

Clearing the Clogs

A plumber's guide to making the reformed LIC-DSF run at scale

Teal Insights · 2 July 2026

CONTENTS

The reforms are only as good as the plumbing that runs them	2
A plumber's assessment	3
Three clogs, one gap: the framework was never written to a shared standard	3
Clog 1: The concepts are documented. How to run the template is not. This is a drain on capacity-development resources and a subsidy for sovereign advisors.	4
Clog 2: A credit assessment this consequential must be testable. Run as undocumented software, it cannot be.	5
Clog 3: The long-term module is the reform's most promising piece, and it stays blocked until two pipes are cleared: the models' calibration, and the connector to the template.	10
Write it down to a standard, so the tool carries the toil and the people do the policy	13
Fix 1: publish the data dictionary, in names a human and a machine can both read .	16
Fix 2: rebuild the workbook clean, and start laying the ground beyond it	17
Fix 3: build the connector and a shared parameter library, starting now	18
Fix 4: a semiannual forum that owns the living standard	18
The fixes are boring, doable, and cheap	20
Sources	21

EXECUTIVE SUMMARY

On paper, the LIC-DSF is a framework. To the boards and creditors who act on it, a risk rating. To the roughly 70 countries that run it, an opaque, cumbersome spreadsheet. We build open-source versions of these tools for a living, so we read this one cell by cell.

So the question your review turns on is not whether the new ideas are good. It is whether the lean Fund teams and finance ministries that run the framework in those 70 countries can make the new ideas run. Every reform adds another input to find and another step to re-run, on plumbing already clogged with toil: the manual work of running the framework that a machine could do, repeated by every team,

every cycle. **And a credit assessment this consequential must be testable. Run as undocumented software, it cannot be.**

The fix is boring, cheap, and overdue: write the framework down to a standard, so a tool carries the rote work and the scarce economists on every side (the Fund's, the Bank's, and the ministries') spend their hours on policy. Today that logic is guild knowledge, held in the formulas and a few experienced heads and written down nowhere, so it is taught by hand, team by team. Write it down once, and one standard serves all ~70 teams every cycle. Four fixes carry it: a data dictionary, a clean rebuild to a published standard, a model connector with a shared parameter library, and a semiannual forum that keeps all three current. Order of magnitude, that is a one-time million-dollar effort against the half a billion a year the Fund already spends on capacity development. The prize is a framework that stops being only a rating and becomes a decision-support tool more countries can run themselves.

We leave the hardest part where it sits, the macroeconomic projections that no consultation response fixes cheaply, and take three clogs a modest effort can clear now. Clear the plumbing first, and the good ideas land.

THE FOUR FIXES

- 1 · a data dictionary
- 2 · a clean rebuild to a published standard
- 3 · a model connector, a shared parameter library
- 4 · a semiannual forum that keeps all three current

The reforms are only as good as the plumbing that runs them

On paper, the LIC-DSF is a framework. To the boards and creditors who act on it, it is a risk rating. To the people who run it, it is an opaque and cumbersome spreadsheet.

So the real question is not whether the new ideas are good. It is whether they can be run in the ~70 countries that use the framework, by the Fund teams and finance ministries who do the work. Running it today already costs more than any invoice shows: the manual rework every team repeats by hand, the week of back-and-forth to make one tool's numbers fit another's, the errors no one can catch. Every reform adds another input to find and another step to re-run. Clear the way first, and the ideas land. That is the case we make here.

We are tool builders. We build SovTech: open-source, philanthropically funded software tools for sovereign-debt analysis, the LIC-DSF among them, made with finance ministries and technical partners. Building a tool means examining every input, every transformation, and every number it returns. Once you have, you see where the pipes join, where they leak, and where they are missing.

This gives us the plumber's perspective, and that is what we share here.

A framework only helps if it moves. Each reform is one more thing to push through a pipe that only takes so much.

~70

countries run the LIC-DSF, resource-constrained lower-income countries and small island states

What backs up the pipe is toil: the manual, repetitive work of running the framework that adds no analytical value and grows with every team and every cycle.¹ Today the operating logic lives as guild knowledge: it is written down nowhere, held in the formulas and in the handful of experts who pass it from team to team. So each team re-derives it by hand. Write it down precisely, and a tool can carry the work instead of the people. That frees the scarce economists on every side (the Fund's, the Bank's, and the ministries') for the policy substance only they can do. It is also the one thing the reforms need. They add the substance everyone wants, but only a framework that has shed its toil has room to run it.

The Fund and the Bank set out to keep the framework comprehensive, transparent, easy to use, and future-proof. We share those goals. The people who must run it are lean Fund and Bank teams and the national authorities, in lower-income countries and small island states. The review's own terms concede the strain: keep it tractable, limit the capacity-development bill. Clearing the existing plumbing is how you keep that promise as the new machinery goes in. The cost of leaving it clogged is high and rising.

Toil

Hands-on work a machine could do, repeated by every team, producing nothing that lasts. It is what a clogged pipe is full of.

A plumber's assessment

Three clogs, one gap: the framework was never written to a shared standard

This assessment does not try to catalog every problem, and it does not take up the hardest one. The slowest step in any debt-sustainability assessment is building the macroeconomic projections beneath it. That is hard work, and no consultation response clears it cheaply, so we leave it where it sits. We take three clogs that throttle the framework now and that a modest, deliberate effort can clear:

- 1. The framework is well documented; how to run the template is not.** The concepts are taught well, but the operating manual for the Excel template was never written, so countries pay to be shown, team by team and cycle by cycle.
- 2. A credit assessment this consequential must be testable and auditable; today's Excel workbook is neither.** That is why the workbook cannot remain the default standard.
- 3. The long-term module will not run at scale until two pipes are cleared.** Its calibration dials and its output connector do not exist yet, so the most promising piece of the reform arrives blocked.

These are three faces of one gap, and each is cheap to clear. Clearing them protects the very tractability the review is trying to keep. Blame the plumbing, not the users. They take heroic care, and the fixes are about the tool they were handed.

¹We borrow a term from software engineering. Toil is work that is manual, repetitive, automatable, and scales linearly as a system grows: the kind worth designing away so people can spend their time on what needs judgment. Google, *Site Reliability Engineering* (O'Reilly, 2017), ch. 5, "Eliminating Toil."

Clog 1: The concepts are documented. How to run the template is not. This is a drain on capacity-development resources and a subsidy for sovereign advisors.

The concepts behind the LIC-DSF are taught, and taught well. The 2018 Guidance Note² and the online course³ lay out the intuition and the logic, the shocks, the thresholds, the ratings, about as well as prose can. The Interactive Guide⁴ goes further and walks a newcomer through the workbook itself. In an hour you come away with the shape of the exercise and its key inputs. For a framework this involved, that is real and useful work.

What was never written is the comprehensive manual for operating the template. Useful as it is, the Interactive Guide points past itself for the rest: its own Welcome page tells users to download a separate, fuller document, the “*User’s Manual for the LIC DSF Template*”, before they begin. We went looking. A good-faith public search turned up no such document.⁵ That manual, the one the institutions’ own guide points to, was, as far as we can tell, never written.

An unwritten manual moves the teaching to its most expensive channel. The concepts you can learn from the public materials. How to operate the template you learn from someone who already knows it. So a team gets there three ways. First, in-person help from the same scarce Fund and Bank economists the reforms mean to free for policy. Second, a hired advisor who has learned the workbook’s quirks. Third, reverse-engineering the cells when a number moves and no note explains why. Each route spends expert time, team by team, cycle after cycle. An opaque template also sustains a market for advisors whose edge is knowing it. The teams run the template with care; what fails them is the tool they were handed.

That cost runs in a chain, from the missing manual to a bill paid every cycle.

²IMF and World Bank, *Guidance Note on the Bank-Fund Debt Sustainability Framework for Low-Income Countries* (2018). <https://www.imf.org/en/Publications/Policy-Papers/Issues/2018/02/14/pp122617-guidance-note-on-lic-dsf>


³IMF and World Bank, *Debt Sustainability Framework for Low-Income Countries* (IMFx online course, edX). <https://www.edx.org/learn/sustainability/the-international-monetary-fund-debt-sustainability-framework-for-low-income-countries>

⁴World Bank and IMF, *Interactive Guide on LIC-DSF* (bundled with the LIC-DSF template download). https://www.imf.org/external/pubs/ft/dsa/tool/LIC_DSF_Download.zip

⁵The Interactive Guide’s Welcome page directs users to a *User’s Manual for the LIC DSF Template*; a good-faith public search (2026) located no such document.

Missing documentation is a standing tax.

The concepts are documented; how to run the template is not, so learning it is a recurring cost, team after team.



Teal Insights analysis.

TEAL INSIGHTS

Figure 1. The tax accrues team by team, cycle after cycle; one document, written once, clears it.

The gap turns the review's own goals against it. The review sets out to keep the framework “comprehensive, transparent, and easy to use” and to limit “capacity development needs.”⁶ Documentation that stops short of operation delivers the reverse: it manufactures the very dependency the review means to shrink. Every reform then adds inputs and steps that no comprehensive manual explains. The institutions already fund capacity development at scale. The scarce resource a missing manual burns is expert time, spent transferring the same operating knowledge by hand.

The cheaper path is the manual that was promised, written once and written to a standard. One document, written once, serves all ~70 teams every cycle, for far less than teaching the same conventions in person, year after year. Written to a standard, the page a person reads is one a tool can read too, so the mechanics need not be taught by hand at all.

Clog 2: A credit assessment this consequential must be testable. Run as undocumented software, it cannot be.

Excel earned its place for good reasons. Excel is installed in every ministry of finance, carries no license cost, keeps sensitive debt data inside the building, and runs on skills officials already have. None of that is in dispute.

Today's Excel template is a Frankenstein's monster, and that is the easy problem. It has grown, tab by tab, into nearly 200,000 formula cells across 86 sheets.⁷ Some

~200K

formula cells across 86 sheets in today's template

⁶IMF and World Bank, *LIC-DSF Review Background Note* (2026), para 7, printed p.11.

⁷Measured on our archived copy of the template (LIC-DSF IDA21, 2025-08-12): 86 worksheets (45 hidden) and 199,437 formula cells, reproducible from the public repository (Teal Insights, *LIC-DSF complexity* repo, <https://github.com/teal-insights/lic-dsf-complexity>). The file measured is the Bank-Fund LIC-DSF template, published by the World Bank on its Debt and Fiscal Risks Toolkit (<https://www.worldbank.org/en/programs/debt-toolkit/dsf>); direct link: <https://thedocs.worldbank.org/en/doc/f0ade6bcf85b6f98dbeb2c39a2b7770c-0360012025/original/LIC-DSF-IDA21-Template-08-12-2025-vf.xlsm>

tabs are vestigial, alive only because no one dares delete them. No one set out to build it this way. It accreted, one well-meant addition at a time, until the structure itself became a source of cost and risk. The reforms now propose to add more. This problem has an obvious answer: rebuild it clean, prune the dead tabs, follow established spreadsheet best practice. That is the easy part.

Best practice is a pipeline. The LIC-DSF is a knot.

A well-built model flows one way, inputs to outputs, so you can build and check it a step at a time. In the LIC-DSF, the sheets reference each other in loops, so the logic doubles back on itself.

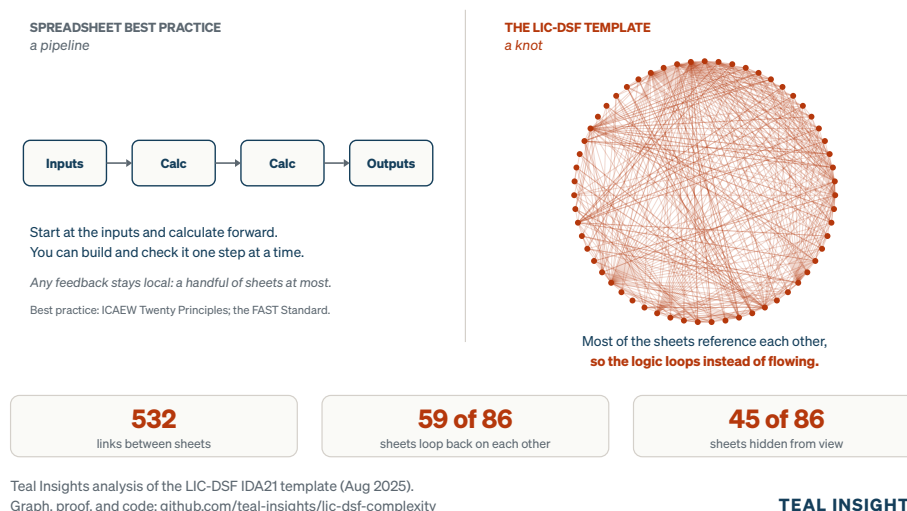


Figure 2. The LIC-DSF's real sheet-to-sheet dependency graph, measured and reproducible from the public repository.

The deeper trouble is not the implementation, it is structural. Excel is a fine instrument for what it was built to do: letting a person play with numbers and watch what happens. A workbook this consequential was never meant to run as software. The LIC-DSF is a core pillar of the international financial architecture, and it plays many roles; we focus on two. First, it helps country authorities navigate the debt-sustainability consequences of their policy choices. Second, the World Bank and the IMF use it as a credit assessment before risking shareholders' capital to lend. A tool of that gravity must be testable, auditable, and built to software-engineering standards. A workbook written to no standard cannot be. Today Excel is the framework's standard, and an undocumented one: **its inputs labeled but not precisely defined**, its logic traceable only through the formulas. Run it as software and the workbook dumps onto the user the work and risk that purpose-built software absorbs: it cannot be systematically tested, and it fails silently. Those two faults impose a hidden tax and a real operational risk, and neither washes out with a fresh coat of paint.

Nobody can audit the logic, so nobody can prove the numbers are right

Excel is untestable at scale. Here is one cell from the external-debt rating path: the borrowing-weighted average interest rate on a country's new external debt, each creditor's rate weighted by how much is borrowed from it.

Can you tell what this formula does?
Can you tell if there is a typo in it?

The same calculation, two ways. Only one can be checked.

Both find the average interest rate on a country's new external borrowing.

WHAT THE SPREADSHEET STORES

one cell - Ext_Debt_Data!F131

```
=IF(F$122=0, "",
  (SUMPRODUCT('Input 4 - External Financing'!$F$10:$F$19, Ext_Debt_Data!F$72:F$81)
  +SUMPRODUCT('Input 4 - External Financing'!$F$21:$F$23, Ext_Debt_Data!F$83:F$85)
  +SUMPRODUCT('Input 4 - External Financing'!$F$26:$F$30, Ext_Debt_Data!F$88:F$92)
  +SUMPRODUCT('Input 4 - External Financing'!$F$32:$F$36, Ext_Debt_Data!F$94:F$98)
  +SUMPRODUCT('Input 4 - External Financing'!$F$38:$F$42, Ext_Debt_Data!F$100:F$104)
  +SUMPRODUCT('Input 4 - External Financing'!$F$49:$F$51, Ext_Debt_Data!F$107:F$109)
  +SUMPRODUCT('Input 4 - External Financing'!$F$54:$F$56, Ext_Debt_Data!F$111:F$113)
  +SUMPRODUCT('Input 4 - External Financing'!$F$59:$F$61, Ext_Debt_Data!F$116:F$118))
/Ext_Debt_Data!F$122*100)
```

Can you tell what it computes? Could you spot a typo?

THE SAME LOGIC, NAMED AND TESTED

an illustrative function in plain Python

```
def avg_rate(rate, borrowed):
    if sum(borrowed) == 0:
        return None # no new debt, no rate
    paid = sum(r * b for r, b in zip(rate, borrowed))
    return paid / sum(borrowed)

# Plain-English tests the computer rechecks on every run:
assert avg_rate([1.25, 7.25], [500, 500]) == 4.25
# borrow equally -> the rate sits halfway
assert avg_rate([1.25, 7.25], [1000, 0]) == 1.25
# zero from one creditor -> it cannot move the average
assert avg_rate([1.25, 7.25], [0, 0]) is None
# no new borrowing -> no rate to report
```

Each test names a truth the computer rechecks for you.

690

characters in this one cell

0

checks Excel can run on it

3

plain-English tests, below

Source: LIC-DSF IDA21 template, cell Ext_Debt_Data!F131 (Aug 2025). Python: Teal Insights.

TEAL INSIGHTS

Figure 3. One cell from the external-debt rating path, and the same calculation as a named, tested function.

Checking that one cell takes patience, sharp eyes, and time, and it is one cell among nearly 200,000. You cannot even tell which formulas matter. Further down the same sheet sits a near-identical formula, 1,529 characters this time, that feeds nothing but a hidden export tab.⁸ By eye it is indistinguishable from the one that helps set the rating.

You do not have to write code to read the alternative. The function beside it is Python, but any scripted language would do. You can read it, follow its logic, and check it. The cell allows none of that.

The tests are the payload. Each comment states, in plain words, something that must be true, and the computer checks it on every run. A creditor you did not borrow from cannot move the average. A year with no new borrowing has no rate to report. Feed the function inputs whose answer you already know, and it passes or fails. The Excel

⁸The 1,529-character cell is Ext_Debt_Data!F410 ("Average discount rate"), a creditor-weighted average that returns a constant and is read only by a hidden export tab, off the rating path. Measured on the same archived copy.

cells carry no such checks, and nothing turns red when one breaks. The careful people who built it have no way to prove they got it right, and no one else can confirm it.

The odds are high that the template contains errors, and this is a statement about arithmetic, not anyone's competence. Audits find errors in at least 86% of real-world spreadsheets. The peer-reviewed audit evidence puts the error rate near 1% of formulas.⁹ Apply that rate to nearly 200,000 formula cells, and you expect roughly 2,000 that carry an error. Copied formulas mean errors cluster rather than scatter, so quibble with the count; even the most conservative rate in the literature leaves close to 800, and an untested workbook this size carries errors no one can see. The most effective check anyone has tested is exhausting: a team reading the workbook cell by cell. Even that misses about one error in five, and a lone reviewer misses half.

The hidden cost of an unauditable model lands hardest on country authorities. Sit in the chair of an economist at a finance ministry whose country's external risk rating moved from moderate to high. Senior management wants to know what drove the shift. The answer hides seventeen formulas deep, on six tabs, two of them hidden, in cells like the formula above, with no audit trail to mark the path. You have three options, all bad. Option 1: reverse-engineer the spreadsheet cell by cell yourself, which can devour days you owe to debt strategy or the budget. Option 2: wait for the Bank or the Fund to fly in. Option 3: hire a sovereign financial advisor, who already holds the knowledge. The upfront cost of Excel is low; the hidden cost of using it as the LIC-DSF's undocumented standard is high.

The black-box worry is real, and the standard is its answer. Capable people at the Fund, at the Bank, and in the ministries read Excel fluently but not Python, so moving the logic into code could, on its face, trade one black box for another. But auditability never came from every user reading the source. It comes from documentation, tests, and independent validation to a standard, the way supervisors already make banks manage model risk. By that test, today's 200,000-cell workbook is the black box. Write the framework down to a standard, and a reference build can test the Excel itself, so the people who keep Excel get a safer Excel. Excel stays a valid option, no one has to leave it, and the path beyond it is steep and multi-year.

Excel fails silently, propagating operational risk

Nothing flashes red when an Excel workbook breaks.

Did every input update this round, or is last cycle's number still sitting in a cell?

Did someone, two years ago, paste a value over a formula?

You changed one assumption between review rounds: can you see what moved, and the reasoning behind it?

~2,000

error-bearing cells expected, at the audit literature's ~1% formula error rate

⁹Field audits find errors in at least 86 percent of operational spreadsheets. Powell, Baker and Lawson's peer-reviewed audit of 50 completed operational spreadsheets concludes that around 1 percent of all formulas are in error (their measured averages run 0.9 to 1.8 percent, by error definition); the lowest rate in the audit literature, about 0.4 percent, still yields close to 800 here. These rates are measured on completed, operational financial and economic workbooks, copied formulas included, and the audits count how errors cluster (about ten affected cells per mistake), so extrapolating by cell rate reflects how real spreadsheets are built. Cell-by-cell inspection is the only check demonstrated effective, and it still misses about one error in five, a lone reviewer about half. Powell, Baker and Lawson, "Errors in Operational Spreadsheets," *Journal of Organizational and End User Computing* 21, no. 3 (2009); Panko, "Spreadsheet Errors: What We Know" (EuSpRIG 2000), <https://arxiv.org/abs/0802.3457>

Often there is no way to know, and that gap taxes country authorities and Fund and Bank teams alike. Well-designed software carries that burden itself, through testing. Excel hands it to the user. The people running these files take heroic care, but heroics are not a control. This is about the system, not the users.

The framework has human controls, and they are real: country teams cross-check, reviewers comment, the Boards see the output. Years of that work build a keen sense for a number that looks off. But a sense is not a test. A test is systematic: it runs every cycle, checks an answer known in advance, and cannot be skipped. Review can inspect the numbers; it cannot exercise the machine that produced them. The workbook is reviewed. It is not tested.

Bank regulators and chief risk officers would blanch at using Excel for a credit assessment this consequential. Supervisors hold the models behind a bank's lending to a clear discipline: document them, test them, validate them independently, and control how they change. The US Federal Reserve and the Bank of England now expect it of banks. The Basel Committee treats a failure of the systems a firm depends on as operational risk in its own right. The LIC-DSF is the credit assessment the Bank and the Fund run before they lend, and when a country cannot pay, the framework anchors a restructuring often worth billions. A credit assessment is a model, whatever software it runs on. This one cannot be tested, and its changes cannot be meaningfully diffed or controlled. So by the supervisors' own test, the LIC-DSF is a high-materiality model held below the bar private banks must clear for far smaller decisions. None of this binds the Fund or the Bank. Yet both assess how well countries supervise their own banks, through the Financial Sector Assessment Program (FSAP), among others, so it is a discipline they already advance.¹⁰

Supervisors make banks test and validate far smaller models. The LIC-DSF gets neither.

EUSPRIG EXCEL HORROR STORIES

Excel has been ubiquitous in economics and finance for decades, but since the global financial crisis its frailty in high-stakes work has drawn growing scrutiny. The European Spreadsheet Risks Interest Group (EuSpRIG) catalogs more than a hundred [Excel horror stories](#). Two for flavor.

JPMorgan's "London Whale" (2012). A risk model run across a chain of Excel spreadsheets divided by a sum where it should have divided by an average, understating volatility by about half. The trading book it was meant to police lost [about 6 billion dollars](#).

The UK Office for Budget Responsibility (2023). A single spreadsheet error in the fiscal watchdog's own monthly figures overstated public-sector borrowing by [9.6 billion pounds](#), from mishandling a transfer between the Treasury and the Bank of England.

¹⁰The recognized supervisory standards for model and operational risk: US interagency guidance SR 26-2, *Revised Guidance on Model Risk Management* (2026); the Bank of England PRA's SS1/23, *Model Risk Management Principles for Banks* (in force since 2024); and the Basel Committee's *Principles for the Sound Management of Operational Risk* (rev. 2021).

Every one of these was a well-resourced institution with more oversight than most debt offices can muster, and a spreadsheet still bit it. The LIC-DSF runs the same kind of model, on the same tool, for higher stakes than most. A credit assessment that cannot be tested, audited, or version-controlled is the opposite of future-proof, and new features bolted onto it will not change that. The durable path moves the calculation onto a foundation that can be tested: the workbook rebuilt clean now against a written standard, and the ground laid for what comes after Excel. Excel stops being the standard and becomes a tool built to one, carrying the testing the careful people now do by hand, and handing those hours back for judgment.

Clog 3: The long-term module is the reform's most promising piece, and it stays blocked until two pipes are cleared: the models' calibration, and the connector to the template.

The module closes an asymmetry in the framework. Until now, a country that invested to defend against a macrocritical climate or development risk got no credit for it: the spending worsened its debt metrics, while the benefit went unmodeled. Picture an agricultural economy, as many low-income countries are, hit by deeper and more frequent droughts that knock points off GDP. The module lets that country weigh both sides of a large irrigation program in one frame: the fiscal cost against a smaller drought hit to growth and revenue.

The module sizes up a resilience investment: its macroeconomic payoff against what it costs to finance. A country negotiating a Fund program is usually told to tighten its belt. Inside the LIC-DSF, rather than in a side analysis, the country can test whether a well-targeted investment improves long-term debt sustainability or adds to the bill. The answer is assessed, not assumed.

Ratings dominate the conversation, and they cut both ways: a high risk rating can raise a country's cost of borrowing, or, for some, unlock the grants that lower it. The more durable value sits underneath the headline.

At its most useful, the LIC-DSF is a decision-support tool, and the module is where that shows. That durable value is the analysis the framework runs: projecting the debt path, stress-testing assumptions, weighing financing options, and finding the policy moves that keep debt sustainable.

And the module sits inside the Fund's own mandate. The Fund's surveillance guidance counts a risk as macrocritical when it "significantly influences present or prospective balance of payments or domestic stability."¹¹ A debt-defining exposure to drought, flood, or a development gap is exactly that.

Two design choices deserve credit. First, the approach is problem-agnostic and model-agnostic: an interface to quantitative tools, not a bolt-on to one model or one challenge. Second, the input surface is small. The module's inputs are expected to

¹¹IMF, *Guidance Note for Surveillance under Article IV Consultations* (2022), Box 2. <https://www.imf.org/-/media/files/publications/pp/2022/english/ppea2022029.pdf>

resemble the simplified “customized scenario” the template already carries, the same template available on the World Bank’s website today.¹²

What follows are not reasons to stop. They are two clogs in the plumbing, and both are cheap to clear. We trace them through the most-cited pathway, a model like DIGNAD feeding a customized scenario. The lesson carries to any model the framework accepts.

The model only answers well if its dials are set right, and there is no shared, documented way to know what “right” is.

Climate-macro models are only as good as their calibration, and that calibration is undocumented. These models ask a country team to set structural parameters that swing the answer. The published sources set them with benchmarks, judgment, and open admissions of thin data. DIGNAD’s own paper puts the dial that most directly governs whether resilience pays, the return on adaptation investment, at 50% while conceding that estimates are “much less abundant.”¹³ So a team runs the climate case by setting the model’s most consequential number with a finger in the air, and nothing written down tells the next team whether that number was researched or guessed. It is guild knowledge again, now in the model’s dials: inherited by hand, written nowhere.

The binding constraint here is capacity, not math. Maintaining sophisticated country-level models for every assessment is beyond what most teams can sustain, which is part of why these tools are not yet used systematically in the LIC-DSF.¹⁴ Some calibration choices are legitimately private, built on confidential country data, and those should stay private. Most are not. They are merely undocumented. The shareable part can be written down: the benchmark values, the ranges, the reasoning behind a given setting. That is a documentation problem, not a research problem.

The plug and the socket are not the same shape.

Plug-and-play is the institutions’ own aspiration, and the right one. The reform materials describe a “Plug & Play interface” drawing on the best climate modeling of the day, naming the World Bank’s CCDR (its country report series) and MFMoD-CC, and the Fund’s Q-CRAFT and DIGNAD.¹⁵ That is the right instinct: models change and country priorities change, and an interface that accepts new tools keeps the framework from locking itself to one.

The connection point is tractable. From conversations with staff, we understand the reformed module will not ask country teams to populate the framework’s full input

¹²IMF and World Bank, *LIC-DSF template* (LIC-DSF IDA21, version 2025-08-12), published via the World Bank Debt and Fiscal Risks Toolkit (<https://www.worldbank.org/en/programs/debt-toolkit/dsf>); direct file: <https://thedocs.worldbank.org/en/doc/f0ade6bcf85b6f98dbeb2c39a2b7770c-0360012025/original/LIC-DSF-IDA21-Template-08-12-2025-vf.xlsm>

¹³Marto, Papageorgiou and Klyuev, *Building Resilience to Natural Disasters: An Application to Small Developing States* (DIGNAD), IMF WP/17/223 (2017), printed p.13. <https://www.imf.org/en/Publications/WP/Issues/2017/10/03/Building-Resilience-to-Natural-Disasters-45295>

¹⁴Ranger, Pasqua and Adam, *Integrating Nature into the IMF-World Bank’s Debt Sustainability Framework for Low Income Countries* (2025): calibrating and maintaining country-level macro-models for the DSF is “beyond current operational capacity,” and DIGNAD “is not yet systematically used within DSAs including the LIC-DSF” (printed p.54).


¹⁵IMF and World Bank, *LIC-DSF Review: Staff Reform Proposals* (presentation, 2026), p.29: “Plug & Play interface... Integration of climate modelling, including from WB (CCDR, MFMoD CC) and IMF (Q-CRAFT and DIGNAD).”

set. It will ask for something close to the template's existing "customized scenario," an interface of 13 inputs entered as changes from a baseline over 20 years.¹⁶ That turns a daunting data exercise into a manageable one.¹⁷ The catch is narrower: of the 13 inputs, exactly one, real GDP growth, reads straight off a DIGNAD run. The other 12 must be built or assumed by hand.¹⁸

DIGNAD is the right place to start: the Fund already uses it, and the reform materials name it. The gap is units and definitions. DIGNAD has already been run in several Fund country programs. But its outputs and the template's inputs do not line up. DIGNAD reports 47 variables, many as deviations from a steady-state path, in its own units.¹⁹ The customized scenario wants something else: nominal levels on national-accounts definitions (revenue and grants, primary expenditure, exports and imports), plus price and exchange-rate series like the GDP deflator and the currency's depreciation, which DIGNAD does not report at all. So the other 12 must be translated, and no published document spells out how.

The plug and the socket are not the same shape.

The customized scenario needs 13 inputs. Only real GDP growth reads straight off a DIGNAD run.



Only 1 of 13 plugs straight in.

The other 12 need standard wiring, not a new model.

The other 12: built or assumed by hand

- Revenue and grants
- Primary expenditure
- Grants
- Public sector assets
- Inflation (GDP deflator)
- Nominal FX depreciation
- Exports
- Imports
- Net transfers, official
- Net transfers, private
- Net FDI
- GDP deflator, USD terms

Source: LIC-DSF IDA21 template, Customized Scenario - public sheet (Aug 2025). Mapping: Teal Insights.

TEAL INSIGHTS

Figure 4. Of the 13 customized-scenario inputs, only real GDP growth reads straight off a DIGNAD run.

¹⁶These 13 are the macro inputs for both the public-debt and the external-debt customized scenarios; six of them (exports, imports, the two net-current-transfer lines, net FDI, and the US-dollar GDP deflator) are the external scenario's variables. The separate "Customized Scenario-External" tab adds no macro inputs: it re-displays that external subset and adds external-debt financing assumptions (the average interest rate on new debt, the US-dollar discount rate, grace period, maturity) and computed debt-dynamics results. Those financing terms are framework inputs, not model outputs, so they sit outside this connection question.

¹⁷One caveat: the reformed template is expected around October 2026 and is not public yet, so our confidence that it resembles the customized scenario rests on those staff conversations, not a published spec. Helping shape it is part of why we are filing this response.

¹⁸Verified by reading the blank LIC-DSF template's "Customized Scenario" sheets against the DIGNAD toolkit output: of the 13 delta inputs, only real GDP growth maps cleanly from a DIGNAD run.

¹⁹DIGNAD Toolkit and 2023 User Manual (Aligishiev et al., IMF TNM/2023/03): the toolkit's variable catalog runs to 47 selectable variables, many reported as deviations from a steady-state baseline. <https://www.imf.org/en/Publications/TNM/Issues/2023/05/29/DIGNAD-Toolkit-533416>

What plug-and-play needs is a universal travel adapter, not a one-off wire. DIGNAD is the stated first target, so start there. But the proposal names several models and promises to accept more. A connection soldered once for DIGNAD, in one team's spreadsheet, fits only that plug, so every team solders it again, the same hand-wiring paid 70 times, the toil the framework could retire once. To be plug-and-play across models, the framework needs the standard a traveler carries: one documented adapter from any model's outputs to the template's inputs. Each of the remaining 12 inputs has an established way to be built or assumed from a model's output. So wiring all 12 is finite, tractable work: serious effort and careful thought, but standard practice, not new research. It is the step still missing.

A model-agnostic module earns that name only if a team can actually calibrate the model and connect its outputs, and today most teams can do neither at scale. The institutions already named the goal: a plug-and-play interface drawing on the best models of the day. Writing down the calibration guidance and the connector, once, to a shared standard, is what lets them keep their own word. Both fixes are cheap, boring, and worth starting now.

Write it down to a standard, so the tool carries the toil and the people do the policy

The reforms set the right goals: a framework that is plug-and-play, future-proof, and runnable by the lean Fund teams and finance ministries who operate it across ~70 countries. None of that is reachable while the framework's inputs, logic, and connections live as guild knowledge, held in the formulas and a handful of experts' heads and written to no standard. Leave it unwritten and the toil never leaves the teams: each re-derives the same mechanics by hand, every cycle. Write it down to a standard, and a tool can carry the toil, freeing the scarce economists on every side for the policy only they can do. Then Excel keeps working and turns testable, other tools become possible, the numbers can be checked, the models finally connect, and the rating engine becomes a tool a country can run itself.

Writing it down means writing it to a standard, not to prose alone. A standard is an agreed interface: names, units, and formats fixed once so both sides of a handoff match. To a standard means three things, written precisely enough for a machine to read:

1. **The inputs.** Every field named, with its units, currency, and permitted values.
2. **The logic.** Every step from input to output, specified so anyone could rebuild it.
3. **The connections.** How each accepted model's outputs map to the template's inputs.

Done is one test, and it doubles as the promise: write it down once, and a person, a machine, and an AI can all read the same page and return the same number. One team can hand-wire its own workbook, but bespoke wiring never generalizes. Across

~70 teams that is the same hand-wiring paid 70 times, the toil the framework could retire once.

WHY STANDARDIZE


The economics of standardizing are compelling. The economics of not standardizing are sobering. A standard is the difference between a guild and an assembly line.

Every economist already trusts this idea, filed under mass production: writing the part down to a standard is what turned the craft guild into the assembly line. The framework is still a guild, each team building its copy by hand. Write down precisely what each input is, how the pieces relate, and the transformations between them, and anyone can build tools on it. A tool absorbs the toil, so the scarce experts spend less time on mechanics and more on policy.

The training does not vanish. What changes is what it buys. Today it buys a walkthrough of the template. Once a tool carries the mechanics, the same spend buys engagement on the policy questions only people can settle.

The standard is the first step; the tools do the rest.

Write down each input and transformation once, and the rote work can be built away.



Teal Insights analysis.

TEAL INSIGHTS

Figure 5. A written standard lets anyone build tools on it, and the tools absorb the rote work, so scarce expert time goes to policy.


Write the framework down to a standard, and you get more than tidier plumbing. You get a different tool. The manual rework each team now repeats becomes work the tool carries once. The week of emails to force one model's numbers into the template becomes a connector written once. Excel keeps working and, for the first time, becomes testable. And the framework stops being only a rating that Fund teams produce and the Bank reviews. It becomes a decision-support tool a finance ministry can run itself, to weigh a real policy choice with the drivers of the debt path in view. Today only a handful of users manage that, and even for them it is cumbersome. An

open standard makes country-led scenario analysis easy for everyone, and it helps most where capacity is thinnest.

That is also what earns the tool trust: the institutions accept it, and the countries that use it can see and control the assumptions that feed it. A written standard is what makes both possible.

One written standard, and everything finally connects.

Write it down once, and a person, a machine, and an AI read the same page and return the same number.



Illustrative schematic. Teal Insights.

TEAL INSIGHTS

Figure 6. Write the framework to one standard, and many inputs plug in, every tool reads it, and the debt path reads out.

These fixes are a down payment on future-proofing, and their price is small against the recurring bill. The heaviest piece, the clean rebuild with its documentation, is a couple of early- to mid-career economists for a year; the rest is lighter still. Size it from the Fund's own number: an early- to mid-career economist (grades A11-A13) earns about \$155,000 at the band midpoint, and the fully loaded cost runs higher.²⁰ Call the heavy piece a one-time million-dollar effort, order of magnitude. Set it against the more than half a billion the Fund spent on capacity development in FY2025, about 31% of what it spends to run itself,²¹ part of it the hands-on, one-on-one country support a written framework would start to shrink. The point is the mismatch: a one-time million against a recurring half-billion, aimed at the root. A written tool teaches everyone, once. Measure twice, cut once.

The sequence is simple, the cheap parts start now, and four fixes carry it: one per clog, plus the body that keeps them current. They are (1) a data dictionary for the inputs, (2) a clean rebuild and a written logic spec, (3) a model-to-template connector and a shared parameter library, and (4) a semiannual operability forum that owns the dictionary, the connector standard, and the parameter library as living documents. The first dictionary, the first crosswalk, a shared parameter sheet, and the first forum session need a mandate, not a budget, and can land by the Annual Meetings. The

~\$1M, once

an order-of-magnitude estimate to fix the root, against half a billion dollars a year in capacity-development spending

²⁰IMF Annual Report 2025, Web Table 3.3 (staff salary structure, effective May 1, 2025): grades A11-A13 (early- to mid-career economists) carry published salaries of roughly \$120,000 to \$200,000, about \$155,000 at the band midpoint (grade midpoints ~\$137k/\$155k/\$177k for A11/A12/A13). Salary only; the fully loaded cost (benefits, pension) is higher. <https://www.imf.org/external/pubs/ft/ar/2025/pdfs/imf-annual-report-2025-web-tables.pdf>

²¹IMF, FY2025 Budget Outturn (PPEA2025027), printed p.16: overall capacity-development spending including Fund-wide overheads totaled \$546 million, 31 percent of gross administrative spending. <https://www.imf.org/en/-/media/files/publications/pp/2025/english/ppea2025027.pdf>

clean rebuild follows the reformed template's expected October release and runs about a year. This is how you widen the pipe before the new machinery goes in.

Fix 1: publish the data dictionary, in names a human and a machine can both read

Clog 1 is a manual that was promised and never written: the framework is taught well, but how to run the Excel template is not. Write it, and write it to a standard.

Many of the template's inputs are ambiguous to anyone who did not build it. Put yourself in the seat the template assumes: you are new at a ministry of finance, and the sheet asks for "Public Investment, constant prices". You open the national accounts: total investment, no government split. You open the government finance statistics: the government's investment spending, on a definition that does not quite match. The series the template wants exists in neither. Every experienced hand knows the construction: the fiscal line is the public piece, the total minus it is the private piece, each deflated to constant prices. The Fund teaches it in its own financial-programming course. Nothing near the template says so. Guild knowledge: obvious once told, invisible until then, and exactly what a machine-readable dictionary would encode.²²

The fix is not more prose. Prose is what we have, and prose is what leaves each field ambiguous. The fix is a structured data dictionary built to the metadata standards already written for this. The Fund and the Bank both sponsor SDMX, the international standard for naming statistical series with their units and definitions, and the Fund already publishes data under it, so this is adoption, not invention. Tag each input with a standardized code, SDMX where the series exists and a defined term where it does not, recorded to a metadata standard that fixes what each element is, its units, and its permitted values.²³ Done once, the ambiguity is gone for three readers at once: a person knows what to enter, deterministic software can read it, and AI can navigate it. The next decade of this work will lean on software and AI alike.

The template is already halfway there: its data sheet ships a series code beside every macro input, wired to the Fund's internal databases, which makes filling the sheet easy for a Fund desk. The codes are assembled from published pieces (the Fund's country numbers, WEO-style mnemonics), but the assembled identifiers work only in the Fund's internal tools; no public service resolves them. The open version exists: the Fund publishes these same indicators through its public SDMX services. Point the

²²The ingredients already have machine-readable identities: gross fixed capital formation is transaction P51G in the SDMX national-accounts codelists the IMF, Eurostat, and OECD share; government investment in nonfinancial assets is item 31 (fixed assets, 311) of the GFSM 2014 Statement of Operations, code G31 in the Fund's public SDMX dataflow (net of disposals where GFCF is gross, one of the conventions a dictionary would state). A dictionary entry could say, machine-readably: public investment = G31; private investment = P51G minus G31; deflate with the investment deflator. Countries in the IMF's e-GDDS already publish national data in SDMX through their National Summary Data Pages, so a template keyed to these codes could draw on what countries already report. The construction is taught in IMF, *Financial Programming and Policies* (FPP:1x/2x, edX).

²³The relevant standards already exist. For statistical series, SDMX (ISO 17369), which also maintains controlled vocabularies (code lists) for common concepts that a dictionary can adopt directly. For the structure of a data element, its definition, units, and permissible values, ISO/IEC 11179 (Metadata Registries). For a machine-readable dictionary, an open schema such as Frictionless Table Schema, with DDI (Data Documentation Initiative) as the heavier official-statistics alternative. We take no view on which a team adopts; we ask only that it conform to a recognized one.

plumbing there, and the ease Fund desks already enjoy extends to every ministry and every tool the framework serves.²⁴

This is the root of it: running the framework takes scarce expert judgment largely because so much of it is unwritten. Much of what the template now demands is not real judgment but undocumented convention someone already settled. Write it down, and scarce judgment goes where the review wants it: the policy questions no document can yet answer.

The ask: publish a data dictionary and calculation guide, every input tagged to a standardized name, unit, and definition under a recognized metadata standard, and shipped machine-readable. A small documentation team owns it, in the Strategy, Policy, and Review Department (SPR) on the Fund side and the Debt unit on the Bank side, working from the live template. Done is when a person, deterministic software, and an AI can all read the same dictionary.

Fix 2: rebuild the workbook clean, and start laying the ground beyond it

Excel will always be on the menu. Writing the framework down keeps it there, makes it testable, and opens the menu to new tools that offer what Excel cannot.

We understand from conversations with staff that the reformed template will be a clean rebuild rather than another layer on the old one. Good: that is half of it. A clean rebuild and its written documentation are one job, not two. Clog 2 showed why: today's workbook cannot be tested or audited. Rebuild it to a published spreadsheet standard (ICAEW and FAST exist for exactly this), with inputs, calculations, and outputs separated and the dead tabs retired. Which standard matters far less than landing on one; this is a coordination problem, like agreeing which side of the road to drive on, where convergence is the prize.²⁵ Document each line in the same standardized names as the dictionary, and the manual writes itself. Left for afterward, it never gets built.

Future-proof means options, not eviction. Excel runs the framework today for real reasons, the ones from the top of Clog 2: it is installed everywhere, costs nothing at the margin, keeps the data in the building, and runs on skills officials already have. It will be here a while, and that is fine. Future-proof means laying the ground, in parallel, for tools beyond it that are testable and auditable by construction, all reading the same written standard. A path that required ~70 teams to switch tools at once would fail the framework's first test: it has to run everywhere it is used. The standard passes that test. The people who keep Excel get an Excel that a reference build can finally check; everyone else gets options. **And building those tools costs less than it did.**

²⁴The data sheet's own instruction reads "To download from DMX or EcXL", the Fund's internal data tools. The codes concatenate the Fund's numeric country code with a series mnemonic (652NGDP_R is Ghana's constant-price GDP, the template's illustrative country). The mnemonics are cousins of the published WEO codes, and the Fund serves the same indicator publicly over SDMX (dataflow IMF.RES:WEO, key GHA.NGDP_R.A); the internal string resolves to nothing there.

²⁵The caveat: a technical standard can lock in a bad choice (QWERTY is the stock case), so "any standard beats none" holds most cleanly for spreadsheet discipline, where any disciplined structure beats an idiosyncratic one.

Agentic AI coding tools have already made well-tested, deterministic software much cheaper to build and maintain. The tool stays deterministic and testable; the AI only helps build it.

The ask: treat the rebuild and its documentation as one job, rebuild to a published spreadsheet-engineering standard (FAST or ICAEW), and publish the logic as an open standard others can build on. This is the funded piece, the couple of economists for a year, and it can follow the October release rather than block it. Done is a workbook that conforms to the standard and a logic spec anyone can re-implement.

Fix 3: build the connector and a shared parameter library, starting now

Clog 3's two missing pipes, the long-term module's calibration and its output connector, have concrete starting points, and both can begin this year. First, the connector. Publish and maintain a single agreed crosswalk from DIGNAD's outputs to the inputs the new template needs, so no team rebuilds it by hand in its own spreadsheet. Use it as the pattern for the next model, then the next; the World Bank can do the same for the models behind its CCDR. The forum below is where "here is how you connect a model" gets decided and written down.

Second, stop sourcing calibration defaults by email. Today they travel inbox to inbox and retire with their owners. Stand up a shared library of the parameters these models need, and let country teams fill it as they run them. Record two things for each: the number, and how it was reached, whether a quick benchmark, a regional average, a figure from a paper, or a serious estimation. Sensitive country data stays out. Over time the library answers what every team now works out cold: when local data is thin, what is a reasonable value, and which parameters even move the results. Many do not, and today each team relearns which.

The ask: publish and maintain a model-to-template connector, starting with DIGNAD, and a shared, machine-readable parameter library that country teams build as they go. The Fund owns the DIGNAD crosswalk, the World Bank its own models' crosswalks; a first version of both can start as a shared sheet at the department level and grow machine-readable from there. Done is when no team solders the connection by hand.

Stop sourcing calibration defaults by email. Today they travel inbox to inbox and retire with their owners.

Fix 4: a semiannual forum that owns the living standard

None of the first three stays fixed without somewhere to keep it current. The workbook grew to 86 sheets because nothing owned it. The forum is the owner. The Global Sovereign Debt Roundtable (GSDR) is the model: it meets twice a year and keeps a living compendium of shared understandings. The analog here, call it the LIC-DSF Operability Forum,²⁶ would meet on the same cadence and own three living documents: the data dictionary, the connector standard, and the parameter library. It is where "how you do this" is written down once, for everyone, instead of rediscover-

The workbook grew to 86 sheets because nothing owned it. The forum is the owner.

²⁶In a field this fond of acronyms, the forum will need one of its own: the LDOF.

ered team by team. It fits the Fund's existing capacity-development architecture, so it needs no new machinery.²⁷ Keeping those documents current is what keeps the framework fit for purpose as models and country needs change.

The ask: stand up a semiannual operability forum that maintains the dictionary, the connector standard, and the parameter library as living documents. Done is a forum that has met twice and three documents that are current.

A vision of the future, in miniature

This is the work we do, and it is why we can read the plumbing so closely: we build SovTech, open-source reimplementations of these tools, with finance ministries and other technical partners. You find where a framework leaks by rebuilding it. We build it for the reason this memo argues: a framework only helps the people who run it if they can actually run it.

One of those reimplementations, the Q-CRAFT Explorer, is a small working picture of what the LIC-DSF could be. It takes the Fund's long-term climate-fiscal tool and rebuilds it for the busy person who has to run it. The barrier to tools like this was never the economics. It was the ergonomics. The few parameters that matter carry their guidance right where you set them. The defaults are sensible and shown. The scenarios line up on one screen, and the numbers export into your own analysis. It is deliberately small, five parameters and one equation, to prove the design before bringing it to the framework that matters most.

When the tool carries the guidance, the ministry does more on its own. That is the payoff this memo has been about: capacity development goes to substantive policy, not spreadsheet training.

Everything we build is philanthropically funded and open source, so our only stake is that it gets used. Take it, borrow it, build something better on it. There is nothing to buy and nothing to sign. If a fuller version of this lifts the busywork off the teams who run the LIC-DSF, whoever builds it, that is the entire point.

²⁷IMF, *Review of the Fund's Capacity Development Strategy: Towards a More Flexible, Integrated, and Tailored Model* (Policy Paper No. 2024/014, April 2024), Overview, printed p.2: the review "proposes modernizing delivery modalities, including by leveraging communications and technology and through blended delivery." <https://www.imf.org/-/media/files/publications/pp/2024/english/ppea2024014.pdf>

Q-CRAFT Explorer: a small working picture of what the LIC-DSF could be.

Point and click, guidance where you set each dial, tested code underneath. Excel stays for whoever wants it; the written standard is what lets tools like this multiply.

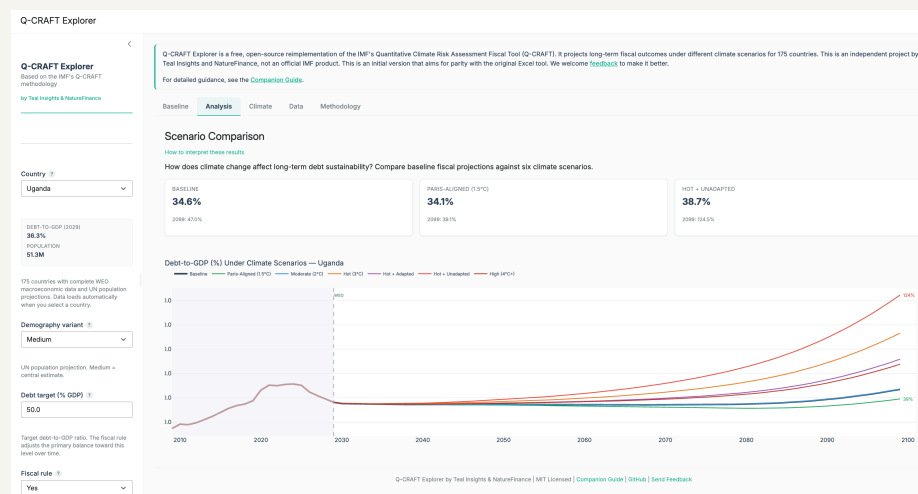


Figure 7. A deliberately small prototype, five parameters and one equation: inputs and their guidance on the left, debt-to-GDP under six climate scenarios on the right.

The fixes are boring, doable, and cheap

Weigh who gains once the toil is gone, and notice it is everyone at the table.

The Fund's and the Bank's economists hate toil. Flying out to walk a team through the template is manual, repetitive work they would trade in a heartbeat for the policy questions they were hired to answer.

The ministry's economists hate toil. Reverse-engineering a rating shift by hand is time stolen from the budget and the debt strategy.

The shareholders who fund all of it hate paying for toil. A capacity-development bill that teaches the same template one room at a time, year after year, is the kind of spend that documentation pays down. And the same fix hands them a credit assessment they can finally audit before they risk billions on it.

So write the framework down to a standard, and start the cheap parts now. We opened with a spreadsheet that is opaque and cumbersome to the people who run it. At its best it is more than that: a decision-support tool a finance ministry runs itself, to weigh a real choice like an irrigation program against a worsening drought. A few already use it this way. Not many, because running a handful of scenarios can eat a full day of copy-and-paste. Write the logic down in the open, let tools carry the toil, and that day's work shrinks to minutes. The teams already doing it get their time back,

and many more can join them. That is the prize beneath all three clogs, and it costs a fraction of the toil it retires. Clear the clogs first, and the reforms finally flow.

Sources

The footnotes carry the short, page-pinned citation at the point of use; this section lists each cited source once, in full, with its link.

IMF and World Bank primary sources

- IMF and World Bank. *LIC-DSF Review Background Note*. 2026. <https://www.imf.org/en/Publications/analytical-notes/Issues/2026/06/02/Review-of-the-Debt-Sustainability-Framework-for-Low-Income-Countries>
- IMF and World Bank. *LIC-DSF Review: Staff Reform Proposals* (presentation). 2026. <https://www.imf.org/en/Publications/analytical-notes/Issues/2026/06/02/Review-of-the-Debt-Sustainability-Framework-for-Low-Income-Countries>
- IMF and World Bank. *Guidance Note on the Bank-Fund Debt Sustainability Framework for Low-Income Countries*. 2018. <https://www.imf.org/en/Publications/Policy-Papers/Issues/2018/02/14/pp122617guidance-note-on-lic-dsf>
- IMF. *Guidance Note for Surveillance under Article IV Consultations*. 2022. <https://www.imf.org/-/media/files/publications/pp/2022/english/ppea2022029.pdf>
- IMF. *Review of the Fund's Capacity Development Strategy: Towards a More Flexible, Integrated, and Tailored Model* (Policy Paper No. 2024/014). 2024. <https://www.imf.org/-/media/files/publications/pp/2024/english/ppea2024014.pdf>
- IMF. *Government Finance Statistics Manual 2014*. 2014. <https://www.imf.org/external/pubs/ft/gfs/manual/2014/gfsfinal.pdf>
- IMF. *FY2025 Budget Outturn* (PPEA2025027). 2025. <https://www.imf.org/en/-/media/files/publications/pp/2025/english/ppea2025027.pdf>
- IMF. *Annual Report 2025, Web Tables* (Web Table 3.3, staff salary structure). 2025. <https://www.imf.org/external/pubs/ft/ar/2025/pdfs/imf-annual-report-2025-web-tables.pdf>

Teaching resources for the LIC-DSF

- IMFx. *Debt Sustainability Framework for Low-Income Countries* (edX online course). <https://www.edx.org/learn/sustainability/the-international-monetary-fund-debt-sustainability-framework-for-low-income-countries>
- IMFx. *Financial Programming and Policies, Parts 1 and 2* (FPP1x/2x, edX online courses). <https://www.edx.org/learn/finance/the-international-monetary-fund-financial-programming-and-policies-part-1-macroeconomic-accounts-analysis> (Part 2: <https://www.edx.org/learn/macroeconomics/the-international-monetary-fund-financial-programming-and-policies-part-2-program-design>)
- World Bank and IMF. *Interactive Guide on LIC-DSF* (bundled with the LIC-DSF template download). https://www.imf.org/external/pubs/ft/dsa/tool/LIC_DSF_Download.zip
- IMF and World Bank. *LIC-DSF template* (LIC-DSF IDA21, version 2025-08-12), published via the World Bank Debt and Fiscal Risks Toolkit. <https://www.worldbank.org/en/programs/debt-toolkit/dsf> (direct file: <https://thedocs.worldbank.org/en/doc/f0ade6bcf85b6f98dbeb2c39a2b7770c-0360012025/original/LIC-DSF-IDA21-Template-08-12-2025-vf.xlsm>)

Cited research

- Marto, Ricardo, Chris Papageorgiou, and Vladimir Klyuev. *Building Resilience to Natural Disasters: An Application to Small Developing States* (the DIGNAD model). IMF Working Paper WP/17/223. 2017. <https://www.imf.org/en/Publications/WP/Issues/2017/10/03/Building-Resilience-to-Natural-Disasters-45295>

- Aligishiev, Zamid, et al. *DIGNAD Toolkit* (and 2023 User Manual). IMF Technical Notes and Manuals TNM/2023/03. 2023. <https://www.imf.org/en/Publications/TNM/Issues/2023/05/29/DIGNAD-Toolkit-533416>
- Ranger, Nicola, Carlo Pasqua, and Christopher Adam. *Integrating Nature into the IMF-World Bank's Debt Sustainability Framework for Low Income Countries*. 2025. <https://www.lse.ac.uk/granthaminstitute/publication/integrating-nature-into-debt-sustainability-analysis/>
- Panko, Raymond R. "Spreadsheet Errors: What We Know." *EuSpRIG 2000* (arXiv 2008). <https://arxiv.org/abs/0802.3457>
- Powell, Stephen G., Kenneth R. Baker, and Barry Lawson. "Errors in Operational Spreadsheets." *Journal of Organizational and End User Computing* 21, no. 3 (2009). https://mba.tuck.dartmouth.edu/spreadsheet/product_pubs_files/errors.pdf
- Google. *Site Reliability Engineering: How Google Runs Production Systems*. O'Reilly, 2017. Chapter 5, "Eliminating Toil." <https://sre.google/sre-book/eliminating-toil/>

Standards

- SDMX (Statistical Data and Metadata eXchange), ISO 17369. <https://sdmx.org/>
- ISO/IEC 11179, Information technology: Metadata registries (MDR). The standard family models data elements with their definitions, units, and permissible values. <https://www.iso.org/standard/78914.html>
- Frictionless Table Schema (open data-package schema). <https://specs.frictionlessdata.io/table-schema/>
- DDI (Data Documentation Initiative), official-statistics metadata standard. <https://ddialliance.org/>
- ICAEW. *Twenty Principles for Good Spreadsheet Practice*, 4th ed. 2024. <https://www.icaew.com/technical/technology/excel-community/20-principles-for-good-spreadsheet-practice-2024-edition>
- FAST Standard (Flexible, Appropriate, Structured, Transparent), v02c. 2019. <https://fast-standard.org/the-fast-standard/>

Supervisory guidance

- Basel Committee on Banking Supervision. *Principles for the Sound Management of Operational Risk*. Revised 2021. <https://www.bis.org/bcbs/publ/d515.pdf>
- Bank of England, Prudential Regulation Authority. *SS1/23: Model Risk Management Principles for Banks*. 2023, in force since 2024. <https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-ss>
- US Federal Reserve and interagency. *SR 26-2: Revised Guidance on Model Risk Management*. 2026. <https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm>

Teal Insights

- Teal Insights. *LIC-DSF complexity* (public reproducible repository measuring the template). <https://github.com/teal-insights/lic-dsf-complexity>
- Teal Insights. *Q-CRAFT Explorer* and open-source app. <https://github.com/Teal-Insights/QCraft-App>