

21 September 2026

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

Part 2: Creating Trust in Valuations of Non-Performing Loans - Reporting Agents

Overview

Table of Contents	
Overview	1
1. Confident investor committee reporting for NPLs	2
2. What external recovery data can establish	3
3. Translate forecast evidence into price boundaries	4
4. Make uncertainty visible in the cash flows	5
5. Give verification and report writing distinct skills	6
6. Structure the IC memo around the bid decision	7
7. Use confidence to improve bidding discipline	9
References	10

Some investors may lose profitable non-performing loan transactions because they lack sufficient confidence in their own forecasts, despite sound valuation methods and appropriate return requirements. Building on Part 1 [1], this article explores how independent challenger valuations, using different methodologies and independently sourced, representative reference data, can help committees distinguish justified caution from unnecessary bid discounts. The unsecured consumer NPL example illustrates how external evidence can strengthen recovery estimates, while timing delays and collection costs can outweigh an apparent valuation uplift [2]. Agreement between valuations earns confidence only when their evidence is sufficiently independent and their assumptions withstand scrutiny.

Confident, comprehensive and accurate reporting makes that evidence actionable. A verification agent should reproduce calculations, challenge forecasts and identify unresolved risks; in this Part 2 we focus on a **report writing agent** who translates those findings into a traceable valuation report and a concise investment committee (IC) memo. The IC memo should lead with the requested decision, proposed price, total funding requirement, downside and approval conditions, supported by recovery analysis, cash-flow reconciliations, sensitivities and an explicit bid ceiling. This approach can help investors compete more effectively without lowering return hurdles or weakening risk discipline. Its success should be measured through forecast accuracy and realised profitability, alongside transaction wins.

1. Confident investment committee reporting for NPL acquisitions

A sound valuation creates a competitive advantage only when the investment committee understands why it is credible, what could invalidate it and how much the investor can safely pay. **IC Reporting** is therefore a critical part of the investment process itself. It converts evidence about uncertain recoveries into an accountable decision about price, funding and risk.

Our Part 1 article described an **independent verification agent** that tests forecasts using reproducible calculations and evidence outside the valuation process. [1] This follow-up article examines the next step: how skilled valuation reporting and investment committee memoranda can make that verification usable when a bid must be authorised.

The hypothesis to test

Our hypothesis is that some investors lose attractive NPL auctions because they discount forecasts they have good reason to trust. Their methodology may be sound and their return hurdle appropriate, yet uncertainty about the evidence leads to an additional reduction in the approved bid. This becomes particularly consequential when an Accuria challenger valuation, using a different estimation method and independently sourced, representative reference data, corroborates the investor's own forecast.

The mechanism is plausible. ECB research links information asymmetry to wide NPL bid–ask spreads and transactions that fail to occur. [3] Experimental auction research also finds that perceived ambiguity varies with strategic complexity and reflective reasoning. [4] These findings support investigating the confidence channel. However, they do not establish that it is the dominant cause of lost institutional NPL bids.

There are competing explanations. Another bidder may have cheaper funding, superior servicing, different tax economics or a lower hurdle. The winner may also have overpaid. Auction theory explains why rational bidders shade bids below their estimates when winning can signal excessive optimism about a shared underlying value. [5] A committee should retain protection against those risks.

The commercial objective is to increase the expected profit from transactions won within the mandate. That requires identifying the portion of conservatism attributable to weak evidence, distinguishing it from genuine downside protection, and changing it only when verification supports the change. A confident report makes that distinction explicit.

2. What external recovery data can establish

Accuria's consumer NPL paper provides an instructive example: 5,000 unsecured consumer accounts with 18 months of collections, an observed recovery of 4.5%, an internally projected ultimate recovery of 11%, and a representative external reference of 14%. Its illustrative credibility weighting gives the internal estimate five-sixths of the weight and the reference one-sixth. The blended estimate is 11.5%. [2, section 7.3]

Recovery measure	Without external data	With external data
Observed collections to date	4.5%	4.5%
Internal lifetime projection	11.0%	11.0%
External reference	Not used	14.0%
Resulting lifetime projection	11.0%	11.5%

Source: Accuria consumer NPL paper, section 7.3, p 13. These are illustrative estimates, not measured investment outcomes. Recovery denominators and lifetime versus remaining recovery must be confirmed before pricing.

The arithmetic is straightforward: $(5,000 \times 11\% + 1,000 \times 14\%) / 6,000 = 11.5\%$. The evidential question is harder. The prior strength of 1,000 is an assumption; the account count alone does not validate the weight. Correlated accounts, limited follow-up, sparse late recoveries and uncertainty in extrapolation can make the effective information content much smaller. The appropriate weight must be justified and tested for sensitivity.

The same paper's closed-account example makes a different point. £540,000 recovered on £6 million of exposure gives 9% for that closed cohort. It is not a lifetime estimate for the entire £25 million pool. Open accounts and selection into closure matters. The subsequent 16% internal estimate and approximately 16.1% blend illustrate how correcting the estimation basis can matter more than the external uplift. [2, section 7.2]

Representativeness must be demonstrated

For unsecured claims, compare product, default definition, seasoning, balance, payment status, legal route, prior collection activity, servicer and economic regime. Global Credit Data distinguishes suitability for ranking risk from suitability for estimating its level; NPL pricing needs evidence about recovery amounts and timing. Characteristic comparisons should be supported by realised outcomes where

available, and weakly matched segments should be filtered, adjusted or excluded. [6]

Carleo, Rocci and Staffa show how recovery curves can accommodate censored histories. Such methods use incomplete observations more carefully, but cannot identify an unobserved long tail without additional assumptions or evidence. [7] A report should therefore show observed collections, the supported forecast horizon and the extrapolated tail separately. Neither a large sample nor a smooth curve establishes long-term accuracy.

3. Translate forecast evidence into a price boundary

To make the example investable, the report must separate historical receipts from cash the buyer will receive. The following extension is a new analytical illustration, not a transaction valuation. Interpret 11% and 11.5% as lifetime gross recoveries on the same original exposure denominator as the historical 4.5%; remaining gross recoveries are then 6.5% and 7.0%. Historical receipts belong to the seller and are excluded.

Normalise original exposure at default to £100. Assume remaining receipts arrive at the end of years 1–5 in proportions 30%, 25%, 20%, 15% and 10%; collection costs are 20% of receipts, acquisition costs are £0.10 at settlement, and the annual unlevered hurdle is 15%. There is no financing, tax, residual tail or additional fixed cost in this simplified illustration. These timing and cost assumptions are introduced here, not supplied by the consumer paper.

Per £100 original exposure	Internal estimate	External blend
Remaining gross collections	£6.50	£7.00
Remaining net collections	£5.20	£5.60
Present value at the 15% hurdle	£3.728	£4.015
Maximum seller price before further risk constraints	£3.628	£3.915
Total outlay at a £3.50 seller price	£3.60	£3.60
Net present value at that price	£0.128	£0.415
Unlevered IRR at that price	16.8%	20.8%

Source: author calculations extending [2, section 7.3].

Discount each annual net receipt by 1.15 raised to its payment year; subtract the £0.10 acquisition cost to solve the seller-price ceiling. IRR uses the same signed cash flows and annual timing. Values shown are rounded.

The external blend adds £0.287 of present value per £100 exposure. This is a change in expected

recovery, not evidence that an uncertainty discount should also disappear. A genuine transaction model must add fixed servicing and legal costs, settlement adjustments, taxes and any contingent payments, counting every item once. Its bid ceiling must also satisfy the investor's downside, liquidity and concentration limits. [10] [11]

Independent corroboration can matter even when the central forecast is unchanged. Holding the blended forecast fixed, an illustrative unexplained 10% deduction from its present value reduces the seller-price ceiling to £3.513. Replacing that deduction with a separately justified £0.20 allowance raises the ceiling to £3.715 without lowering the 15% hurdle. The £0.201 difference is potential bidding capacity. These allowances are assumptions, not recommended haircuts; releasing either requires evidence that the remaining allowance covers the residual risk without double counting it.

4. Make uncertainty visible in the cash flows

A committee cannot assess a complex forecast from its average recovery rate. It needs the contribution, timing and fragility of the cash flows that support the bid. The detailed valuation report should make four analyses readily traceable. [11, VR03–VR11]

Portfolio composition and evidence coverage. Reconcile accounts, borrowers and balances to the model population, then compare each segment's share of exposure with its share of value. Split regular payers, irregular payers and non-payers; distinguish settlements, litigation and insolvency routes. Show what proportion of value is supported by relevant observations, by external evidence or primarily by assumptions, including the unrepresented remainder.

Recovery behaviour and execution. Show payment-plan adherence, missed payments, redefaults and transitions into paying status, with denominators and observation periods. Link those rates to collections and costs. Translate workload into staff, productivity, legal capacity, onboarding time and operating expenditure. A forecast assuming immediate full servicing capacity must carry the cost and timing consequences of achieving it.

Validation and attribution. Compare original business plans with actual collections on fixed cohorts, and separately show performance against current plans. Test genuinely held-out vintages and report bias, absolute cash error and net present value error by segment and horizon. Reconcile investor–challenger differences into recovery amount, timing, costs, segmentation and discount rate; disclose interaction effects and the order of the bridge.

Timing and downside. Report collections within 12, 24 and 36 months, cash-flow-weighted life, payback and the value of the tail. Combine recovery shortfalls with delay and cost inflation. In the pricing

illustration, a one-year delay reduces the blended seller-price ceiling from £3.915 to £3.391, more than offsetting the £0.287 benefit from external data.

Scenario at fixed £3.60 outlay	Seller-price ceiling at 15%	Investor IRR
Base external blend	£3.915	20.8%
Remaining gross receipts fall 10%	£3.513	15.2%
All receipts delayed one year	£3.391	13.9%
10% lower receipts, one-year delay, cost rate rises to 25%	£2.845	8.2%

Source: author calculations using page 3 assumptions. Variable costs follow stressed receipts; acquisition cost remains £0.10 at settlement. The ceiling is re-solved for each scenario, while IRR holds the original £3.50 seller price fixed. Scenarios are judgemental, with no assigned probabilities.

At the base timing and cost rate, remaining gross receipts can fall only 10.3% before the 15% hurdle is lost. Nominal capital recovery survives a larger 35.7% fall. That difference matters: recovering the investment is a weaker test than earning its required return. A financed transaction must also disclose peak equity funding, reserves, debt service and covenant breaches under the same stresses.

5. Give verification and report writing distinct skills

A reporting skill should define a repeatable analytical task: its inputs, approved calculations, required evidence, exceptions, output and release checks. An instruction to “write a convincing IC memo” is inadequate. NIST identifies confidently false content and fabricated explanations as risks of generative AI; polished prose can therefore increase misplaced reliance unless it is tied to verified evidence. [9]

The verification agent establishes warranted reliance

The verifier should reproduce the valuation mechanically, test the estimation method and judge the adequacy of its supporting evidence. It freezes input and model versions, checks dates and balances, and challenges assumptions using held-out outcomes. Financial calculations run in deterministic tools; the language model selects tests and explains their outputs. It records Pass, Fail, Not tested or Not applicable against explicit checks. Material failures prevent unqualified release and require human resolution; supported analysis can still proceed.

Accuria’s challenger can add information by using, for example, a separately calibrated cohort recovery

curve to challenge an investor's payment-state transition model, alongside independently sourced comparable collection histories. The comparison must align scope, costs, servicing strategy and discount conventions. Agreement is more informative when the models also reproduce observed behaviour by segment and horizon. A similar aggregate value may otherwise conceal offsetting errors.

Data independence, methodological diversity and freedom to challenge need separate assessment. If a reference curve was used to build the forecast, it cannot simultaneously verify that forecast independently. Shared borrowers, servicers, calibration periods or vendor assumptions may create correlated errors. The verifier should disclose overlap, use reserved evidence where possible, preserve dissent and report to the accountable human reviewer independently of the valuation lead. [1]

The report writing agent makes the evidence usable

The writer should produce two linked documents. The **valuation report** explains value, methods and permitted reliance; the **IC memo** requests a particular investment decision. Both use the same versioned evidence and calculations, while the memo separately explains proposed price, financing and transaction costs. Material qualifications must travel with every reused headline.

Its skills must include a metric dictionary, source-linked exhibits, distinction between facts and judgement, cross-page arithmetic checks and disciplined treatment of missing information. For every material variance it should state the amount, cause, supporting evidence, price consequence and required action. A unit mismatch or a loan average labelled as a borrower average should trigger correction before the committee reads the recommendation. The checklist's historical source checks illustrate these failure modes.

A human sets materiality, methodology, hurdles, limits and exceptions, and retains investment authority. US model-risk guidance revised in April 2026 supports proportionate validation, outcomes analysis and documentation for conventional quantitative models; it expressly excludes generative and agentic AI from its scope. Applying complementary controls to the agents is a design recommendation here, not a claim of regulatory certification. [8] A reporting specification establishes requirements. Deployment confidence must then be earned through executed acceptance checks of the calculations, evidence and resulting documents.

6. Structure the IC memo around the bid decision

The committee memo must be concise and stand alone. Put the decision, downside and unresolved conditions on page one; retain model derivations and granular schedules in the linked valuation report

and appendices. The reporting specifications support this structure.

Memo section	Required decision evidence
1 Decision and authority	Bid stage, proposed seller price, total funding, currency and dates; base and downside returns; price ceiling, requested authority and unresolved conditions. Distinguish entry bid, binding mandate and walk-away limit.
2 Thesis and transaction	Why this pool is attractive at the proposed economics, what would invalidate the thesis, feasible alternatives, acquired and excluded assets or entities, and tender milestones.
3 Portfolio and confidence	Reconciled perimeter; exposure and value by segment; history depth; diligence completed, pending and excluded by count, balance and value; reference suitability and material data limitations.
4 Recovery and servicing	Collection routes and non-overlapping proceeds; plan transitions; workload-to-staff-to-cost bridge; ramp-up, incentives, service standards, remedy and replacement rights.
5 Valuation and economics	Gross-to-net collections, dated cash flows, investor–challenger gap, asset value to seller price to settlement outlay. IRR, cash multiple, payback and NPV with explicit bases; bid-price sensitivities.
6 Downside and conditions	Joint recovery, timing and cost stresses; hurdle and capital-loss breakpoints; peak funding and financing failure. Material risks with impact, mitigation, residual exposure, owner and deadline.
7 Execution and monitoring	Diligence budget and dependencies; next-stage approval conditions; authority to revise a bid. Original underwriting baselines, monitoring triggers, response owners and a preserved decision record.

Structure is adaptable to transaction complexity.

Cross seller price with remaining recovery and, separately, delay. The illustration gives:

Seller price per £100 exposure	Base IRR	IRR with 10% lower gross receipts
£3.40	22.3%	16.6%
£3.50	20.8%	15.2%
£3.60	19.3%	13.8%

Author calculations using page 3 assumptions. Acquisition costs of £0.10 are additional to each seller price.

At a 15% hurdle, this downside supports a seller-price ceiling of approximately £3.513. The memo must state whether the mandate requires the downside itself to meet the hurdle.

Where relevant, add financing terms and a no-financing case; joint-venture contributions, capital calls

and distributions; sponsor and servicer economics; and earn-outs with returns before and after payment. Deduct operating costs from servicing fee revenue. Secured pools also need priority waterfalls and debt-to-property ownership transitions. These conditional analyses are outside the unsecured illustration.

7. Use confidence to improve bidding discipline

The authorised ceiling and the tactical bid serve different purposes. The ceiling identifies the economic boundary under the mandate. The bid reflects the auction stage, competitive evidence and the value of preserving room to negotiate. A well-supported memo should pre-authorise a bounded price range, the conditions for moving within it, and the evidence that would require a fresh decision. This can reduce delays between diligence and a binding offer.

For each proposed bid increase, show a bridge from the previously authorised price. Separate new cash-flow evidence, changed financing or costs, released uncertainty allowances and tactical use of existing headroom. Hold the return hurdle constant unless the committee explicitly changes it. Avoid imposing the same recovery risk in lower collections, a higher discount rate and an extra price haircut without explaining their distinct purposes.

Competitor intelligence belongs in that discussion with its source, date and confidence level. Where bid-history data are credible, compare probability of winning with expected profit conditional on winning, after bid costs, rather than optimising win rate alone. Conditioning on winning matters because the highest estimate can be the most optimistic one. [5] Without adequate auction data, use transparent ranges and scenarios instead of a fabricated probability of success.

Test whether underconfidence is actually costing deals

The hypothesis requires a transaction-level record. Before the outcome is known, preserve the investor's original forecast, Accuria's independent challenger, verification findings, return hurdle, every conservatism adjustment, committee price ceiling and submitted bid. Record a small set of reasons for reductions or overrides, including forecast distrust, financing, execution risk and genuine mandate constraints.

Then track auction stage reached, approval time, bid rank or clearing price were disclosed, and comparable realised collections were lawfully available. Compare similar transactions with and without a completed verification and reporting package, controlling for seller, jurisdiction, vintage, servicing, funding and market conditions. A staggered introduction of the process, with outcomes and

comparison rules specified in advance, can make the evidence more persuasive.

Measure approval-to-model price gaps, unexplained haircut reductions, forecast bias and dispersion, override outcomes and profit after costs. Validate whether high-confidence segments actually miss less often and whether stated prediction intervals achieve their claimed coverage. Preserve original business plans when current forecasts change, so revisions do not erase misses.

Strong performance on acquired portfolios alone cannot establish profitable opportunities forgone: winners are a selected sample, and lost portfolios may remain unobservable. Support for the hypothesis would require narrower unsupported discounts and more transactions within unchanged risk limits, without deterioration in forecast calibration or realised profitability. Evidence that losses are mainly explained by funding disadvantages or overoptimistic forecasts would weaken it.

At Accuria, our trust proposition is strongest when independent challenge gives committees a defensible reason to act on their own analysis. The reporting layer should show exactly which cash flows merit confidence, which risks remain and what each means for price. That is how an investor can use better evidence to compete more effectively while preserving the discipline that makes an NPL acquisition profitable.

References

- [1] Accuria (10 August 2026a). Creating Trust in Agentic Valuation of Non-Performing Loans. Part 1, especially pp 3–10 on verification independence, evidence and decision rights. Supplied article.
- [2] Accuria (30 April 2026b). Valuing Non-Performing Consumer Claims with External Reference Data and Curves. Sections 4–5 and worked examples 7.2–7.3, p 13. Supplied as Valuation of Consumer NPL with External Data(2).pdf.
- [3] Fell, J., Moldovan, C. and O'Brien, E. (2017). Resolving non-performing loans: a role for securitisation and other financial structures? ECB Financial Stability Review, May, pp 158–174. Historical evidence on market frictions, not a current estimate of bid discounts.
- [4] Deck, C., Jindapon, P., Melkonyan, T. and Schneider, M. (2024). Ambiguity, Cognitive Reflection, and Strategic Complexity across Auctions. Chapman University, ESI Working Paper 24-04. Experimental and theoretical evidence; not an institutional NPL study.
- [5] Royal Swedish Academy of Sciences (12 October 2020). The Prize in Economic Sciences 2020. Official account of Milgrom and Wilson's auction research, including rational bid shading and the winner's curse.
- [6] Global Credit Data (4 June 2025). Guidelines on assessing representativeness in credit risk modelling, version 1.0. Executive summary and sections 1.4–1.5 and 3. Working-paper framework adapted here to acquisition analysis.
- [7] Carleo, A., Rocci, R. and Staffa, M. S. (2023). Measuring the Recovery Performance of a Portfolio of NPLs. Computation 11(2), 29. Sections 2–3 on censored recovery curves, smoothing and empirical applications.
- [8] Federal Reserve, OCC and FDIC (17 April 2026). Revised Guidance on Model Risk Management, issued with SR 26-2, replacing SR 11-7. Sections V–VII; footnote 3 excludes generative and agentic AI. Cited as a governance reference for quantitative models.
- [9] NIST (July 2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. Sections 2.2 and 2.7 on confabulation and human–AI configuration.

This report is prepared with the help of Accuria's credit research agent and is not investment advice.

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.