD on the closed cohort = 9% — below the reference (still consistent because the closed cohort skews "easier"). Using a Beta-Binomial update with this implied 9% over 1,500 effective observations:

Z = 1,500 / (1,500 + 200) = 0.88 posterior REC* ≈ 0.88 × 9% + 0.12 × 18% ≈ 10.1%

But the closed-case bias means we should not feed 9% directly as *REC*_data. A better approach uses the Kaplan-Meier conditional rates from the *full* 5,000-account at-risk set, projected forward via the parametric exponential model, giving x_data = 16% with effective sample size ≈ 5,000:

Z = 5,000 / (5,000 + 200) = 0.96 posterior REC* ≈ 0.96 × 16% + 0.04 × 18% ≈ 16.1%

The valuation moves materially. After 36 months — when most early recoveries are observed — Z is essentially 1 and external influence vanishes.

7.3 Worked example — new vintage with 18 months of data

A 2024-vintage UK unsecured consumer NPL pool sold in mid-2025 has 18 months of Template-5 cashflows. With heavy censoring, raw recovery is 4.5% but the Kaplan-Meier-projected REC* (using a fitted exponential) is 11% with wide error bands. The investor uses a representative external reference REC_ref = 14%, κ_0 = 1,000 (high concentration, reflecting strong qualitative/quantitative representativeness):

Z = 5,000 / (5,000 + 1,000) = 0.83 posterior REC* ≈ 0.83 × 11% + 0.17 × 14% ≈ 11.5%

A year later, with 30 months of data and a refined KM estimate of 12%, Z rises and the posterior is recalibrated. This dynamic re-estimation is standard practice for forward-flow purchase agreements.

8. Conclusion and Best Practices

External credit-collection data and reference curves are not optional inputs for institutional NPL investors operating in UK and EU consumer markets — they are the basis on which most modern transactions are credibly priced. The framework set out in this report has three interlocking pillars:

First, a rigorous representativeness assessment following for instance the GCD Guidelines, distinguishing risk differentiation from quantification, applying both qualitative dimensions (scope, default definition, lending and recovery policies, economic context) and quantitative tools (PSI, Hellinger, KS, Kruskal-Wallis), and reinforcing them with model-performance testing in the Kruger-Schutte-Verster (2021) tradition where direct distributional comparison is insufficient.

Second, a censoring-aware recovery-curve methodology following Carleo (2023), using product-limit/Kaplan-Meier-style estimators on EBA Template 5 cashflows to extract unbiased conditional recovery rates, smoothing them via splines, generalised additive models, or simple parametric models.

Third, a Bayesian-credibility blending step that combines the seller's data with a representative external reference using conjugate priors (Beta-Binomial for REC*, Gamma-Exponential for WAL) and credibility weights $Z = n/(n + \kappa\_0)$ — calibrated to the strength of representativeness evidence and stress-tested for sensitivity.

The decision rule between Case 1 (rich seller history) and Case 2 (EBA-template only or new vintage/legal status) operationalises the framework: in Case 1, Z is high, internal data dominates and external curves serve as benchmarks; in Case 2, Z is low, the external reference dominates, segmentation hierarchy and indirect representativeness analysis become critical, and posterior uncertainty must be wide enough to support an explicit margin of conservatism.

Three best practices close the report:

- **Documentation**: every step, from segmentation choice through κ_0 calibration to MoC, should be auditable. UK investors should align with FCA Consumer Duty governance expectations on board-level engagement and record-keeping; EU investors should anticipate national-competent-authority scrutiny under the NPL Directive 2021/2167.
- **Continuous validation**: as cash flows accumulate post-acquisition, the posterior recovery estimate should be refreshed quarterly. The Bayesian framework provides the natural mechanism: yesterday's posterior becomes today's prior. Drift between revealed performance and the reference curve is itself a representativeness signal and should trigger model review.
- **Humility about regime shifts**: post-COVID, cost-of-living, BNPL regulation, DRO liberalisation, and the FCA Consumer Duty have all moved the recovery distribution within the past five years. Reference curves are conditional on a regime; rigorous valuation makes the conditioning explicit and stresses the alternatives.

For institutional investors with the data engineering, statistical capability and governance to deploy this framework end-to-end, external data is more than a fallback for missing information. It is the connective tissue that turns isolated portfolios into a comparable, defensible asset class — and the discipline that makes NPL pricing repeatable across cycles, sellers and product innovations.

© Accuria Ltd. 2026

References

- Accuria 2026. https://accuria.com/valuation-of-large-loans-anchored-to-historical-recoveries/
- Accuria 2024. https://accuria.com/performance-update-for-italian-npl-securitisations-july-2024/
- Global Credit Data GCD (2025). *Guidelines on Assessing Representativeness in Credit Risk Modelling*, v1.0, June 2025. globalcreditdata.org
- Carleo, A., Rocci, R., Staffa, M. S. (2023). "Measuring the Recovery Performance of a Portfolio of NPLs." *Computation* 11(2), 29. mdpi.com
- Kruger, C., Schutte, W. D., Verster, T. (2021). "Using Model Performance to Assess the Representativeness of Data for Model Development and Calibration in Financial Institutions." *Risks* 9(11), 204. mdpi.com
- European Banking Authority (2022). *Final Report on Draft ITS on NPL Transaction Data Templates*, EBA/ITS/2022/05. eba.europa.eu
- Directive (EU) 2021/2167 on credit servicers and credit purchasers. eur-lex.europa.eu
- UK Insolvency Service. *Individual Insolvency Statistics* and *IVA Outcomes and Providers 2024*. gov.uk
- Scope Ratings (2025). *Non-Performing Loan ABS Rating Methodology*. scopegroup.com
- Joubert, M., Verster, T., Raubenheimer, H., Schutte, W. D. (2021). "Adapting the Default Weighted Survival Analysis Modelling Approach to Model IFRS 9 LGD." *Risks* 9(6), 103. mdpi.com
- European Central Bank (2024–2025). *Supervisory Banking Statistics; Annual Report on Supervisory Activities*. bankingsupervision.europa.eu

This report is prepared for institutional investor due-diligence teams and is not investment advice.
All data are as available at the time of writing in April 2026.

About Accuria

Accuria is a cutting-edge credit portfolio management platform that helps clients trade and monitor loan portfolios using a series of domain expert AI agents to automate the processing of data, documents and transactions. Accuria offers automated due diligence, data migration, valuation and reporting services for performing and non performing assets across 28 jurisdictions.

With the help of its proprietary data mapping and transformation tool Accuria helps financial institutions to map their data to a variety of data formats such as those defined by EBA for NPL transactions, EBA for the valuation in resolution, and by ESMA for securitisation disclosures. Once standardised and validated, the loan-level data can be uploaded to the Accuria valuation tool to conduct a detailed discounted cash flow analysis using pre-populated pricing parameters in different macroeconomic scenarios across all major asset classes.

Disclaimer

This paper contains confidential information about Accuria, current at the date hereof. This presentation is not intended to provide the sole basis for evaluating Accuria 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.

© Accuria Ltd. 2026