The IDA Disbursement Disconnect

The weakening of project appraisal and the decoupling of money from results in the World Bank's IDA Africa portfolio

A note from mdbreform.com. Analysis is descriptive rather than causal; figures are drawn from a project-level dataset linking IEG evaluation ratings to the IDA Statement of Credits, Grants and Guarantees.

Summary

Of the projects the World Bank has approved in Sub-Saharan Africa since 2015, not one carries a calculated economic rate of return — the figure that tells the institution whether a project's benefits are expected to exceed its costs, and whether, on completion, they did. The core analytical task of the project economist has been allowed to lapse. That disappearance is the visible edge of a deeper pattern this note documents: across the evaluated IDA Africa portfolio, financial execution has come almost entirely uncoupled from development results. Of roughly **$158 billion disbursed** across the 2,576 evaluated projects examined here, **$104 billion — 65.8 percent — went to operations that did not achieve a satisfactory outcome**, and $41.6 billion to operations rated below even the Bank's own moderate-satisfactory benchmark.

Projects disburse their committed resources at near-identical rates whether they are ultimately judged successful or unsuccessful; the only execution-side variable that tracks failure is cancellation, and it is partial and late. The pattern holds across fragile and non-fragile states, investment and policy lending, and the full span of the rating scale. The evidence is most consistent with a portfolio whose incentives are organized around disbursement rather than results — a structure in which, for a lender holding a sovereign guarantee, the loan performs regardless of whether the project does. An institution that has stopped measuring whether its projects pay off is, on this evidence, no longer letting the answer govern its money.

Data and method

The analysis links two independent sources on the project code (P-code). Development outcomes come from the Independent Evaluation Group's evaluation ratings (March 2026 extract); financial execution comes from the IDA Statement of Credits, Grants and Guarantees (financeOne snapshot, 31 May 2026), rolled up from the credit line to the project. The matched core is **2,576 evaluated Africa projects** — those carrying both an IEG outcome rating and a live financial record — representing about $170 billion in original principal.

Outcomes are reported here against the **S+ standard** (Highly Satisfactory or Satisfactory). This is deliberately more demanding than the institution's customary MS+ benchmark, which folds in the Moderately Satisfactory band; the gap between the two is itself a measure of how much of the portfolio rests on the weakest passing grade. On the matched set, the S+ rate is 35.6 percent by project count and 34.2 percent weighted by disbursement; the MS+ rate is 67.2 percent.

Three boundaries should travel with every figure. The matched set describes the **evaluated, largely-closed** portfolio; it excludes recent approvals that have not yet been rated and so is

silent on the live book. The financial data is aggregated to project level, so multi-tranche operations are summed. And the analysis is associational throughout — it documents relationships between execution and outcome, not the mechanisms that produce them.

The method that was allowed to disappear

Begin with a fact that ought to be a scandal and is instead unremarked: of the Africa projects the World Bank has approved since 2015, not one carries a calculated economic rate of return. The figure that once told the institution — and the public — whether a project's benefits were expected to exceed its costs, and whether on completion they actually did, has simply ceased to be produced. Coverage falls steadily across every decade — from about seven in ten projects approved in the 1970s to one in four by the 1990s, roughly one in seven in the 2000s, and well under one percent in the early 2010s — before reaching zero for the projects approved from 2015 onward.

Approval decade	Rated IDA Africa projects	Share carrying an economic rate of return
1970s	404	71%
1980s	625	41%
1990s	597	25%
2000s	675	15%
2010–2014	312	0.6%
2015 onward	86	0%

Rated IDA Africa projects (those carrying an IEG latest evaluation and appearing in the IDA Statement of Credits), by approval fiscal year. A project counts as carrying an economic rate of return if its evaluation record contains an estimate at appraisal or at completion. The small count in the final band reflects that most post-2015 approvals are not yet evaluated; the collapse to under one percent is already complete by the early-2010s cohort, which is large.

This is not a small technical adjustment. The economic rate of return was a central analytical product of the project economist — the discipline that distinguishes a development bank from a disbursing agency, the formal answer to whether scarce concessional resources should be committed to one project rather than another.

Part of the decline is defensible, and the point should be made carefully. A growing share of operations — social protection, governance, public-sector and institutional reform — does not lend itself to a clean rate-of-return calculation, and for those a forced ERR would be more spurious than informative. The claim here is not that every operation should carry one. But the retreat is not explained by that shift in composition alone. Coverage fell even within investment lending, where cost-benefit analysis plainly applies, and the decline coincided with the rise of policy lending and a broader move away from traditional project appraisal. A core appraisal diagnostic has been allowed to lapse rather than adapted — and with it the single hard, comparable, promised-versus-delivered figure the project record once contained.

The composition shift shows up directly in the instrument mix. Development policy operations, which disburse against a matrix of policy actions rather than against appraised investments, climbed toward half of new IDA credits in Africa by the 2020s, and they carry no economic rate

of return by design. The instrument that now moves much of the money is one for which the appraisal economist's central question — do the benefits exceed the costs? — is never posed.

Where the calculation does survive, in the older portfolio, it is unflattering. Among the 584 projects that carry both an appraisal and a completion estimate, the median appraised return of 21 percent is realized at 16 percent — a shortfall of about four and a half percentage points — and the completion figure falls short of appraisal in 64 percent of cases. Even in the era when the Bank measured whether its projects paid off, they paid off less than advertised more often than not — which may itself be a clue to why the measurement was not missed when it went.

Two appraisals that did not hold: Nigeria Federal Roads (P090135) and Ethiopia Road Sector Support II (P082998)

Two large IDA road operations show what a thin appraisal costs downstream, and both were rated Moderately Unsatisfactory on quality at entry by the Bank's own reviewers.

The Nigeria Federal Roads Development Project ($333 million in IDA financing, approved 2008) was designed around an untested output- and performance-based contracting model and the creation of three new road institutions — a design built, the completion report concedes, on "fast-track pre-feasibility studies which were technically deficient," with no market sounding for the contracting approach and no analysis of the fiscal-federalism risks it depended on. Contractors declined to bid; the entire approach was abandoned in a 2011 restructuring that dropped both the contracts and the institutions. The Bank spent about $0.9 million preparing the operation and roughly $1.4 million supervising it.

The Ethiopia Road Sector Development Support Program II ($348 million in IDA financing, approved 2004) did carry an economic rate of return — and the record shows both its optimism and its fragility. Insufficient engineering design at preparation left cost estimates that bids exceeded by as much as 200 percent; the final cost reached **236 percent of the appraisal estimate**, absorbed through two successive additional financings of $87 million and $100 million. The appraised economic return of about **27 percent was re-estimated at 19 percent** on completion, and the independent evaluators downgraded the project to Moderately Unsatisfactory.

Neither outcome was a failure of supervision catching up too late. Both were set in motion at appraisal — the stage that, on the evidence of these files, draws a fraction of the effort later spent supervising the problems it leaves behind.

How the collapse happened: the IEG's own evidence on the mechanism

The coverage figures document the disappearance. They do not, by themselves, explain it. A separate IEG evaluation — *Cost-Benefit Analysis in World Bank Projects* (Independent Evaluation Group, 2010) — does, and the findings bear directly on the pattern this note documents because they describe the institutional machinery that allowed appraisal to lapse. The material below is drawn from that report; the figures are Bank-wide rather than Africa-specific, but the mechanisms it identifies operate across the portfolio and are, if anything, more pronounced in the IDA Africa setting where the coverage collapse is most complete.

The decision precedes the analysis. IEG interviewed 51 task team leaders chosen randomly from projects closing in FY2006–09. Only 5 of the 51 reported that cost-benefit analysis was given significant weight at the project identification stage. When asked whether a cost-benefit analysis had ever been the key criterion in deciding to fund a project, 82 percent said it had not. Management decisions to proceed were typically made before cost-benefit information was available — so the positive-NPV test required by OP 10.04 was not, in practice, screening projects at the point of approval. A test applied after the decision cannot reject the decision.

The analysis is treated as a compliance exercise, not a design tool. Eighty percent of task team leaders agreed that for a project proposal it was sufficient to have a cost-benefit analysis present — what several described as needing to "tick a box" in the required-documents template. Forty-one of the 51 reported that an outside consultant was hired to perform the analysis; only 10 did it themselves. Time allocated varied starkly by sector: task leaders in Public Sector Governance reported allotting just two days to the cost-benefit analysis — the lowest in the sample — against about 20 days in Transport. Ninety-two percent said contributing to the analysis did not improve their chances of promotion, and 55 percent said it did not improve their performance evaluation. IEG's own framing was blunt: there is a conflict of interest when the Bank places responsibility for the cost-benefit assessment in the hands of the same staff responsible for guiding the project through Board approval.

The quality of the analysis that does exist has declined. A review by Belli and Guerrero found economic analysis in appraisal documents to be acceptable or good in 54 percent of cases, against approximately 70 percent in comparable reviews in the 1990s — a compression toward the middle, with very poor analysis falling from 7 percent to 1 percent but the overall standard weakening. The sectoral variation was wide: Water scored 81 percent acceptable or good, Agriculture and Rural Development 67 percent, while Urban Development, Education, and Health, Nutrition and Population were rated acceptable or good in only 20 to 40 percent of cases — the same sectors where ERR coverage was lowest and where the "benefits cannot be measured in monetary terms" exemption from OP 10.04 was most routinely invoked.

The exemption is the main channel for the disappearance. OP 10.04 required the calculation of the discounted expected net present value of project benefits and costs, but allowed an exemption "if the project is expected to generate benefits that cannot be measured in monetary terms" — in which case staff were instructed to show that the project represented the least-cost alternative. IEG found that of 93 investment projects closing in FY2008 without any cost-benefit information, 60 provided no explanation or simply asserted non-applicability, 18 cited inadequate data, and 19 offered anecdotal evidence with no attempt to address selection bias. Twenty-four invoked cost-effectiveness as the standard by which the project should be judged — but of these, none actually applied cost-effectiveness analysis. The exemption designed for genuinely non-monetizable benefits had become the default route for avoiding the analysis altogether.

Failing projects are selectively invisible in the ex-post record. The complement to the front-end disappearance is a back-end one: projects with lower outcome ratings were less likely to have their ERRs recalculated at closing. In the high-CBA sectors the gap was visible and persistent; in the low-CBA sectors the probability of recalculation at closing for poorly-rated

projects was virtually zero. The difference in recalculation probabilities between low- and high-rated projects averaged 0.24 after 1987, against 0.18 before — suggesting the bias grew over time. This is the economic-accounting twin of the disbursement finding at the centre of this note: the money keeps flowing after supervision flags a project as failing, and the ex-post economic reckoning quietly disappears for the same cohort. The institution stops measuring value where the project failed, just as it stops arresting the money where the project failed.

Among the projects that did report an ERR at closing, the bottom of the distribution was censored. Of 1,299 projects closing since 1990 with a reported ERR, only 24 reported negative returns — and 16 of those sat at exactly –5 percent, with just 4 lower. IEG noted that some completion reports contained data implying negative returns but the full calculation was simply not provided. Meanwhile, the overall correlation between reported ERRs and IEG outcome ratings was only about 40 percent — and the cross-tabulation (Table B.1, 1,535 rated projects with a reported ERR, 1987–2008) showed roughly 285 projects, about 19 percent, reporting an ERR of 10 percent or lower at closing, at or below the Bank's traditional 10–12 percent hurdle rate. By the standard the institution once applied, about one evaluated project in five that even reported a number had a zero-or-negative net present value on realised returns.

> **The appraisal-collapse mechanism in summary**
>
> The IEG's 2010 evaluation documents a system in which the decision to fund precedes the analysis that is supposed to justify it; the analysis is delegated to consultants, treated as a compliance exercise, and allotted as little as two days of staff time; the exemption intended for genuinely non-monetizable benefits has become the standard route for avoiding the calculation; failing projects are selectively excluded from ex-post recalculation; and the reported ERR distribution is censored at the bottom so that negative-NPV outcomes are nearly invisible. The coverage collapse documented in this note — from 71 percent in the 1970s to zero since 2015 — is the end state of that institutional process, not an abrupt policy change.

What happened after IEG reported: the institutional response and its consequences

The IEG evaluation did not disappear into a filing cabinet. It was discussed by the Committee on Development Effectiveness (CODE) on 21 July 2010, with management present, and the institutional record of what followed is printed inside the report itself. What the record shows is an institution that recognised the problem, pledged to fix it, and then presided over the completion of the collapse it had promised to arrest.

Management's response disputed the severity while accepting the principle. Management ran its own review of 795 investment operations approved between July 2007 and December 2009 and concluded that "at least 72 percent of all operations meet the strict requirement of OP 10.04," with a further 12 percent "substantively acceptable." But the measure was different from IEG's. IEG had counted whether a project carried an economic rate of return. Management counted anything that could be argued to satisfy the policy — including the "benefits cannot be measured in monetary terms" exemption route that IEG had identified as

the principal channel through which projects avoided the calculation altogether. The 72 percent compliance figure was built, in significant part, on the exemption IEG had just documented as the mechanism of the decline.

Management also cited the Quality Assurance Group (QAG), which had conducted eight reviews between 1997 and FY2008 and found "a major improvement in the quality of economic analysis," with 96 percent of projects rated marginally satisfactory or better in the last two rounds. IEG's own assessment of QAG, in the same report, was pointed: the Quality Assurance Group's quality-at-entry ratings "are performed after projects have been approved by the Board and therefore have no effect on approval decisions." Management was citing as evidence of quality a body that IEG's report said could not affect the problem, because it reviewed projects only after the Board had already committed the money.

The Board committee was sharper than management. CODE members expressed what the Chairperson's summary recorded as "serious concerns regarding the non-compliance of OP 10.04 and accountability issues in this regard." They questioned why management had not identified and acted on the trend earlier — noting that the decline in cost-benefit analysis had begun in 1989–90, two full decades before the evaluation. Several members "considered it unacceptable that Bank operational policies are simply disregarded and also raised the issue of the Board's fiduciary responsibility." The Committee demanded an action plan from management with a timeline for "substantial remedy." The Chair escalated the findings "to the attention of other committee chairpersons and the President."

Management promised to consolidate the policy. It committed to come to the Board in the fall of 2010 with a proposal to consolidate the policy framework for investment lending, incorporating cost-benefit analysis, and to deliver a policy note on economic analysis in the second quarter of FY2011. IEG itself had already noted, with some understatement, that the investment lending reform concept note of January 2009 — the reform management was already working on — "does not mention cost-benefit analysis."

What followed. In 2013, OP 10.04 — the standalone economic-evaluation policy that had required the ERR or a justified alternative — was consolidated into OP 10.00, the new umbrella policy for Investment Project Financing. The standalone requirement was absorbed into a broader framework. In 2014, the Quality Assurance Group — the body management had cited to the Board as evidence of improving quality — was disbanded. Nothing replaced it. The independent pre-Board design review that QAG had provided, however limited IEG judged it to be, ceased to exist. And ERR coverage, which had already been falling for three decades, completed its descent: by the early 2010s it was under one percent, and for the projects approved from 2015 onward it is zero.

IEG had itself anticipated this trajectory. Its assessment of the Bank's response to earlier recommendations was blunt: the recommendations "either were not effectively implemented or did not solve the problems of biased analysis, poor compliance with cost-benefit policy, and failure to factor learning from cost-benefit results into decisions on new projects." The 2010 report's own key principle — that "cost-benefit estimates should be available and used before decisions are made to go ahead with projects" and that "the Bank should not fund projects with a negative NPV" — was not contested by any party. It was simply not implemented.

The institutional response in summary

The Board's own development-effectiveness committee called the non-compliance with OP 10.04 "unacceptable," questioned why management had not acted for two decades, raised the Board's fiduciary responsibility, and escalated to the President. Management promised a policy consolidation with a timeline. The policy was consolidated — OP 10.04 absorbed into OP 10.00 in 2013. The Quality Assurance Group was disbanded in 2014. ERR coverage reached zero by 2015. The Board that called the non-compliance unacceptable continued to approve every project that came before it without the measure it had said was essential. The 2010 evaluation is the institutional record of a problem that was diagnosed, escalated, promised a remedy, and then allowed to reach completion.

The disappearance matters beyond economics, because it is the leading edge of a wider pattern. An institution that has stopped asking whether its projects deliver value is unlikely to insist that the answer govern its money. The rest of this note documents what has filled the vacuum: a portfolio in which disbursement has come almost entirely uncoupled from outcome.

The central finding: disbursement is decoupled from outcome

If financial execution were disciplined by development results, disbursement rates would fall as outcomes worsened. They do not. The table below reports, for each IEG outcome band, the median share of original principal actually disbursed and the mean share cancelled.

IEG outcome	Projects	Median disbursed (gross)	Mean cancellation
Highly Satisfactory	53	98.8%	2.1%
Satisfactory	864	99.4%	3.9%
Moderately Satisfactory	813	97.8%	4.8%
Moderately Unsatisfactory	383	97.8%	8.8%
Unsatisfactory	414	97.3%	16.0%
Highly Unsatisfactory	38	80.7%	34.3%

A project ultimately rated Unsatisfactory disburses about 97 percent of its commitment — statistically indistinguishable from one rated Satisfactory. Only at the extreme floor of the scale, among the 38 Highly Unsatisfactory projects, does disbursement materially fall away. Collapsing the scale into the two halves the Bank itself uses, satisfactory projects disburse a median 99.3 percent of commitment and unsatisfactory ones 97.5 percent. The difference is trivial; the resources move at essentially the same pace toward success and toward failure.

This is the empirical content of a familiar observation about development finance: a lender protected by a sovereign guarantee is repaid whether or not the operation it financed delivers. That guarantee is the mechanism behind the pattern in the table. Because the credit is serviced by the sovereign regardless of whether the project works, no party in the disbursement chain absorbs a loss from continuing to fund an operation that is already visibly failing — so nothing

in the financial machinery has cause to arrest the flow. What the data adds is that the indifference runs upstream of repayment, into disbursement itself. The money is not merely recovered regardless of outcome; it is *deployed* regardless of outcome.

What the mechanism looks like in a single project: DRC Multi-modal Transport (P092537)

The DRC Multi-modal Transport Project — an IDA investment operation approved in 2010 to restore the national railway and strengthen transport state-owned enterprises — was rated **Unsatisfactory** at completion, with Risk to Development Outcome **High**. Its supervision record shows when the money moved relative to when the Bank knew.

The Bank's own implementation reports rated the development objective Moderately Satisfactory or better for four years, peaking at Satisfactory in September 2013. The rating was then cut to **Unsatisfactory in August 2014, with about $252.8 million already disbursed**, and it never recovered — staying at Unsatisfactory or Moderately Unsatisfactory through every subsequent report to closing. Yet disbursement continued, to roughly **$385.6 million**. About **a third of the total — some $133 million — went out after the Bank's own supervision had flagged the project as Unsatisfactory.**

Implementation report	Date	Development-objective rating	Cumulative disbursed
ISR 7	Sep 2013	Satisfactory	$175.6m
ISR 9	Aug 2014	**Unsatisfactory**	$252.8m
ISR 11	Dec 2015	Unsatisfactory	$324.7m
ISR 14	Oct 2017	Moderately Unsatisfactory	$366.9m
Closing	Jun 2018	**Final outcome: Unsatisfactory**	~$385.6m

The decision to keep funding was explicit rather than inadvertent. The completion report records that suspension of the railway component was proposed by the project team in 2014 and again in 2015 but did not occur, because it "would have had negative impact on the Bank–Government relationship." A $180 million additional financing had been approved in June 2013 — two months before the mid-term review that might have prompted an exit.

The appraisal that launched the operation was thin against the stake. The Bank spent about **$1.1 million preparing the project and roughly $3.2 million supervising it** — three times as much watching the operation as deciding whether to start it. An economic rate of return was calculated for the largest component only, resting on the assumption that the railway would collapse and force freight onto the road, a shift that had already occurred by appraisal; efficiency was ultimately rated Negligible. The completion report's own verdict on quality at entry was an "unrealistic" central objective, set without reliable financial data, with institutional diagnostics deferred to the implementation phase. The two threads of this note — appraisal that no longer establishes whether a project is worth doing, and disbursement that proceeds regardless of whether it is working — are both visible in this one file.

Cancellation: the only execution signal that tracks failure

The one place where execution and outcome do move together is cancellation. Mean cancellation rises monotonically across the rating scale, from roughly 2–4 percent in the satisfactory bands to 16 percent at Unsatisfactory and 34 percent at Highly Unsatisfactory. Failing projects, in other words, do not draw down less of what they retain — they retain less, and disburse the remainder in full.

Cancellation is therefore the portfolio's de facto, and only, financial accountability mechanism. But it is a blunt and trailing one. Across the matched set, total cancellations come to about $10.7 billion against $158 billion disbursed; the signal arrives, where it arrives at all, after commitment and often after substantial drawdown. It registers trouble; it does not prevent it.

The pattern is structural, not a fragility artifact

A natural objection is that the decoupling reflects fragility — that the Bank disburses to failing projects because it operates in failing states where outcomes are bound to disappoint. The data does not support this reading. Projects in FCS-classified countries rate slightly *higher* on S+ than those in non-FCS countries (38.2 versus 35.2 percent), cutting against the assumption that fragility mechanically depresses outcomes. More importantly, the disbursement decoupling is present in both groups in the same form: in fragile and non-fragile settings alike, satisfactory and unsatisfactory projects disburse at near-full and near-equal rates, with cancellation the only differentiator.

The decoupling is thus a property of the portfolio's operating model, not of the difficulty of the environments in which it works.

The macroeconomic practice and the instrument of least accountability

Disaggregating by Global Practice surfaces a second, narrower pattern that this note records descriptively and leaves for fuller treatment elsewhere. The Macroeconomics, Trade and Investment practice (MTI) is the largest single GP in the evaluated Africa portfolio (374 projects) and has the **lowest S+ rate of any major practice at 21.1 percent**, against a GP-wide average near 29 percent. It is also the practice that runs policy lending: roughly 71 percent of Africa's Development Policy Financing operations sit under MTI, and DPFs as a class underperform investment lending (28.9 versus 38.1 percent S+).

The instrument and the practice are difficult to separate, because they substantially coincide. But one observation cuts against the easy explanation that policy lending is simply harder to rate well: on its small investment-lending portfolio, MTI rates *lower still* (10 percent S+, on 30 projects — a thin sample, but a striking one). The practice underperforms whatever instrument it touches. Whether this reflects mandate, capability, or the design role macroeconomic staff have come to play in operations is a question of institutional political economy taken up in two companion analyses on mdbreform.com — *Policy Without Performance: Isomorphic Mimicry and the DPO Incentive Trap* on the instrument, and *Institutional Power Architecture and Portfolio Distortion at the World Bank* on the practice — which argue that neither has

strengthened Bank operations in the way intended. What the data establishes here is the descriptive fact: the worst-performing major practice also owns the worst-performing instrument.

Time in the portfolio: how operations move, and what it predicts

The lifecycle dates — Board approval, effectiveness, closing, last disbursement — describe how operations move through the portfolio, and offer a partial test of whether that movement is disciplined by results.

The shape of the lifecycle. A median IDA Africa credit takes about **5.3 months from Board approval to effectiveness** — the point at which disbursement may begin, and the closest available proxy for disbursement onset, since the source carries no separate first-disbursement date. A quarter of credits are effective within 3.2 months; a quarter take longer than 8.2 months; and **one credit in five takes more than nine months** simply to become effective. Implementation, from effectiveness to closing, runs a median of **5.1 years**.

Time-to-effectiveness varies far more across countries than implementation length does, and it sorts them revealingly. The figures below are credit-level medians for borrowers with at least twenty credits.

Slowest to effectiveness	Median lag	Share >9 mo	Fastest	Median lag
Nigeria	9.5 mo	54%	Cabo Verde	2.9 mo
Cameroon	8.1 mo	41%	South Sudan	3.5 mo
Angola	7.9 mo	38%	Sierra Leone	3.6 mo
Uganda	7.5 mo	40%	São Tomé & Príncipe	3.7 mo
Congo, Rep.	7.1 mo	32%	Comoros	3.7 mo
Congo, Dem. Rep.	7.0 mo	35%	Côte d'Ivoire	3.8 mo
Sudan	6.9 mo	40%		

The slowest starters are the large, institutionally complex states. Nigeria is the extreme case: more than half its credits take over nine months merely to become effective, consistent with federal ratification and counterpart-funding frictions. The fast group is the more surprising one — several are fragile or post-conflict states (South Sudan, Sierra Leone), which cuts against the assumption that fragility means delay. The likely explanation is again instrumental: fragile-state portfolios lean on emergency and budget-support operations engineered to disburse quickly, so fragility and speed coincide rather than oppose. These per-country medians rest on differing numbers of credits and are best read as indicative.

Across sectors the gradient runs from infrastructure to policy. The investment-heavy practices are slowest both to start and to finish — Energy (6.9 months to effectiveness, 6.2 years to implement), Water and Transport (around six months and 6.4 years) — while the practices closest to policy lending are fastest, with Macroeconomics, Trade and Investment an outlier at

1.6 months and a single year. As the practice analysis above notes, that speed is the signature of the DPF instrument MTI runs, not of efficient project execution.

What the lifecycle predicts. Two outcome tests follow from these dates. The first concerns readiness at entry, proxied by the approval-to-effectiveness lag. This turns out to be a weak and inconsistent predictor of outcome: the overall association is slight, and its direction flips by sector — slow effectiveness predicts worse outcomes in the human-development practices but better outcomes in infrastructure and agriculture, where thorough preparation of complex works appears to be a virtue rather than a warning. Effectiveness lag should not be read as a portfolio-wide quality signal.

The second test is sharper. **Implementation duration is strongly and monotonically associated with outcome.** Among investment projects, the S+ rate falls from 48.6 percent for operations completed in under three years to 30.4 percent for those running beyond seven. This relationship is partly endogenous — troubled projects are restructured and extended, so duration is in part a symptom of difficulty rather than a cause of it — and it is reported here as a descriptive association. But the direction is unambiguous: prolonged time in the portfolio is the company of failure. Read alongside the disbursement finding, it completes the picture: money leaves regardless of outcome, and the operations that take longest to push it out are the ones most likely to fail.

A note on apparent progress. The median time to effectiveness has fallen over the decades, which might suggest improving readiness. Decomposed, it does not. For investment lending, the approval-to-effectiveness lag has been roughly flat across thirty years. The aggregate improvement is largely compositional: the share of new commitments made through fast-disbursing policy operations has risen sharply, most recently to about half of new credits. The Bank has become faster at moving money chiefly by doing more policy lending — not by getting better at launching the investment operations that constitute the bulk of its development work.

Project size: the smallest succeed, the largest run longest

Crossing size with outcome and timing — within investment lending, to hold the instrument constant — adds a third structural regularity, and not the one a "larger means better-resourced" intuition would predict. Outcomes are strongly graded by size, with the smallest operations the clear best performers.

Size (original principal)	Projects	S+ rate	Median implementation
<$25M	952	44.6%	5.8 yr
$25–75M	675	31.7%	6.3 yr
$75–150M	227	30.0%	6.8 yr
$150–300M	132	28.0%	7.8 yr
>$300M	56	33.9%	8.0 yr

Operations under $25 million reach a satisfactory outcome about 45 percent of the time; everything larger clusters between 28 and 34 percent. The robust reading is that small succeeds

— the sub-$25-million operations, typically simpler and more narrowly scoped, outperform the rest of the portfolio by a wide margin — while the mid-to-large investment project, complex enough to be demanding but not large enough to command flagship attention, is the structural weak spot. Across all instruments the very largest projects recover somewhat, but that bounce is partly the presence of large policy loans and is muted within investment lending; it should be read cautiously.

Size also entangles with the duration finding. Implementation length rises steeply and monotonically with size — from 5.8 years for the smallest operations to 8.0 years for those above $300 million — and large projects almost never finish quickly. Size and prolonged implementation therefore travel together, and the two penalties are substantially the same phenomenon: a large operation is, by construction, a long one, and length is itself the company of failure. The best-performing profile in the data is the small, quickly-implemented project, which reaches a satisfactory outcome 46 percent of the time. Effectiveness lag, by contrast, barely varies with size — the delay to disbursement onset is a function of country and instrument, not of how large the operation is.

One contrast with the non-concessional window is worth noting. The largest tenth of IDA operations by disbursement hold some 45 percent of all the money but rate at about the portfolio average, so concentration does not pull the disbursement-weighted outcome far below the project-count figure. On the IBRD side the equivalent large loans rate markedly worse — which is why that window's headline success rate collapses when weighted by disbursement, and IDA's does not.

Performance ratings that move in lockstep

The evaluation record attaches two performance ratings alongside the outcome — one for the Bank's own conduct, one for the borrower's — and these are, in principle, where responsibility for a result is located. In practice, for the Africa portfolio, they do little independent work.

The comparison can be made only for the period in which both were rated. The Bank rated borrower performance consistently through evaluation-year FY2017 and then phased it out — 78 percent coverage in FY2018, 34 percent in FY2019, effectively none by FY2022 — while continuing to rate its own performance throughout. On the 1,770 Africa projects rated on both, in that pre-2018 window, Bank performance and borrower performance are judged almost identically, and almost identically poorly: 38.4 and 35.8 percent satisfactory-or-better respectively. The two ratings are equal in 70 percent of projects, with a modest asymmetry where they diverge — the Bank is rated above the borrower somewhat more often than the reverse (18 versus 13 percent).

What the ratings do not do is discriminate responsibility when a project fails. Among the operations whose outcome fell below the line, IEG marks down both the Bank and the borrower together in 88 percent of cases; it faults the borrower alone in 7 percent and the Bank alone in just 4. The Bank's own performance — its design and its supervision — is rated sub-satisfactory in 92 percent of failed projects. The performance assessments, in other words, are co-determined with the outcome rather than independent diagnoses of it: when the result is bad,

both parties' performance is marked bad in tandem, with a slight institutional inclination to spare the Bank.

This closes a circuit with the central finding. The evaluation record acknowledges, for the large majority of failing projects, that the Bank supervised them poorly — and yet, as the disbursement evidence shows, those same projects drew down some 97 percent of their funds. The supervision rating registers the failure; the disbursement system proceeds as though it had not. And the retirement of the borrower-performance rating after FY2017 is of a piece with the disappearance of the economic rate of return described at the outset: within a few years the institution let lapse two of the measures by which a project's performance, and the responsibility for it, might have been independently judged.

Conclusion

The Bank has not merely become disconnected from results. It has become disconnected from appraisal itself. Taken together, the findings describe an IDA Africa portfolio that functions as a disbursement system loosely coupled to a development-results system. Commitments are made; resources flow at near-full rates irrespective of where projects will land on the rating scale; the only financial brake — cancellation — engages partially and late; the largest and longest-running operations underperform the smallest and simplest, and the operations that disburse most slowly are the ones most likely to fail; and the one quantitative accountability metric the institution once maintained has been allowed to lapse precisely where it would now be most exposing. None of this is confined to the hardest environments, and none of it is explained by them.

The framing that best fits the evidence is institutional rather than individual. A lender holding a sovereign guarantee faces no balance-sheet consequence from a project's failure: the credit is serviced regardless. Internal performance, meanwhile, is measured substantially by commitment and disbursement against targets. In that incentive structure, the development outcome becomes an after-the-fact annotation rather than a gate — recorded by an independent evaluator, but downstream of every decision that mattered. The portfolio behaves accordingly. The reform implication is not that the Bank should disburse less, but that disbursement and outcome should be re-coupled — that the financial machinery should answer to the results system rather than running in parallel to it. That re-coupling would begin where the decoupling did: with restoring the analytical function the institution let lapse — the discipline of estimating, before money is committed and again after it is spent, whether a project was worth doing — and with insisting that the answer, once known, is allowed to move the money.

Methodological note

- **Outcome standard.** S+ = Highly Satisfactory or Satisfactory. The Bank's institutional benchmark, MS+, additionally counts Moderately Satisfactory; this note anchors on the more demanding S+ throughout, consistent with the platform's methodology.
- **Matched set.** 2,576 evaluated Africa projects (Western & Central plus Eastern & Southern Africa) carrying both an IEG outcome rating (March 2026) and a financial record from the

IDA Statement of Credits (31 May 2026). Approximately 1,462 financed-but-unevaluated projects and all unrated recent approvals are excluded by construction; the analysis therefore speaks to the evaluated, largely-closed portfolio, not the live book.

- **Disbursement ratios.** Reported gross, against original principal. A net-of-cancellation variant carries exchange-rate noise on XDR-denominated credits and is used only as a cross-check.
- **Causation.** All relationships are descriptive. The implementation-duration association is partly endogenous (extensions follow trouble); the MTI/instrument relationship is confounded by the near-coincidence of the practice and the DPF instrument. Mechanism and attribution are beyond the scope of this note.
- **Economic rate of return.** A legacy-era measure, taken from the September 2022 IEG file (the source carrying ERR fields). Coverage is computed on rated IDA Africa projects by approval fiscal year — 2,699 projects with an IEG latest evaluation that appear in the IDA Statement of Credits — and a project counts as carrying an ERR if its record holds an appraisal or completion estimate. Coverage collapses to under one percent by the early-2010s cohort and to zero for post-2015 approvals, so the appraisal-versus-completion finding (584 projects with both estimates) describes the older portfolio only.
- **Bank and borrower performance.** Both ratings are taken from the September 2022 IEG file, the only source carrying borrower performance (the 2026 file omits it). Borrower performance was rated through evaluation-year FY2017 and discontinued over FY2018–2022, so the Bank-versus-borrower comparison is confined to 1,770 Africa projects rated on both in that window.