White Papers

Short Supply Chains and Agri-Food Linkages: Geographic Information Systems for Local Economic Development

By **Diego F. Parra** · Updated 2026-07-07 · Social Impact

QUICK VERDICT

Verdict: Traditional credit bureaus reject 62% of food-service SMEs for lack of formal history, not for real risk. Scoring built on operational data —food cost, prime cost and daily-measured waste— predicts default with 2.3 times the discriminating power of classic financial scoring. Multilateral banks that integrate GIS and short supply chains cut expected portfolio loss by 180 to 260 basis points. Operational data is the collateral that was missing.

White Paper Technical document · C-Suite & multilateral banking · 21 min read · 2026-07-07

The independent restaurant is the most numerous and most opaque economic unit in Latin America and the Caribbean. It generates formal and informal employment at massive scale, yet its five-year mortality rate exceeds 60% and its access to formal credit is marginal. The problem is not that the business is bad: it is that it is invisible to the financial system. There are no audited statements, no mortgage collateral, no sufficient bureau history. The loan officer sees an information void and, when in doubt, rejects. ECLAC estimates that MSMEs generate close to 60% of regional employment yet receive only a marginal share of business credit; in food service that gap widens because the main asset —the daily operation— appears on no financial statement.

This white paper argues that the void is an instrumentation failure, not a risk one. Every day the restaurant generates an operational data footprint —food cost per dish, prime cost, inventory turnover, waste— that predicts its solvency more accurately than any annual financial statement. SATE Institute, operating the Twin Ecosystem Model with Masterestaurant S.A.S. as technology partner, proposes a framework where that data becomes verifiable collateral, and where the geography of sourcing —short supply chains mapped with GIS— becomes both a risk-mitigation variable and a lever for local economic development. The document develops the framework in six technical chapters: macroeconomic context, the failure of the traditional approach, the scoring methodology, the solution architecture, stress-scenario simulation, and the implementation path with its KPIs and board-level ROI.

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	TRADITIONAL BUREAU	SATE OPERATIONAL SCORING
Data source	✕ Annual financials, bureau history (12-24 month lag)	✓ Daily operational telemetry: food cost, prime cost, food loss, avg ticket
SME rejection rate	✕ 62% for insufficient history	✓ 27% (real operational risk only)
Discriminating power (Gini)	✕ Gini 0.31 on food portfolio	✓ Gini 0.71 (2.3x higher)
Signal latency	✕ 12-24 months (stale balance)	✓ 24-72 hours (daily close)
Collateral required	✕ Mortgage guarantee or 120% surety	✓ Verified operational flow + GIS mitigant (local SSC)
Expected portfolio loss	✕ Base 100 (reference)	✓ -180 to -260 bps vs. base
Origination cost per operation	✕ USD 380-620 (manual file review)	✓ USD 90-140 (certified telemetry + score)
SDG alignment	✕ Not measured	✓ SDG 8, 9 and 12 with built-in M&E

Chapter 1 · The restaurant as an invisible economic unit in LAC

The independent restaurant is the most numerous and most opaque economic unit in Latin America and the Caribbean, and that opacity —not its real risk— explains its credit exclusion. ECLAC documents that MSMEs generate around 60% of regional employment but capture a marginal share of business credit; the Inter-American Development Bank estimates an MSME financing gap above USD 1.2 trillion for the region. In food service the phenomenon worsens: the business leaves a daily cash footprint but never translates it into audited statements. Diego F. Parra has seen it across dozens of operations he audits with Masterestaurant: healthy cash flow, five-year mortality above 60%, and still an automatic bank rejection. The structural problem is not solvency. It is that the financial system cannot read the asset that already exists. The food-credit gap is a macroeconomic problem, not an isolated case of entrepreneurial bad luck. The ILO reports that labor informality in LAC services hovers near 50%, and the independent restaurant concentrates much of that first-job employment.

Chapter 1 · The macro gap this framework aims to close

Without formal history or collateral, 62% of food SMEs are left out of credit by instrumentation, not by risk. The aggregate cost is twofold: formalization stalls and the employment multiplier each opening generates in its territory is lost. SATE Institute argues that instrumenting operational data closes that gap without diluting credit rigor. Daily telemetry —food cost, prime cost, waste— turns the living operation into a verifiable signal. It is the lever local economic development needed to make the region's micro food-service enterprise bankable again. For the operator the reading is direct: credit exclusion is not a verdict on the business but on the absence of instrumentation. Measuring food cost per dish and prime cost daily is not an accounting chore: it is building the asset the bank never knew to ask for. The local economic development multilateral banking pursues —SDG 8 decent work— depends on the micro food enterprise becoming legible.

Chapter 1 · Implications for the operator

Diego F. Parra insists at Masterestaurant that the first step for any operator seeking expansion capital is to stop improvising the register and start recording it daily. That record, certified by a neutral third party like SATE Institute, is what turns an automatic rejection into an approved line. Instrumentation precedes credit; never the other way around. The traditional credit bureau rejects 62% of Latin America's food SMEs for lack of formal history, not for real default risk. I have audited dozens of restaurants with healthy cash flow that were denied working capital because there were no audited statements and no mortgage collateral. The credit officer sees an information gap and, faced with doubt, rejects: it is an instrumentation decision, not a solvency one. The structural problem is that the independent restaurant is the region's most opaque economic unit, with a five-year mortality rate above 60%. Diego F.

Chapter 2 · Why does the credit bureau reject a profitable restaurant?

Parra sums it up at Masterestaurant: the business is not bad, it is invisible. The operational telemetry —food cost, prime cost, waste— exists every single day, yet no one reads it as a credit signal.

That is the error we correct. Inaction has a measurable cost: every well-underwritten loan the bureau rejects destroys value for both operator and financier. In the traditional model, manual review of a file costs USD 380 to 620, a fixed cost that makes serving the micro segment loss-making and pushes the bank to reject by default. The operator, in turn, migrates to informal credit at rates that in LAC often exceed 60% effective annual, draining the margin meant to fund growth. Expected portfolio loss under classic financial scoring, with a Gini of just 0.31, forces high provisions that make the whole portfolio expensive. The result is a small, costly, poorly discriminated book.

Chapter 2 · The quantified cost of inaction

The cost of not instrumenting operational data is paid in lost approvals, excess provisions and local development that never happens. The operator must understand that a bureau rejection is not negotiable with better rhetoric: it is negotiable with better data. As long as solvency depends on an annual statement born stale, the operator competes at a disadvantage against the bank's risk analysis. The way out is not to wait to accumulate history —which can take years— but to replace the missing history with certified present-day telemetry. Diego F. Parra frames it at Masterestaurant as a change of asset: the operator stops begging for a mortgage they lack and starts presenting the operational discipline they demonstrably prove month after month. For the micro segment, where manual origination cost makes formal credit unviable, this logic is the only realistic door. Instrumenting today is cheapening tomorrow's credit, for lender and operator alike.

Chapter 3 · Operational scoring predicts default before any balance sheet

Scoring built on operational data predicts default with 2.3 times the discriminant power of bureau history, because it measures the variable that governs solvency: cost control today. A restaurant with food cost stable at 28% and waste measured below 4% of purchases has a demonstrably better repayment profile than one with older financial statements but prime cost out of control. Telemetry captures operational maturity 12 to 24 months ahead of any annual financial statement. A balance sheet describes the past; prime cost describes the trajectory. When an operator holds prime cost below 60% for three consecutive months, they are proving cash discipline that no bureau data reveals. That is the signal a well-instrumented risk model must capitalize on, and the one

the classic financial system ignores entirely. The score combines four operational variables and a geographic mitigant, each with an explicit calibration assumption. The first is food cost per dish, capped at 32% maximum per dish under the Masterestaurant costing rule.

Chapter 3 · The model's variables and assumptions

The second is prime cost as a P&L ratio, with an alert cut at 60% and a risk cut at 65%. The third is food loss and waste, with a 4%-of-purchases threshold as the border between mature and at-risk operation. The fourth is average-ticket volatility, which proxies demand stability. The GIS mitigant weights the share of purchases within the short radius. The governing assumption is that telemetry is captured at source and reconciled across purchases, sales and recipes; without that integrity, the score is invalid. Each variable enters normalized by operation segment. Probability of default is estimated with a logistic regression over the normalized operational variables, expressible in text as PD = f(food cost, prime cost, food loss, ticket volatility, GIS mitigant). Discriminant power is measured with the Gini coefficient, which in the pilot portfolio reached 0.71 versus 0.31 for classic financial scoring —hence the 2.3x factor.

Chapter 3 · How probability of default is computed

A second formula governs the alert signal: operational variance, defined as (real cost − theoretical cost) / sales, triggers review above 3%. The geographic mitigant applies as a rate discount: up to 80 basis points when the operator verifies ≥60% of purchases in the short radius. These are not lab parameters: they are recalibrated against observed 30-, 60- and 90-day arrears. The model is honest about its limit: it predicts trajectory, not deliberate fraud, which is controlled by reconciliation. For the operator, the scoring model is a roadmap of what to measure and in what order. They need not understand logistic regression; they need to understand that four numbers —food cost, prime cost, food loss and ticket stability— determine whether they are bankable. Diego F. Parra teaches these at Masterestaurant as the indicators the owner should review before any sales report: they govern the cash and, now, the credit. Keeping prime cost below 60% and food loss below 4% is not just good operation; it is actively building a good score.

Chapter 3 · Implications for the operator

The operator who internalizes this stops reacting to the bank and starts deliberately managing their risk profile. Every clean daily close is a contribution to the information collateral that separates them from automatic rejection. Operational data becomes verifiable collateral when a neutral third party certifies it in real time, and that is the model's leap. A restaurant without mortgage collateral does own an asset: its daily footprint of food cost, prime cost, turnover and waste, generated by the point of sale and validated by the management system. When SATE Institute certifies that telemetry under the Twin Ecosystem Model, the operator presents the lender an audited signal that replaces the missing financial statement. I have seen businesses access working capital for the first time with this approach, not because risk dropped, but because it finally became legible. Collateral stops being a building and becomes proven operational discipline, month after month.

Chapter 4 · Operational data as verifiable collateral

For the 62% left outside today, that legibility is the difference between growing and dying at the five-year mark. The framework's architecture has three layers that operate as a digital twin of the restaurant's physical operation. The capture layer takes telemetry at source: POS for sales and ticket, inventory for turnover and food loss, costed recipes for food cost and prime cost. The reconciliation layer crosses those three flows and computes the

gap between theoretical and real cost, which acts as an integrity control. The certification layer, run by SATE Institute, validates the data, applies the scoring model and issues the audited signal to the financier. Masterestaurant S.A.S. provides the instrumentation —MTIE, Dashboard, Gastronomic Radar— and SATE sets the agenda, calibrates risk and runs the M&E. The role separation is deliberate: the technology partner is not judge and party to the score. That independence is what makes the process auditable for multilateral banking.

Chapter 4 · Quantified mini-case: from three rejected locations to an approved line

A three-location short chain quantifies the leap the framework produces. Before: average food cost of 33.5%, prime cost at 64%, food loss unmeasured, zero access to expansion credit and automatic bank rejection for lack of a mortgage. Intervention: 60 days of telemetry instrumentation, recipe correction bringing food cost to 27.8%, prime cost to 58% and food loss to 3.5% of purchases, plus sourcing geo-referencing with 64% of purchases in the short radius. After: a score with an implicit Gini of 0.69, a line approved 340 basis points below the informal rate on offer, and a 210-bps reduction in expected loss versus the control group. The collateral was not a building: it was the clean daily close. That is the pattern the framework seeks to replicate by cohorts. The mini-case leaves a clear operational lesson: certifying the data is worth as much as the data itself.

Chapter 4 · Implications for the operator

An operator may have healthy cash, but if it is not recorded with integrity and validated by a third party, they remain invisible to the bank. Diego F. Parra works this at Masterestaurant as the difference between knowing you make money and being able to prove it with auditable evidence. Sequence matters: first instrument, then correct food cost and prime cost to threshold, then geo-reference sourcing, and only then seek credit. Skipping instrumentation is returning to rejection. The operator who follows this order converts 60 days of discipline into a rate discount of hundreds of basis points. It is the highest-return investment an invisible food business can make: becoming legible before asking for capital. The second difference from classic scoring is geographic: the bureau treats every restaurant in a city as equivalent in its exposure to input shocks, and GIS proves they are not.

Chapter 5 · Sourcing geography as a risk mitigant

An operator who sources 60% of purchases within a 50-kilometer radius —short supply chains mapped with Geographic Information Systems— absorbs a logistics spike or a route disruption with far less damage than one dependent on suppliers 400 kilometers away. That proximity is a measurable mitigation variable, not decoration. SATE Institute, operating the Twin Ecosystem Model with Masterestaurant S.A.S. as technology partner, integrates sourcing geography into the risk profile. The result: two restaurants with identical prime cost can carry very different exposure, and the model distinguishes them. The short chain lowers cost volatility and, in the process, activates verifiable local economic development around each operator. The framework's resilience is tested with a three-level input-inflation stress simulation, comparing a short chain against a long chain with an identical 58% base prime cost. In the conservative scenario (+5% inflation), the long chain sees prime cost rise to ~61% and the short one to ~59.5%, thanks to lower logistics cost and stronger local bargaining.

Chapter 5 · Input-stress scenario simulation (5%, 12%, 20%)

In the base scenario (+12%), the long chain crosses the risk threshold at ~65% while the short stays at ~62%, still in a manageable alert zone. In the stress scenario (+20%), the long chain jumps to ~70% —probable-default territory— and the short reaches ~66%. The vulnerability differential widens with the shock: the worse the environment, the more the short chain is worth. That is why the GIS mitigant weights more in inflationary scenarios,

precisely when the portfolio needs it most. The combined effect of operational data and the geographic mitigant cuts expected portfolio loss by 180 to 260 basis points over the base, and that saving has a measurable social reflection. Lower PD translates into lower provisions, which cheapens credit for the whole cohort and makes the micro segment bankable. In parallel, each point of purchasing shifted into the short radius anchors agri-employment in the territory and cuts the logistics food-loss footprint, aligning the operation with SDG 9 and 12.

Chapter 5 · The impact on expected loss and local development

The IDB, through its #SinDesperdicio initiative, estimates that food loss in LAC equals a significant fraction of regional production; cutting it in the last mile of sourcing is at once risk mitigation and local economic development. The framework does not choose between profitability and impact: it makes them the same variable. For the operator, the stress simulation turns the short chain from an ethical preference into a financial decision. Sourcing locally is not only sustainability: it is buying insurance against imported-input inflation that, in the +20% scenario, can be the difference between surviving and falling into default. Diego F. Parra works this at Masterestaurant as margin protection: the operator who diversifies toward short-radius suppliers reduces structural vulnerability and, in passing, improves their score and rate. The practical recommendation is to map today what share of purchases comes from under 50 kilometers and set the 60% target. That number, verified with GIS, is worth up to 80 basis points of credit discount.

Chapter 5 · Implications for the operator

Few purchasing decisions carry such simultaneous financial and territorial return. Instrumenting the model starts by measuring food cost per dish and prime cost daily, not once a year, because without that cadence there is no credit signal to capitalize on. The first step is to close the register with clean data: a point of sale that logs every sale and an inventory that quantifies waste below 4% of purchases. The second is to keep prime cost below 60% for at least three months to prove sustained discipline. The third is to map sourcing with GIS and identify what percentage of purchases comes from local suppliers inside the short radius. The fourth is to certify that telemetry with a third party like SATE Institute so it becomes verifiable collateral. Diego F. Parra insists at Masterestaurant: this is not technology for fashion, it is turning the data you already generate into the asset the bank never asked for because it could not read it.

Chapter 6 · Tracking KPIs and board-level ROI

The framework is governed by a 3-, 6- and 12-month KPI dashboard that translates operation into board language. At 3 months it measures approval rate (target: from ~38% to ≥70%) and score Gini (target ≥0.65). At 6 months it tracks early arrears at 30 and 60 days (cap 4.5%) and the share of the cohort with a verified short chain (target ≥50%). At 12 months it consolidates expected-loss reduction (target 180-260 bps), origination cost per operation (target USD 90-140, versus USD 380-620 manual) and SDG 8, 9 and 12 indicators for multilateral reporting. The financier's ROI rests on three levers: lower provisions, lower origination cost and a larger bankable book. Each cohort leaves a learning curve that cheapens the next, turning the framework into an asset that appreciates with scale. This framework is rigorous but not infallible, and its rigor demands stating its limits.

Chapter 6 · Limitations and assumptions of the analysis

First assumption: telemetry is captured at source with integrity; without reconciliation across purchases, sales and recipes, the score loses validity and must not be originated. Second: the pilot figures —Gini 0.71, 180-260 bps reduction, 2.3x factor— come from bounded initial cohorts and must be revalidated by market and portfolio

size before generalizing. Third: the model predicts solvency trajectory; it does not eliminate deliberate fraud, which reconciliation mitigates but does not remove. Fourth: the GIS mitigant assumes real supplier availability in the short radius, which varies by territory and cannot be forced. Fifth: the thresholds (prime cost 60%, food loss 4%, food cost 32% maximum per dish) are calibration points, not universal truths; they adjust by segment. SATE Institute reports these limits because a framework credible for multilateral banking rests on the honesty of its primary source. The traditional bureau answers the wrong question: 'did this business have formal history?' when the question that predicts default is 'does this operator control its prime cost today?'.

Chapter 19 — The structural difference no bureau captures

A restaurant with food cost stable at 28% and measured waste below 4% of purchases has a demonstrably stronger repayment profile than one with older financial statements but prime cost out of control. Operational telemetry captures operational maturity —the variable that actually governs solvency— 12 to 24 months ahead of any balance sheet. The second difference is geographic. Classic scoring treats every restaurant in a city as equivalent in its exposure to input shocks. GIS reveals the opposite: an operator sourcing 60% of purchases within a 50 km radius —a real short supply chain— has lower structural vulnerability to imported-input inflation and a smaller logistics food-loss footprint. That proximity, mapped and verified, is at once a credit-risk mitigant and a local-economic-development lever. The bureau does not see it; the SATE framework measures it and rewards it with pricing.

POINT BY POINT

Comparative analysis: bureau vs. operational scoring

BASIS OF THE CREDIT DECISION

A · TRADITIONAL BUREAU	B · MASTERESTAURANT
Audited history and mortgage collateral	Verified operational flow and geographic mitigant

Verdict: Operational scoring approves the good operator the bureau rejects for original informality, without assuming more real risk. In the pilot, approval rose from 38% to 73% while 90-day arrears stayed under 4.1%.

SPEED OF THE RISK SIGNAL

A · TRADITIONAL BUREAU	B · MASTERESTAURANT
12-24 months (balance sheet)	24-72 hours (daily close)

Verdict: Latency drops from months to hours: the framework detects deterioration while it is still reversible, not once default is here. Prime cost crossing 62% triggers an early alert six months ahead of the missed payment.

ORIGINATION COST AND SCALABILITY

A · TRADITIONAL BUREAU	B · MASTERESTAURANT
USD 380-620 per manual file	USD 90-140 with certified telemetry

Verdict: The automated score cuts unit origination cost by roughly 75% and makes the micro segment viable, which manual review renders loss-making. That differential is what allows scaling by cohorts.

IMPACT ON LOCAL ECONOMIC DEVELOPMENT

A · TRADITIONAL BUREAU	B · MASTERESTAURANT
Neutral or not measured	Rewards SSC and cuts food loss with M&E

Verdict: The framework turns every loan into a lever for SDG 8, 9 and 12, measurable for the multilateral financier. Each point of purchasing shifted into the short radius cuts logistics footprint and anchors agri-employment locally.

SIDE-BY-SIDE COMPARISON

The traditional approach (financial bureau) STRUCTURAL FAILURE

- ✕ Assesses the audited past, not the living operation of the business
- ✕ Penalizes original informality, not real repayment risk
- ✕ 12-24 month lag: the signal arrives once default is already here
- ✕ Demands mortgage collateral that 71% of food-service SMEs lack
- ✕ Ignores the geography of sourcing and its effect on vulnerability

Operational-data scoring + GIS MASTERESTAURANT

- ✓ Measures food cost, prime cost and food loss in near real time (24-72 h)
- ✓ Turns verified operational flow into information collateral
- ✓ Maps short supply chains with GIS as a risk mitigant
- ✓ Embeds M&E and traceability for multilateral-bank reporting
- ✓ Aligns the micro-operation with SDG 8, 9 and 12, measurably

SIDE-BY-SIDE COMPARISON

Side-by-side comparison

	TRADITIONAL BUREAU	SATE OPERATIONAL SCORING
Data source	✕ Annual financials, bureau history (12-24 month lag)	✓ Daily operational telemetry: food cost, prime cost, food loss, avg ticket
SME rejection rate	✕ 62% for insufficient history	✓ 27% (real operational risk only)
Discriminating power (Gini)	✕ Gini 0.31 on food portfolio	✓ Gini 0.71 (2.3x higher)
Signal latency	✕ 12-24 months (stale balance)	✓ 24-72 hours (daily close)
Collateral required	✕ Mortgage guarantee or 120% surety	✓ Verified operational flow + GIS mitigant (local SSC)
Expected portfolio loss	✕ Base 100 (reference)	✓ -180 to -260 bps vs. base

	TRADITIONAL BUREAU	SATE OPERATIONAL SCORING
Origination cost per operation	✕ USD 380-620 (manual file review)	✓ USD 90-140 (certified telemetry + score)
SDG alignment	✕ Not measured	✓ SDG 8, 9 and 12 with built-in M&E

THE NUMBERS THAT MATTER

Figures that define the problem and the framework

62%
Food SMEs rejected by the bureau for insufficient history, not real risk

2.3x
Discriminating power of operational vs. classic financial scoring (Gini 0.71 vs 0.31)

260 bps
Maximum reduction in expected portfolio loss with operational data + GIS mitigant

72h
Operational signal latency vs. the 12-24 months of the audited balance

4%
Food-loss-on-purchases threshold separating a mature operation from one at risk

90
DAYS
Implementation roadmap for the operational-scoring framework per cohort

REAL CASE

“We had three profitable locations and zero access to expansion credit: the bank asked for a mortgage we did not have. When we connected operational telemetry to the SATE pilot, our 27.8% average food cost and sub-3.5% waste became our collateral. We went from an automatic rejection to a line approved 340 basis points below the rate we were offered informally.”

— CFO of a three-location short supply chain, SATE-Masterestaurant pilot cohort

HOW TO APPLY IT IN YOUR RESTAURANT

90-day implementation roadmap

1. **Days 1-20 · Instrumentation and baseline**
   Connect the cohort's operational telemetry (POS, inventory, costed recipes) and set the baseline for food cost, prime cost and food loss per operation. No clean data, no scoring: this phase cleans and validates the primary source. Each supplier is geo-referenced for the GIS map of short supply chains, and the food-loss threshold is set at 4% of purchases as a data-quality cut-off.

2. **Days 21-45 · Score modeling**
   Calibrate the risk model with operational variables and the geographic mitigant. The cut-off threshold is defined, Gini is validated against the existing portfolio (target ≥0.65) and assumptions are documented. The partner bank's risk committee approves the credit policy derived from the score and sets the rate discount for a verified short chain (up to 80 bps).

3. **Days 46-70 · Origination pilot**
   Originate the first credit batch under the new score with daily M&E. Approval rates, pricing and early behavior are compared against a control group served by the traditional bureau. The model is fine-tuned with real origination data and the GIS mitigant weight is recalibrated against observed 30- and 60-day arrears.

4. **Days 71-90 · Scaling and M&E**
   Consolidate the indicator dashboard for multilateral-bank reporting (SDG 8, 9, 12), certify team competencies with Open Badges micro-credentials and define the KPI to track at 3, 6 and 12 months. ROI is documented for the board —180-260 bps of expected-loss reduction— and the next cohort is prepared with the learning curve already internalized.

FAQ

FAQ for credit and program officers

Why does operational data predict default better than the financial statement?

Because it measures the living operation 12-24 months ahead of the balance sheet. Prime cost out of control anticipates cash-flow deterioration long before it appears in an annual audited statement, which is already history by the time it reaches the bureau.

How is data integrity verified to prevent manipulation?

Telemetry is captured at source (POS and inventory) with cross-reconciliation between purchases, sales and costed recipes. The gap between theoretical and real cost acts as an internal control: gaming the score requires distorting three flows at once, which the M&E can detect.

What does GIS add to credit risk, beyond local development?

GIS quantifies the structural vulnerability of sourcing. An operator with short supply chains has lower exposure to imported-input shocks and less logistics food loss, which reduces margin volatility and therefore probability of default.

Is the framework replicable beyond the pilot for multilateral banking?

Yes. It was designed to scale by cohorts with built-in M&E and reporting aligned to SDG 8, 9 and 12. Open Badges micro-credentials certify the operating team, standardizing data quality and making the process auditable for the financier.

DATA & SOURCES

Sector data 2026 (official sources)

Verifiable industry benchmarks from official, non-commercial sources (government, industry associations, market research) - not competitors.

Metric	Benchmark 2026	Source
Informalidad laboral en ALC	**≈140 millones de trabajadores informales (~la mitad del empleo regional)**	OIT
Desempleo juvenil en ALC	**13,8% en 2024 — casi el triple que el de los adultos**	OIT — Panorama Laboral 2024
Informalidad juvenil	**≈6 de cada 10 jóvenes ocupados de ALC trabajan en la informalidad**	OIT
Peso de las pymes en la economía	**≈90% de las empresas y >50% del empleo a nivel mundial**	Banco Mundial — SME Finance
Tejido empresarial mipyme en ALC	**>99% de las empresas y ≈60% del empleo formal, con baja productividad estructural**	CAF
Barreras de adopción digital mipyme	**financiamiento, habilidades tecnológicas e infraestructura: las tres barreras críticas**	CAF — Conectividad y transformación digital