

UNDERWRITING & CREDIT RISK

The bureau tells you where the customer stands. Cash flow tells you what is changing.

Why transaction data can add information to bureau scoring — and how a GCC lender tests whether that information is worth acting on.

SEPTEMBER 2026 DUBAI & RIYADH

01 Our position

Cash flow is not a better bureau. It is a different observation of the same customer, and the value is in what the two together support that neither supports alone. The objective is not a better score. It is a better decision.

A credit officer is not ranking a customer. They are deciding whether to extend an exposure, at what price and limit, to someone whose ability to repay may change over its life. The bureau is very good at one part of that — what has happened in the customer's credit relationships. Transaction data addresses another: what is happening in their financial position now.

A bureau file records reported credit behaviour, and it carries a good deal: balances, utilisation, limits, account age, enquiries, status. Its strongest *deterioration* signals, though, are generally realised credit events — the missed payment, the delinquency — and those arrive after the financial deterioration that produced them. A customer whose income is becoming less regular and whose month-end buffer is disappearing can remain current until the day a payment fails. The payment failure is the bureau event. The deterioration came earlier.


So the economic question is not whether cash flow beats the bureau. It is narrower and more answerable: **does cash flow contain information the current decision does not?** If it does, it has to show up somewhere a credit officer already recognises — more approvals at the same loss, fewer bad approvals at the same approval rate, better limits and prices, earlier deterioration signals.

Neither signal is the whole picture, and neither is the only alternative to the other. A consumer lending decision in this region draws on at least four observations of the same person.

FIGURE 1 · FOUR OBSERVATIONS, ONE DECISION

SCHEMATIC — NOT MEASURED DATA

Four observations of one customer



Each layer answers a different question, and the economics sit in the decision rather than in any one layer. Adding a layer is worth doing only where it changes what the decision would otherwise have been.

This paper is about the second layer, and it matters more as the file thins. Cash flow does not solve thin files automatically — transaction data can be incomplete, fragmented or misleading — but it may be the richest observable signal available where credit history is not.

Three propositions sit underneath it, and they are not equally proven.

One — cash-flow data carries information the bureau does not. Well supported by published research, and the basis of everything that follows.

Two — it can signal deterioration earlier than delinquency-based reporting. Strongly supported in principle and in adjacent evidence: a signal defined by the credit event lags the financial deterioration that produces it.

Three — it materially improves decisions on a Gulf consumer portfolio. Untested in public. A hypothesis, and the last sections of this paper are how a lender tests it on its own book rather than taking our word for it.

WHAT CASH-FLOW UNDERWRITING IS

Underwriting a customer on the movement of money through their bank account, rather than on a record of how they have repaid credit elsewhere.

THE DATA

Six to twelve months of transaction records from the accounts a customer actually banks with — every credit and debit, with date, amount, counterparty and channel. Movements, not balances. A single month-end balance is exactly the point-in-time snapshot this technique exists to replace.

HOW IT IS OBTAINED

Consent-based API access under the UAE Open Finance Circular 03/2025 and the SAMA Open Banking Framework in Saudi Arabia. For a lender's own customers, the core banking system already holds it. Uploaded or parsed PDF statements work and are materially worse — narrower coverage, weaker classification, no refresh.

WHAT IS BUILT

Raw transactions are classified — income against transfer, recurring against one-off, committed against discretionary — and then reduced to four constructs: **income continuity, obligation coverage, volatility and buffer**. Those four feed the credit decision and, refreshed monthly, the monitoring of the account afterwards. Same feed, two uses.

WHAT IT IS NOT

Not a bureau replacement — it answers a different question and the two combine. Not earned wage access or a salary advance, which are products rather than an underwriting method. Not income verification, which confirms a number at a moment; this measures its behaviour over time.

02 Where it changes a decision


One build, seven uses. The business case should not rest on a single underwriting application. The same constructs, from the same feed, run across the credit lifecycle.

APPLICATION	WHAT CHANGES	CONSTRUCTS	ECONOMIC EFFECT
Income verification and affordability	A declared figure and a salary certificate are replaced by observed inflow. Variable, multi-source and self-employed income becomes describable.	Continuity · Volatility	APPROVALS THE POLICY WAS DECLINING ON DOCUMENTATION RATHER THAN ON RISK
Debt burden ratio	Obligations the customer is actually servicing enter the calculation, including commitments no bureau reports and no form asks about.	Coverage	FEWER OVER-LIMIT APPROVALS; FEWER WRONGLY CAPPED GOOD CUSTOMERS
Thin-file and new-to-credit approval	A decision becomes possible where credit history does not exist.	All four	ADDRESSABLE POPULATION

APPLICATION	WHAT CHANGES	CONSTRUCTS	ECONOMIC EFFECT
Limit and price at origination	Exposure is sized against demonstrated surplus and cover rather than a percentage of a declared income.	Coverage · Buffer	EXPECTED LOSS PER ACCOUNT; RETURN ON CAPITAL ASSIGNED
Early warning on the existing book	Deterioration becomes visible while the account is current, ahead of any days-past-due bucket.	Buffer · Coverage, as directions	INTERVENTION TIME, WHICH IS WHAT MAKES TREATMENT CHEAPER
Collections prioritisation	Cannot-pay is separated from will-not-pay before a treatment is chosen.	Buffer · Coverage	COST TO COLLECT; CURE RATE PER CONTACT
First-payment-default screening	Circular transfers and one-off top-ups staged ahead of an application are visible in the ledger.	Continuity · Volatility	EARLY-VINTAGE LOSS

Five sit at origination and two on the existing book, and they can be built from the same underlying transaction-data layer. A lender that scopes cash flow as an underwriting project buys one of the seven and pays for all of them.

Early warning is the one most often left on the table. Refreshed monthly on accounts you already hold, the same four constructs produce a deterioration signal ahead of the first missed payment. That is not a better score. It is time, and time makes every downstream treatment cheaper.



Both hold the same bureau score today and the bureau is not wrong — neither has missed a payment. One of them is nine months into a deterioration that a record of realised events has no way to express.

THE GOVERNING CONSTRAINT

The objective is not a better score. It is a better decision.

03 What the evidence shows

Three kinds of evidence get used interchangeably here and are not interchangeable. Adoption says a technique is institutionally accepted. Predictive evidence says it ranks risk. Complementarity evidence says it adds something to what you already run — and only the third justifies paying for a second signal.

Adoption. This is not frontier practice. The most conservative underwriting infrastructure in the United States runs it, as do the bureaus that sell the scores it supplements.

Freddie Mac

LOAN PRODUCT ADVISOR · NOV
2022

The agency automated underwriting system assesses twelve or more months of checking, savings and investment activity. Announcing it, Freddie Mac said the addition of positive monthly cash flow data lets the system *"help with more accurately predicting a borrower's ability to pay."*

Fannie Mae

DESKTOP UNDERWRITER

Desktop Underwriter *"will assess a borrower's cashflow by reviewing transaction patterns, balance trends, and other observations over time in their accounts"* — transaction patterns and trends, in the agency's own words.

FICO

CASH FLOW ULTRAFICO · NOV
2025

The score vendor now builds directly on *"historical and current information about the money flowing into and out of a consumer's transaction accounts,"* seven years after the first UltraFICO pilot with Experian.

Equifax

CASHFLOW INSIGHTS


A bureau selling consumer-permissioned bank data to *"verify account ownership, income, assets, expenses, and cash flow."* The clearest available admission that a credit file alone does not answer the question.

Not fintechs taking a position: the agency mortgage system, the score vendors and the bureaus themselves. It establishes acceptance, not performance — and certainly not performance here.

Predictive power. The primary study is FinRegLab's 2019 work across six non-bank lenders: Accion, Brigit, Kabbage, LendUp, Oportun and Petal. In the populations and products studied, cash-flow variables and scores were *"generally at least as strong as the traditional credit scores and credit bureau attributes studied."*

The qualification matters and we will not drop it: six US non-bank lenders, specific products, specific populations. It establishes that cash-flow data is a serious signal, not a transferable performance figure.

The finding with commercial consequences is the second. Cash-flow data "*frequently improved the ability to predict credit risk among borrowers that are scored by traditional systems as presenting similar risk of default.*" That is the credit officer's actual question: among customers my system treats as the same, does this tell me who is different?



A bureau band is a claim that these applicants are the same risk. A second, weakly correlated signal tests the claim. The spread on the right is where a constant-loss approval rate is bought from.


Complementarity. The clearest quantification of why a second signal pays does not come from cash-flow data at all. Berg, Burg, Gombović and Puri studied digital footprints — not transactions — across 254,808 customers of a German e-commerce lender. Bureau score alone: 68.3 percent AUC. Digital footprint alone: 69.6. Both together: 73.6. The correlation between the two scores was approximately 10 percent.

We cite it with the distinction intact. This is not evidence about cash-flow underwriting. It is evidence for the principle underneath it: a new signal does not need to be stronger than the incumbent to create value — it needs to contain information the incumbent does not. That is why correlation matters, and why a weaker but different signal can improve a decision.

FIGURE 4 · WHY WEAK CORRELATION IS THE POINT

PUBLISHED DATA — DIGITAL FOOTPRINTS, NOT CASH FLOW

Correlation between the two scores: approximately 10%



The combined gain comes from the correlation, not the strength. It changes the question to put to any vendor: not what your model's Gini is, but how correlated your score is with the one we already run.

Berg, Burg, Gombović & Puri (2020), Review of Financial Studies 33(7). AUC values as reported in NBER Working Paper 24551. $n = 254,808$ scorable customers, German e-commerce deferred payment, Oct 2015 – Dec 2016. Digital footprint variables, not transaction data.

On deterioration, the nearest published support is BIS Working Paper 834: a model on machine learning and non-traditional data was *"better able to predict losses and defaults than traditional models in the presence of a negative shock to the aggregate credit supply,"* with a comparative advantage that *"tends to decline for borrowers with a longer credit history."* The advantage is largest where history is shortest.

04 What the evidence does not show

There is no published body of evidence demonstrating the performance of transaction-based cash-flow underwriting across a representative UAE or Saudi consumer-credit population. None.

Which means the following are hypotheses to test, not facts to declare: that cash flow will outperform bureau scoring; that it will produce higher approval rates at constant loss; that it will predict deterioration months earlier in every portfolio; that it will work equally well across all GCC segments.

The absence of Gulf evidence is not an argument against the technique. It is the reason to run the experiment. The Gulf question is empirical, not philosophical — and it is answerable, on your own book, in a quarter.

That gap is also the opportunity. The advantage in this technique does not sit in the model, which is replicable. It sits in the performance data a lender accumulates by originating and monitoring against it — data that cannot be bought and that compounds. The first institution here to build it owns it.

And an AUC gain is not an economic result. Three points of discrimination tells you nothing about where the improvement sits relative to your cut-off, whether it moves any decision at all, whether the additional customers are profitable at your price, or what the integration costs to run.

THE DISTINCTION THAT MATTERS

A model can rank better and change nothing.

05 How it is built

Cash-flow underwriting fails most easily in construction. The data permits thousands of variables; that is not a reason to put thousands into a decision. The standard failure is four hundred attributes built because the feed allows it, no economic intuition behind any of them, and noise fitted across ninety days.

Build four. Not because four is a magic number, but because these are the dimensions on which a repayment obligation actually fails: the income stops, the obligations outgrow it, the pattern becomes unreliable, or the reserve runs out. Each is measured twice — as a level and as a direction. The level underwrites an application; the direction monitors a customer. A healthy buffer today can still be deteriorating; a modest one can be entirely stable.

Income continuity

Regularity and persistence of credits, not size. A salary arriving in 11 of 12 months is a different risk from one arriving in 7 at the same annual total. A falling month-count registers before any payment is missed: the credit stops arriving before the obligation fails.

Obligation coverage

Committed outflows against reliable inflows. The work is classification, not arithmetic: a recurring debit to a finance company is committed, one to a supermarket is not, and an inbound transfer from a related account is not income. The hardest case in this region is the standing remittance — it behaves as a committed obligation, the customer treats it as one, and no bureau sees it.

Volatility

Dispersion of the net monthly position. Two customers can share an average monthly surplus while one is stable and the other swings between surplus and deficit. Level and stability are separate variables, and a model carrying only the first has discarded half of what the account says.

Buffer

Days of cover at the observed burn rate, measured at the low point of the month rather than at month end when salary has just landed. Buffer converts a shock into a missed payment, which makes its trajectory the shortest-horizon warning of the four.

The construction problem is classification, and it is a credit question rather than a data-science one. A transaction feed does not arrive labelled. Is an inbound transfer salary, or a movement between the customer's own accounts? Is a recurring remittance discretionary or effectively committed? Is a large cash withdrawal ordinary spending, or income leaving the observable system? Is the account you can see the primary one? A credit committee can answer those. A feature-selection routine cannot.

Four constructs with a stated economic intuition also clear a governance bar four hundred features do not. CBUAE Model Management Standards require economic intuition behind variables, independent validation and ongoing monitoring.

The model is then proved forward, not backward: windows fixed and recorded before the data is examined, scored on an application population rather than a booked one, and tested for stability as well as discrimination. Both challenger and incumbent face that test.

06 Where it fails

A credible proposition describes where its signal breaks. Cash flow degrades in specific and predictable ways, and a lender that cannot name them has not built it yet. Six matter here.

01 **The money leaves before you see it.** A worker paid through the Wage Protection System who withdraws and remits in cash presents a buffer that is real in the account and absent in life. One of the Gulf's most important distortions, and one that concentrates in exactly the segment the technique is meant to serve.

02 **The relationship is fragmented.** You see one account; the customer runs three. Single-bank visibility understates income and obligation asymmetrically, and the direction of the error depends on which account you happen to hold.

03 **Transfers are read as income.** A credit from a spouse, a related account or a savings drawdown looks like an inflow until it is classified. Categorisation quality, not data access, decides whether income continuity means anything.

04 **The window is too short, and seasonality fills it.** Ninety days establishes presence, not persistence, and cannot see an annual cycle at all. Ramadan, Eid and bonus cycles move all four constructs; a model built on the wrong three months will find a pattern that is not there.

05 **Employer risk dominates.** Transaction behaviour can stay pristine until the day employment ends. Cash flow observes the customer; it does not observe the employer.

06 **The attribute layer becomes a black box.** More attributes do not produce better decisions, and the explainability obligation will be applied to the challenger long before anything else on the stack.

None of these is an argument against the technique. They are the reason a controlled test on a local population is not optional, and the reason an off-the-shelf vendor score is a weaker proposition here than an attribute layer built on the lender's own data.

07 What is different in the Gulf


The rails exist. Coverage does not yet follow. The UAE Open Finance Regulation, now Circular 03/2025, licenses and mandates access — and onboards mandated participants in phases, beginning with banks and insurers, with later phases announced by the Central Bank. SAMA announced the commencement of licensing fintechs to provide open banking services on 26 March 2026. The question has moved from whether a lender can access the data to whether it can obtain enough coverage, quality, consent and refresh to make it useful — an implementation question, not a policy one.

Affordability is already a regulatory requirement. SAMA's Responsible Lending Principles require a clear, transparent and documented scientific method for evaluating creditworthiness and ability to repay. Nothing in that needs changing. The open question is whether the data used to answer it is observed, current and complete — and a salary certificate plus a bureau pull is a thin answer.

A debt burden ratio is only as good as the income and obligations underneath it.

CBUAE caps DBR at "50 percent of gross salary and any regular income from a defined and specific source," and the same article instructs institutions not to apply that maximum automatically but to assess each borrower's circumstances. Both halves are hard to satisfy on a declared figure: a salary certificate states an income number, but not the standing remittance, the informal obligation or the variable second income. The opportunity is not to replace the affordability framework. It is to improve the information used inside it.

Representativeness has moved the compliance burden. The CBUAE AI and Machine Learning Guidance of 11 February 2026 requires training data to be "*sufficiently accurate, relevant and representative of the customer populations to which the models will be applied.*" Against a thin-file applicant pool, an imported vendor scorecard is the harder position to defend.



A scorecard developed on thick-file salaried borrowers, or imported from a market with a mature bureau, is being applied to a population it was not trained on. The burden of demonstrating otherwise sits with the institution, not the vendor.

Which inverts the usual committee conversation: the cash-flow model is the newer object in the room and attracts the scrutiny, but it is the one trained on the population it will be applied to.

The Wage Protection System is a distinctive Gulf asset. For customers paid through WPS, income continuity can be partly anchored against a formal payroll rail rather than inferred entirely from account behaviour. Payroll plus transaction behaviour plus bureau plus internal relationship data is a stronger stack than cash flow alone, and that advantage is going unused.

And income here stops discontinuously, because it stops with residency. A customer can be financially healthy today and face an abrupt change if the employment relationship ends. The bureau may not see it. The transaction account may not see it until the salary stops. The employer axis therefore belongs alongside cash flow rather than beneath it, and it is the subject of Paper III.

08 From data to decision economics

This is where the conversation should land. The chain runs in one direction, and the last link is the only one that appears in a set of accounts.

DATA → SIGNAL → DECISION → ACTION → ECONOMICS

Transaction data becomes income continuity, obligation coverage, volatility and buffer. Those combine with the bureau into a decision; the decision becomes an action — approve, decline, limit, price, monitor, collect — and the action resolves into expected loss, revenue, capital and customer value. A measurement framework that stops at discrimination has stopped three links early.

AT ORIGINATION

- How many additional customers can be approved at the incumbent loss rate?
- How much loss can be removed at the incumbent approval rate?
- What happens to expected loss, and to risk-adjusted contribution?
- Does the signal change the optimal limit, or the price?

ON THE EXISTING BOOK

- How many currently-current customers move onto a deteriorating trajectory?
- How many months of lead time before delinquency?
- How stable is the signal, and what share of alerts are actionable?
- Which treatment produces the best economic outcome?

Then by segment — thin-file, new-to-credit, expatriate, self-employed, variable-income, high-remittance. The value of cash flow should differ across those populations. That is not a problem with the technique; it is the reason for measuring it.

09 Five tests a credit function can run now

None of this requires a programme. It requires a controlled experiment.

One — run the swap set on your declines. Take twelve months of declined applications where transaction data is held or obtainable with consent, build the four constructs, and count how many sit in the same bureau band as approved accounts that subsequently performed. Then run it in reverse: which existing approvals would the combined signal have flagged? The objective is not to prove the model, but to establish whether the incumbent decision contains measurable room.

Two — rank the existing book by trajectory, not by score. Refresh the four constructs monthly and look at customers who are current, not yet delinquent, and deteriorating on buffer or coverage. Then measure how far the signal leads the first missed payment. The question is not whether the model predicts delinquency; it is how much time it gives the credit function to act. Time is an economic variable.

Three — set the validation windows before you look. Observation window, performance window, application population, outcome definition, exclusions, missing-data treatment — written down and dated before the result is examined. Otherwise validation becomes a search for a favourable window after the fact.

Four — put the incumbent and the challenger on the same terms. The bureau score is a decisioning input; the existing credit policy is a decisioning system. Both should face the same population, performance window and economic outcomes. A challenger should not have to clear a bar the incumbent never faced.

Five — price the incremental signal. The first four test whether the signal ranks. This one tests whether it pays, and most evaluations skip it.

THE MEASUREMENT, NOT A FORECAST

Build three decisions on the same population — incumbent, challenger, combined — then ask the question that carries the money:

At your existing loss threshold, what proportion of declined applications can the challenger move across the decision boundary — and what is that worth?

The answer has a fixed shape, and every term is observable on your own book:

```
incremental value
= (accounts moved across the boundary × contribution per account)
  - incremental expected loss
  - incremental acquisition and servicing cost
  - cost of capital on the incremental exposure
```

The controlled test establishes the first term; the economics of the other terms should already be observable on the book. The count of accounts that clear the boundary — not the statistic that identified them — is what reaches the P&L. If the figure is negative, a three-point AUC gain is irrelevant. If it is strongly positive, discrimination was never the argument.

What you do with a deterioration warning — limit action, repricing, pre-delinquency treatment, retention — is Paper IV. The prerequisite is having the warning at all.

10 The credit officer's question

Cash-flow underwriting is often presented as a technology question. It is a credit question enabled by better data.

The bureau answers: what has this customer done with credit? Cash flow: what is happening in their financial position? Employer data: how durable is the income? Internal behaviour: what is changing in the relationship with us?

The credit function's job is not to choose one, but to decide which combination produces the best decision for the exposure in front of it. Which is why the strongest implementation will not be a standalone cash-flow score. It will be a decisioning layer in which each signal has a defined economic role.

The question is not whether cash flow works. Published evidence and institutional adoption have largely settled that elsewhere. The question is whether it contains information your current decision does not, and whether acting on it creates value on your book. Both are measurable.

Run the test.

AUTHOR	PUBLISHED	PRACTICE AREA
NYK Advisory Dubai & Riyadh	September 2026 Paper II of IV	02 — Underwriting & credit risk 04 — Decision economics & risk-based pricing

GET IN TOUCH

If a number in your book is moving and nobody can tell you why, that is usually the right moment for a conversation.

CONTACT	PRACTICE	THE SERIES
nykadvisory.io hello@nykadvisory.io	Credit risk & decisioning Dubai & Riyadh	Four papers through Q4 2026

REFERENCES

1. FinRegLab (2019), *The Use of Cash-Flow Data in Underwriting Credit: Empirical Research Findings*.
2. Berg, T., Burg, V., Gombović, A. & Puri, M. (2020), "On the Rise of FinTechs: Credit Scoring Using Digital Footprints", *Review of Financial Studies* 33(7), 2845–2897. AUC values as reported in [NBER Working Paper 24551](#).
3. Gambacorta, L., Huang, Y., Qiu, H. & Wang, J. (19 December 2019), "[How do machine learning and non-traditional data affect credit scoring?](#)", *BIS Working Paper* 834.
4. Board of Governors of the Federal Reserve System, CFPB, FDIC, NCUA and OCC (3 December 2019), [Interagency Statement on the Use of Alternative Data in Credit Underwriting](#).
5. Freddie Mac (17 October 2022), [Underwriting Innovation to Help Lenders Qualify More Borrowers](#). LPA cash-flow assessment effective 6 November 2022.
6. Fannie Mae, [Consider On-Time Rent Payments and Cashflow](#), Desktop Underwriter.
7. FICO (20 November 2025), [FICO Partners with Plaid to Launch Next-Generation Cash Flow UltraFICO Score](#).
8. Equifax, [Cashflow Insights](#) product description.
9. Central Bank of the UAE, [Open Finance Regulation](#), Circular 03/2025, replacing C 7/2023; phased onboarding.
10. Saudi Central Bank (26 March 2026), [SAMA Commences Licensing of Fintech Companies for Open Banking](#).
11. Central Bank of the UAE, [Regulations Regarding Bank Loans and Other Services Offered to Individual Customers](#), Article 3.
12. Central Bank of the UAE (11 February 2026), [Guidance on the Use of Artificial Intelligence and Machine Learning](#).
13. Central Bank of the UAE (2022), [Model Management Standards](#), Notice 5052/2022.
14. Saudi Central Bank, [Open Banking Policy](#).
15. Saudi Central Bank (2018), [Responsible Lending Principles for Individual Customers](#).