

RESEARCH AND INSTITUTIONAL DESIGN MONOGRAPH — CONDENSED EDITION

The Humanity Compact

A functional accountability architecture for cross-border and intergenerational effects

The Global Accountability Chain

Nizar YAICHE

Former Minister of Finance, Republic of Tunisia — independent practitioner-researcher

Condensed edition — July 2026

A concise gateway to the full research monograph (350 pp.). It presents the complete argument; the full edition contains the detailed demonstration, the evidence, the protocols and the annexes. Section references (§) point to the edition of record.

How to read this short version

This is not a summary. A summary says less; this edition preserves the whole architecture of the argument — problem, diagnosis, theory, instantiation, testing, decision — and delegates to the full edition the evidence, the protocols, the detailed objections and the annexes.

Substantive claims are anchored. A sentence that asserts carries the reference of the section that establishes it. The hurried reader can stop here; the sceptical reader traces the claim back in one step.

Nothing is loosely paraphrased. Definitions, mandates and safeguards are reproduced from the full edition rather than restated. Where a formulation differs, the full edition governs.

Epistemic discipline is preserved. Evidential statuses, the non-additivity rule and the epistemic rungs survive intact. Without them a short version becomes advocacy — which the work refuses to be.

1. The argument in one page

The problem. Monetary, digital, trade and climate decisions cross borders; accountability stops at the threshold of the jurisdiction that decides.

The diagnosis. The international system is not empty: it is dense but discontinuous. Between measuring a harm, qualifying it, drawing a consequence, financing it and opening a remedy, the evidence is lost from one link to the next.

The theory. The Global Accountability Chain states what complete accountability requires: functional completeness (every necessary function is covered), relational completeness (each is tied to an identifiable forum) and procedural completeness (evidence and uncertainty travel across the interfaces), ordered by functions and grounded in rights.

The instantiation: the 7+3. Seven functions — measure, qualify, price where applicable, adjust, finance, certify, adjudicate — and three instruments: the Carthage Charter, the Global Accountability Framework and the Call of Carthage. This is a candidate instantiation, not the theory itself, and emphatically not necessarily seven new organisations: a function may be hosted by an existing institution, a protocol or a federated network.

The testing. The chain is tested on four grounds: an integrated demonstrator (an artificial-intelligence data centre), monetary policy spillovers, development finance (Sevilla) and preventive peace.

The decision. A minimum viable pilot, graduated and non-binding adoption routes, and explicit criteria for maintaining, simplifying or rejecting the architecture in light of what it produces.

2. The landscape: twenty families, five transmission modes, one order of magnitude

Before any theory, the scale. The reference points below come exclusively from third-party sources; they are not outputs of the research programme. They measure different objects and do not add up — yet each lens, independent of the others, arrives at the same order of magnitude (§2.2).

What the lens measures	Order of magnitude	Third-party source
Fossil-fuel subsidies — including the under-pricing of environmental and health damage	≈ USD 7.4 trillion per year (2024), of which 6.7 implicit	IMF (Black et al., 2025)
Sustainable Development Goals financing gap	≈ USD 4 trillion per year	United Nations (2025a)
Illicit financial flows out of developing countries	≈ USD 1 trillion per year	Global Financial Integrity (2021); Tax Justice Network (2023); FACTI (2021)
Asymmetry in sovereign ratings (mid-2024)	68 developing countries below investment grade, against 24 within it	UNCTAD (2024)
Financial and technological mechanisms: high-frequency trading, tax havens, algorithmic bias, data extraction	Tens to hundreds of billions of dollars per year, each	BIS (2023); UNCTAD (2023); OECD (2023); ILO (2023)
What is not counted in dollars — ambient and household air pollution	≈ 7 million premature deaths per year	World Health Organization (2024)

None of these figures is a sum, and none is claimed by the work. They suffice to pose the question that everything else answers: these effects cross borders and generations; which institution measures them, qualifies them, draws a consequence, and before which forum can they be contested?

The programme models twenty families of externalities across five transmission modes (§2.3-2.4). They are neither twenty autonomous theories nor twenty additive magnitudes: they form a portfolio of phenomena to which one and the same accountability grammar applies, with different data, rights, causal mechanisms and forums.

The work does retain an aggregate exposure envelope, but subjects it to a strict non-additivity rule and keeps it out of all communication material until the underlying workbook is published and replicated. That rule applies to this short version as well: the aggregate does not appear here.

3. An invoice with no sender

A documented episode illustrates how these abstractions combine in a lower-middle-income country. The monetary tightening undertaken by major central banks from 2022 — a legitimate response, under a domestic mandate, to inflation — mechanically raised debt-service and refinancing costs for countries that had no voice in the decision. For Tunisia, the author estimated this imported cost at TND 4.2-4.5 billion (roughly EUR 1.3-1.4 billion; illustrative, §F.16) at that point in the cycle — the same order of magnitude as the IMF programme of about USD 1.9 billion the country was then negotiating. The external assistance under discussion and the external cost imposed were, in effect, offsetting entries — except that one came with conditionality and the other with no address to which the invoice could be sent (§2.1).

I did not arrive at this architecture through theory; I arrived at it through a ledger. As Minister of Finance of Tunisia, I watched costs I had not chosen and could not vote on land in my country's accounts — lawful, measurable, and with no address to which to send the invoice. The lesson was not about managing a crisis. It was that taking account of one's effect on others cannot be left to goodwill, because goodwill runs against the short-term interest of the actor who would have to exercise it: what a structure does not require does not survive contact with incentives.

Nothing in this observation implies bad faith: central banks were executing their mandates. That is precisely the point — the harm is systemic, lawful, measurable, and without an identified responsible party. No claim in the work depends on this case: it is one family out of twenty.

Positional disclosure. The author served as Minister of Finance of the Republic of Tunisia at the onset of the Covid-19 crisis, and the case study concerns that same country. This proximity is both a source of field knowledge and a legitimate ground for scrutiny: the reader is entitled to ask whether the choice of case, the counterfactual or the attribution bounds lean in favour of the receiving country. The work answers by structure rather than by assurance — conservative zeros, a strictly domestic restriction on one channel, a sign convention that requires a favourable component to be carried alongside the unfavourable one, no netting between agents without a documented transfer, and pre-registration of the joint attribution set. The decisive mitigation is external: the protocol is designed to be executed and replicated by parties unconnected with the author, and the programme is already committed to a second receiving country of a different profile (Annex I).

4. The diagnosis: a dense but discontinuous system

The deficit is not primarily informational: the IPCC exists, the System of National Accounts was revised in 2025, a “beyond GDP” dashboard has been proposed. Nor is it primarily financial: financing is the most crowded function in the present landscape. The deficit is architectural — no arrangement brings measurement, rating, pricing, border adjustment, financing, certification and adjudication into a single mutually coherent chain. Each existing initiative occupies at most one or two links and stops precisely where bindingness would begin (§2.5).

Definition. A functional accountability deficit exists where a necessary function is assigned to no actor simultaneously holding a mandate, a capacity, a duty to justify, a forum of judgment and a regime of consequences.

Definition. An accountability deficit towards affected parties exists where the persons bearing the effects of a decision have no effective avenue of information, representation, contestation or consequence with respect to the holder of power.

4.1 What the register of United Nations sources shows

The programme assembled 136 primary United Nations source documents bearing on externalities, measurement, financing and accountability, deduplicated into 99 distinct primary instruments. Tested against the grid of six routes to binding force, only six bind identifiable subjects; eight are consultative and eighty-five political or non-binding (§3.1). This count concerns the programme's recent United Nations corpus, not international law as a whole; it remains a programme estimate, coded by a single coder, whose nominal register is published in the evidence pack.

Function of the chain	Binding instruments in the corpus	Reading
Humanity Court — adjudicate (§5.8)	Statute of the ICJ; Rome Statute	Covered
Call of Carthage — participate (§6.3)	CRPD, art. 29; ICCPR, art. 25	Covered
Carthage Charter — found and revise (§6.1)	UN Charter, arts. 108-109	Covered
BFM — adjust (§5.5)	Paris Agreement, arts. 2 and 4	Indirectly covered: procedural obligation, not a border measure
HulB — measure (§5.2)	none	Silence
GIRA — qualify (§5.3)	none	Silence
GFEA — price (§5.4)	none	Silence
GSF — finance (§5.6)	none	Silence
IPDA — certify (§5.7)	none	Silence
Preventive loop — prevent (§15)	none	Silence

What this establishes. Six functions out of ten are covered by no binding instrument in the corpus — and they are exactly those the work identifies as the core of the deficit: measure, rate, price, finance, certify, prevent. Bindingness exists where the system knows how to adjudicate and how to enable participation; it disappears the moment an effect must be measured and a consequence drawn from it.

The silence that activity conceals. Before the legal recoding of the register, pricing and financing appeared to be covered — by two General Assembly resolutions on international tax cooperation and by the decision establishing the loss and damage fund. These texts open negotiations and create funds financed by voluntary contributions; they impose no identifiable obligation on any determinate subject. It is precisely where the system has produced the most recent activity that it has produced no obligation at all.

4.2 How a norm becomes binding — the six routes

The diagnosis only makes sense if “binding” is defined. International legal obligation does not flow from a general world authority: it results from a title of competence and a rule of attachment (§3.1.1).

Route	Basis	Principal limit
Treaty in force	Consent to be bound; entry into force	Scope of the treaty; reservations and withdrawal
Security Council decision	Charter, art. 25	Competence and terms; interplay with art. 103
Organisational regulation by tacit acceptance	Constituent instrument (WHO Constitution, arts. 21-22)	Opt-out within the period; scope of the mandate
Judicial decision	Statute, art. 59; Charter, art. 94	Inter partes; enforcement not automatic
Contractual and arbitral consent	Written agreement, arbitration clause	Applicable law, public policy, immunities

Route	Basis	Principal limit
National or regional incorporation	Statute, regulation, authorisation, public contract	Constitutional and judicial review
Customary or peremptory rule	General practice accepted as law	The instrument records or codifies; it does not create

A resolution of the General Assembly satisfies none of these conditions. This is the grid that led to recoding the programme's register and bringing the count of binding instruments down from twelve to six.

5. The theory: three completenesses

The central proposition is the Global Accountability Chain. It defines accountability as a property of the complete architecture rather than of any isolated body. Functional completeness requires that the operations necessary to handle a case be covered. Relational completeness requires that each holder of a function answer before an identifiable forum able to demand information, explanation, justification, correction and, where the law allows, consequence. Procedural completeness requires that interfaces carry not a bare number or a final decision, but the source data, the assumptions, the model version, the uncertainty, the objections, the dissents and the avenues for review (§1.4).

The theory sits within established literatures — actor-forum accountability, global administrative law, rational institutional design, regime complexes and polycentric governance, orchestration and experimentalist governance. The claimed contribution is to connect them into a multi-stage chain, to raise the interface to the status of an object of governance, and to combine negative mandates, stopping rules, comparison of architectures and procedural preservation of evidence (§4.1.2).

The cardinal distinction. The theory is what must survive the abandonment of the instantiation. Refuting the 7+3 does not refute the chain; refuting one case — the monetary one is a single family out of twenty — refutes neither (§4.7).

5.1 The nine design principles

The principles form the normative specification of any instantiation: they apply to the 7+3 as to any rival architecture. The full edition states nine, at §4.2; a tenth, governing prevention, is added at §15.5. They are reproduced here in that order.

P1 — Complete the chain, at a minimum. One accountable institutional capacity per missing function, and no more: measurement (HuIB), rating (GIRA), pricing (GFEA), border adjustment (BFM), financing (GSF), democratic certification of digital systems (IPDA), adjudication (Humanity Court). Seven is not numerology; it is the number of distinct functions in the chain. Nor is it deduced from first principles: it derives from the functional decomposition — one function, one holder, one negative mandate — and must satisfy the eight design requirements (§4.3). The words “at a minimum” state a subsidiarity rule: the architecture duplicates no existing regime that works — the World Trade Organization, the International Monetary Fund and the Financial Stability Board, the Financial Action Task Force, the World Health Organization — it completes and connects them through the measurement layer and the border, where they stop. This decomposition remains a refutable hypothesis, subject to the test of exhaustiveness, inclusion and merger of functions (§4.3.1).

P2 — Single mandates, negative mandates. Each institution does one thing, and its charter states expressly what it does not do. The measurement board does not rate; the rating authority sets no price; the pricing authority levies no tax itself; the fund neither measures nor taxes; the Court guards the guardians, it does not command them — it reviews legality, never expediency, and never substitutes its own judgment for an institution's policy discretion. Negative mandates are the architecture's first line of defence, against mission drift and against the “world government” objection alike. To this separation of functions corresponds a separation of accountability modes: no institution may be held to a conception of accountability incompatible with its function. The Court's independence from State instruction, like the measurement board's fidelity to scientific standards alone, are not accountability deficits but its appropriate forms.

P3 — Jurisdiction by consent. No component binds anyone by its own force. A rating is a public and contestable impact opinion; it acquires legal effect only where a statute, a contract, an authority or a court incorporates it — exactly as credit ratings, accounting standards and technical standards acquire their effect today. The Court's contentious jurisdiction is acquired by consent: treaty accession, contractual arbitration clauses, or derivation through national law for private actors. This is not a weakness to apologise for: consent and incorporation are the principal mechanisms through which international arrangements acquire durable legal effect across jurisdictions.

P4 — Non-substitution. The architecture fills the blind spots of the multilateral system — opposable measurement, consequence, participation — and replaces nothing. It does not amend the Charter of the United Nations, does not place itself above constitutions or treaties, and expressly recognises existing mandates: the independence of monetary policy, the law of the World Trade Organization, the human rights treaty bodies. Its posture towards the United Nations is one of continuity: from “We the peoples” to Monterrey, Addis Ababa and Sevilla, each generation has added the layer its predecessors could not lay.

P5 — One measurement layer; rating is not pricing. Every quantitative claim originates from a single scientific institution, with a public methodology, uncertainty ranges and declared gaps; every downstream institution consumes those measurements and none re-measures. Downstream, the functions remain separate: the rating authority converts measurements into a severity coefficient per externality and a class per entity, but never into a monetary amount; the pricing authority converts coefficient and class into a

price; the border mechanism applies that price to market access. One figure, measured once, transformed at each link by a single institution — this is what makes the chain auditable, and what prevents double counting and double penalty by construction.

P6 — Non-circular budgets. No institution's operating budget may depend on the flows it generates or on which it rules. The fund's initial endowment is independent of the border mechanism's receipts; the Court is financed through an autonomous trust arrangement that no single party can alter; and the rating authority is never remunerated by the entities it rates.

P7 — Sequence: the pilot before the regime. Each component is specified in phases — a pilot perimeter with volunteer jurisdictions and sectors, then extension, then a regime — and the Court itself is staged: advisory chamber, then binding internal review for adherents, then treaty-based jurisdiction. Nothing in the design presupposes wholesale adoption, and each phase is useful even if the next never comes.

P8 — Some things are not for sale. Where an externality arises from gravely unlawful conduct — forced labour, or activity carried out in contexts internationally recognised as unlawful — the response is non-compensable exclusion, never pricing: to put a price on exploitation would be to authorise it for a fee. For such cases the severity coefficient carries a non-compensable cap, and the qualification of unlawfulness itself belongs to law and to courts, never to rating bodies. These situations are treated generically, by category of conduct, never naming jurisdictions. Pricing a harm constitutes neither moral authorisation, nor a presumption of lawfulness, nor a right to buy the possibility of causing harm: it charges the cost without releasing anyone from legal liability, and some flows remain non-compensable.

P9 — Do no harm: the architecture holds itself to its own standard. An architecture designed to correct injustice must first guard against reproducing it; these safeguards are not aspirations but tests, whose failure warrants simplification or withdrawal (§4.4). It must not add complexity it does not offset: failing a demonstration that the separation of functions yields more than it costs, a simpler design would be preferable. It must not cost exorbitantly: its budgets are modest, non-circular and published, and a function that cannot be financed at minimum viable scale is abandoned. It must not create new harm nor aggravate existing harm: every instrument is tested by pilot before any regime, and the risk of government failure is treated as real rather than denied. It must not produce injustice: special and differential treatment protects low-capacity actors, and the burden may never exclude those it means to defend. It must offer no purchase to predators who would divert it: governance parity, the prohibition of issuer-pays, transparent-box methods and non-solidarity between institutions answer capture structurally. And it must never substitute itself for the sovereignty of peoples: nothing binds without consent, primacy remains strictly internal, and no algorithm stands above the will of peoples. Should any one of these safeguards give way without remedy, it is not the rule that bends but the offending organ — or the entire architecture — that must give way.

P10 — Prevent without governing by forecast. Where a cross-border or intergenerational cascade risk is reasonably foreseeable and a window for action remains, the architecture must be able to anticipate it, qualify it, notify it and obtain a reasoned response. An alert never transfers the decision-maker's competence to the producer of the forecast; it creates neither presumption of fault, nor pricing, nor sanction (§15.5).

Four cross-cutting rules. Four formulations recur across the chain without being principles in their own right. They are stated here because they are the shortest way to hold the discipline in mind.

Evidence travels with the file. An interface transmits source data, assumptions, model versions, uncertainty, objections and dissents — not a bare number. This is procedural completeness, stated at §1.4.

Separate measurement, causation, equity and law. Four operations, four thresholds of proof; conflating them is the most common failure of accountability arrangements (§9.6).

The right not to conclude. Insufficient evidence suspends the inference rather than weakening it: stopping rules are a competence, not an admission of failure (§1.7).

Human authority is named. Every public decision is attributable to an identifiable human authority, named in the file (§7.3).

6. The safeguards: what the architecture forbids itself

Negative mandates are the principal safeguard: the measurer neither qualifies nor prices; the qualifier does not reconstruct the science; the consequence authority does not collect on its own account; the financier does not alter the base; the certifier does not govern by algorithm; the adjudicator does not design the policy it reviews. This separation limits role conflicts but creates interface costs. It must therefore be justified by a demonstrable benefit (§1.7, §4.2).

Stopping rules give the architecture the right not to conclude. Insufficient data, an unidentifiable counterfactual, an unreconciled overlap, absent competence, or a disclosure liable to cause disproportionate harm must suspend the inference or downgrade the status of the result.

The S0-S6 classification organises routing: closure, domestic mitigation, capacity, solidarity, efficiency correction or legal remedy. It prevents measurement from automatically becoming a price, and equity from being presented as causation (§9.7).

The legal lock. The most serious class — that of a legal breach — requires the identification of an applicable primary obligation, attribution of conduct, a finding of violation, established injury and a legally sufficient causal link, together with a subject entitled to invoke responsibility before a competent forum. Failing any one of these, the case is routed elsewhere (§5.8.1).

6.1 Interfaces: what travels between functions

The seven institutions are worth nothing unless they are connected. A single measurement layer that no one consumes produces reports; a pricing authority that receives a number without its uncertainty produces arbitrary decisions. The interface is therefore treated as an institutional object in its own right: it has a content requirement, an integrity requirement and a review requirement (§5.9).

Integrity. Interface integrity means that no consequence may be attached to an output that does not carry its justification file. In practice: the receiving function must be able to reconstruct how the figure was produced, under which assumptions, with what uncertainty, against which objections — and must be able to return the file rather than act on it.

Why it matters. This is where most existing arrangements fail. Information flows between international institutions today in the form of published conclusions, not of contestable files. The receiving institution inherits a number it cannot audit, and the affected party inherits a decision it cannot challenge on its merits.

7. Two loops: respond and prevent

A chain that begins only after harm has materialised remains insufficient where propagation mechanisms are reasonably foreseeable. In interdependent systems, an initial decision changes the conditions to which other actors react; those reactions become shocks in turn. A cascade can thus result from locally justified decisions without having been globally intended (§1.5).

The work therefore distinguishes two loops. The Global Accountability Chain remains the principal loop: it handles an observed or sufficiently measurable effect, from evidence to remedy. The Preventive Accountability Loop is transversal: anticipate trajectories, detect exposures and thresholds, qualify an alert, notify decision-makers, obtain a reasoned response, prepare anticipatory action and evaluate ex post the quality of both alert and response. Prevention is neither a guarantee against harm nor a transfer of competence to the forecaster. It is a procedural duty of diligence: a material, cross-border or intergenerational risk, reasonably foreseeable and still actionable, must be able to enter the decision.

The two loops share a versioned case file. No alert automatically creates a price, a sanction or a presumption of liability (§15).

7.1 The preventive loop in detail

Definition. Preventive accountability exists where an actor or a system is required, before harm has fully materialised, to examine a foreseeable risk, to notify the competent authority, to give a reasoned answer and to record that answer in a file that can be challenged. The duty bears on the act, never on the accuracy of the forecast (§15.2).

Seven operations. The loop performs seven operations: anticipate trajectories; detect exposures and thresholds; qualify an alert by level; notify the competent decision-makers; obtain a reasoned response within a fixed period; prepare anticipatory action where warranted; and evaluate ex post the quality of both the alert and the response — including false positives and the effects of the communication itself. Graduated disclosure distinguishes internal watch, coordinated confidential alert, public alert and urgent alert, weighing the need for action, the right to information, data protection, legitimate secrecy and the risk that publication itself triggers the cascade it announces (§15.3, Annex L).

What the loop must not become. Prevention that governs by forecast would reproduce the pathology it seeks to avoid: it would transfer competence to the forecaster and convert uncertainty into authority. The safeguard is that no alert creates a consequence. It creates procedural duties — examine, notify, give reasons — and nothing more. The alerting authority's aggregate performance is itself published, so that its record is auditable (§15.5).

Its binding precedent. The precedent exists in positive law. The International Health Regulations require notification to the World Health Organization within twenty-four hours of any event that may constitute a public health emergency of international concern, permanently reachable national focal points, and core detection and response capacities. Transposed, this is the anticipate-detect-notify-respond duty of the preventive loop: it asks of economic decision-makers nothing more exorbitant than what health law already requires of every State (§15.8).

8. The candidate instantiation: seven functions, three instruments

The work does not assert that these seven functions require seven new organisations. A function may be hosted by an existing institution, a protocol or a federated network; what the theory requires is that the function be held by someone, that its holder answer before an identifiable forum, and that the interface carry the file (§5.1).

Function	Mandate	What it forbids itself	§
Measure — Human Impacts Board (HulB)	An independent scientific body that makes human impacts measurable: it maintains the shared taxonomy of externalities (twenty families), sets the methods, and publishes measurements under an explicit evidential status.	Does not qualify, does not price, does not adjudicate	§5.2
Qualify — Global Impact Rating Authority (GIRA)	Condenses HulB measurements into decision-usable signals: a dimensionless severity coefficient per externality, bounded and neutral at one, and an entity-level rating.	Does not measure (it consumes HulB); sets no monetary price; never rates individuals	§5.3

Function	Mandate	What it forbids itself	§
Price — Global Financial Equity Authority (GFEA)	Never presumes that an impact must be priced. It first determines whether a financial consequence is legally and economically admissible given the class of the case, the quality of identification, net loss and benefits.	Neither measures nor rates, and never recalculates a rating; does not collect — there is no world tax administration in this architecture	§5.4
Adjust — Border Fairness Mechanism (BFM)	Makes the border a point of equity: where a cross-border flow embodies a measured, rated and unpriced human externality, the mechanism applies a correction to market access — always based on the measured gap, never on nationality.	Sets no general tax doctrine; produces no science of its own; decides no use of funds; is not a tribunal	§5.5
Finance — Global Solidarity Fund (GSF)	Turns the architecture's receipts into protection: a Stability window for rapid disbursement after a shock, and a Solidarity window for programmed recovery — grants first, and never debt imposed on the most vulnerable.	Neither taxes nor collects; neither measures nor rates — it commissions HuIB to measure its own results	§5.6
Certify — digital certification authority (IPDA)	The democratic chamber of the architecture. Mandated by national parliaments, it issues transparency standards for algorithmic systems, enforceable through ratification.	Does not measure, rate, tax or adjudicate; regulates no trade; certifies processes and transparency, never content	§5.7
Adjudicate — Humanity Court	The keystone of legality: it reviews the legality — never the expediency — of the acts of the six other institutions against the Charter, and hears claims of breach brought by States, institutions and, subject to admissibility, affected persons.	Does not design the policy it reviews	§5.8

8.1 Measure — Human Impacts Board (HuIB)

Mandate. An independent scientific body that makes human impacts measurable. It maintains the shared taxonomy of externalities — twenty families, from climate and air quality to digital harms and financial instability — sets the methods, and publishes measurements as ranges with confidence grades, never as point estimates carrying false precision. Its data contract specifies, for every published measurement, the metric, the range, the evidential grade and the quality level, under persistent identifiers. Its founding axiom: one can only correct what one can measure.

Negative mandate. It does not qualify, does not price, does not adjudicate. Measurement that grades its own severity has already begun to decide.

8.2 Qualify — Global Impact Rating Authority (GIRA)

Mandate. GIRA condenses HulB measurements into signals usable for decision, and rates exactly two objects. Per externality, it issues a dimensionless severity coefficient — bounded, neutral at one — used downstream to modulate fiscal and customs consequences. Per entity, it issues a class (A+ to E) with an outlook, produced by transparent aggregation under a published and versioned methodology, in which certain red-line findings cap the class in a non-compensable way. It inverts the perspective of conventional ESG: it rates the harm an actor imposes on humanity — impact materiality — not the risk the world poses to that actor's portfolio.

Negative mandate. GIRA does not measure (it consumes HulB); it sets no monetary price (that is GFEA's function — the coefficient is a pure number, never an amount); it never rates individual persons (no social scoring); it never rates acts as such; and it never attributes legal responsibility — fault falls within the exclusive competence of the Court. The three rungs of the rating ladder are strictly ordered: measured (HulB), rated (GIRA), adjudicated (Court); a rated entity is not thereby a liable entity. A corollary: an entity's human impact ledger feeds HulB measurement upstream, where it is verified adversarially — not the rating. GIRA never rates self-declared data, only certified measurements, and it is the interposition of an independent measurer that separates it from ESG self-reporting.

8.3 Price — Global Financial Equity Authority (GFEA)

Mandate. The Global Financial Equity Authority never presumes that an impact must be priced. It first determines whether a financial consequence is legally and economically admissible, given the class of the case, the quality of identification, net loss, offsetting benefits, domestic vulnerabilities and proportionality. It may conclude that no consequence is due, recommend a non-financial measure, route the case towards solidarity, set a reference price for an efficiency correction, or find that a question belongs to applicable law rather than to pricing. It does not measure, does not requalify the facts, and never turns a decision of equity into proof of liability.

Negative mandate. GFEA neither measures nor rates — and never recalculates a rating. It does not collect: there is no world tax administration in this architecture; collection remains with the national systems that incorporate the schedule (fiscal subsidiarity, P3). It does not redistribute (the fund's function) and does not adjust at borders (the mechanism's function).

8.4 Adjust — Border Fairness Mechanism (BFM)

Mandate. The Border Fairness Mechanism makes the border a point of equity: where a cross-border flow embodies a human externality that has been measured, rated and left unpriced, the mechanism applies a correction to market access — always grounded in the measured gap, never in nationality. Its response is calibrated by category of case, not uniform: a price adjustment for measurable environmental externalities (the CBAM logic, generalised); a non-compensable exclusion for gravely unlawful conduct; a levy on destabilising financial flows, in the tradition of the Tobin tax — the pricing doctrine belongs to GFEA, the mechanism merely applies at the point of passage the schedule already set; and mandatory disclosure for financial services and products. Receipts are assigned to the solidarity fund.

Negative mandate. The mechanism sets no general tax doctrine (GFEA's function), produces no science of its own (it consumes HulB and GIRA), decides no use of funds, is not a tribunal (first-instance disputes go to its

contentious chamber, appeals to the Court), and is never an instrument of geopolitical retaliation — nationality neutrality is absolute.

8.5 Finance — Global Solidarity Fund (GSF)

Mandate. The fund turns the architecture's receipts into protection: a Stability window for rapid disbursement after a shock, and a Solidarity window for programmed financing of durable recovery — grants first, and never debt imposed on the most vulnerable. Allocation follows the equity measurements produced upstream — exposure measured by HuLB, priority set by GIRA — under rules issued by the pricing authority, with full traceability from receipt to result. The exposure measurement distinguishes the suffered, structural and eligible share from the reasonably avoidable share arising from a policy choice: the fund repairs and protects, it does not subsidise imprudence.

Negative mandate. The fund neither taxes nor collects (GFEA and the mechanism feed it); it neither measures nor rates — it commissions HuLB to measure its own results, a deliberate inversion of the usual pathology of self-evaluation; it is not a development bank seeking a return; and it imposes no austerity conditionality — its conditionality is traceability, not the surrender of policy sovereignty. This discipline applies differently by window: the Stability window disburses unconditionally on parametric triggers, because one does not condition emergency relief; the Solidarity window directs its financing towards a recovery that reduces future exposure, so that support ends as measured exposure falls — it is success, not failure, that closes it. In both cases the only requirement is traceability from receipt to the result measured by HuLB, never a policy prescription. This separation — the fund finances but does not rate — is precisely what distinguishes it from a lender that is also its own judge: the condition is set by a neutral measurement, not by the creditor.

8.6 Certify — the digital certification authority (IPDA)

Mandate. The digital authority is the democratic chamber of the architecture. Mandated by national parliaments — the institutions that already hold the legitimacy the digital sphere lacks — it issues transparency standards for algorithmic systems, enforceable through inter-parliamentary ratification; it qualifies AI systems of high human impact against thresholds co-defined with the measurement body; it accredits parliamentary delegations; and it certifies the deliberative infrastructures the architecture itself uses. Its doctrine on artificial intelligence fits in one sentence: AI signals; elected humans decide.

Negative mandate. It does not measure, rate, tax or adjudicate; it regulates no trade; it is not a global ministry of truth — it certifies processes and transparency, never content, and censors nothing; and it is not “an AI that governs”. Yet platform abuse escapes no consequence: detected and, where warranted, decertified by the digital authority, it is downgraded by GIRA, priced by GFEA and the Border Fairness Mechanism, adjudicated by the Court and enforced by the national authorities that incorporated the standard — the sanction is the loss of certified market access, not a verdict of the authority itself, which neither judges nor sanctions.

8.7 Adjudicate — Humanity Court

Mandate. The Court is the keystone of legality: it reviews the legality — never the expediency — of the acts of the six institutions against the Charter; it rules on claims of breach brought by States, by institutions and, subject to admissibility, by directly affected persons; it settles conflicts of competence; it issues preliminary rulings on interface disputes and preventive advisory opinions. Its remedies are graduated: declaration,

injunction, suspension, annulment with remand — and, as soon as a claim is admissible, suspensive effect on the contested downstream act. Its cardinal principle, and the signature of the architecture: the Court guards the guardians; it does not command them. It may annul an unlawful rating or adjustment; it may never dictate the content of the act that replaces it.

What is still missing. Two techniques are still missing from its conventional regime, and the work identifies them as such: complementarity on the model of the Rome Statute — the Court hears a case only where no competent national, regional or sectoral forum has examined it under equivalent guarantees — and declarable exclusions by category, on the model of the ICJ Statute and UNCLOS, which may never bear on the non-derogable floors of the Charter: dignity, non-discrimination and access to remedy (§5.8).

8.8 Comparison, and who pays

Comparison is mandatory. Seven is not a result: it is a comparative hypothesis. The work requires the instantiation to be compared with an improved status quo, a minimal protocol and a federated configuration, and requires the lightest arrangement to be retained at equivalent guarantees (§4.3-4.6, Annex E). Functional separation is retained only where it improves precision, legibility or independence — and it must pay for its own interface costs (§4.4).

Who pays. One principle bounds the whole design: finance the measurement — but the rated never pays the rater (§5.10). Three costs must be kept apart. Keeping the impact ledger is a compliance cost borne by the entity, like its accounts or its tax returns. Verifying and rating that ledger can never be paid for by the entity verified: the issuer-pays arrangement is the documented source of the Enron failures, of the capture of credit rating agencies and of the limits of ESG — whoever pays the assessor shapes the verdict. And priced harm flows downstream to the fund: it is the consequence, not the cost of measuring. Independent verification is financed instead by a mutualised levy on a neutral formula, on the model of the PCAOB, the FDIC and ESMA.

8.9 The three instruments

The Carthage Charter. The Carthage Charter is the legal foundation of the architecture. It sets out the rights, guarantees, decision principles and limits applicable to the institutions of the chain: dignity and non-discrimination; information, participation and counter-expertise; the requirement to give reasons; reversibility; and access to remedy. Three of these are non-derogable floors that no reservation, no emergency and no declarable exclusion may reach — dignity, non-discrimination and access to remedy. The Charter creates no automatic consequence: measuring an effect entails neither fault, nor pricing, nor financing, nor sanction. Its primacy is strictly internal: it prevails within the order of the architecture, and expressly subordinates itself to the Charter of the United Nations, to constitutions and to applicable treaties (§6.1, Annex N).

The Global Accountability Framework. The Global Accountability Framework is the operational framework agreement. It organises the mandates, interfaces, data, procedures, financing, audits, alerts, correction mechanisms and avenues of remedy required for the consented application of the chain and of the preventive loop. It is modular: each party designates, in a public declaration, the functions, sectors and territories it activates, and no competence is inferred from silence, from attendance at a meeting or from the supply of data. It distinguishes what is amended — mandates, competences, negative mandates, guarantees — from what is adjusted: quantitative parameters, thresholds, coefficients, taxonomies and file formats, which are revised by qualified decision with automatic effect, subject to an objection period and to the right of a refusing party to withdraw from that module alone (§6.2, Annex O).

The Call of Carthage. The Call of Carthage is the civic instrument. It organises four channels — adherence, qualified petition, deliberation by sortition, and transmission — with their thresholds, deadlines and guarantees. It is consultative: it may feed the agenda, participation and contestation, but it modifies no competence unless the Framework incorporates it. Its precedent is not hypothetical: the compliance committee of the Aarhus Convention, open to communications from members of the public including non-governmental organisations, has operated since 2002 — the objection of impossibility does not hold (§6.3, Annex P).

Source clause and no presumption of bindingness. Nothing in these instruments binds an entity not otherwise legally bound. Participation in a reference standard, a consultation, a public register, a simulation or a shadow-mode pilot constitutes neither recognition of liability, nor consent to a jurisdiction, nor acceptance of a levy, a compensation or a border adjustment. Any binding effect states expressly its legal vehicle, its addressee, its scope, its date of applicability, the forum of review and the enforcement regime.

9. The tooling: institutions assisted by AI, governed by humans

Authority matrix. Algorithmic authority is classified by verb: search, summarise, recommend, validate, decide. Only the first three may be delegated to a system; validation and decision remain with an identifiable human authority, named in the file (§7.3).

Three requirements frame this tooling: an identifiable human authority for every public decision; contractual plurality of models, so that no vendor error propagates from one link to the next; and the capacity to operate without artificial intelligence, in degraded mode (§7).

10. The methodological laboratory: monetary spillovers

The monetary programme is not the thesis of the work: it is the ground on which the attribution discipline is pushed furthest, on one family of externalities out of twenty (§9).

When major central banks tighten, the cost reaches the receiving country through at least fifteen cash-flow mechanisms grouped in four families — plus one balance-sheet memo item that must never be added to the others. The decomposition is disciplined but not exhaustive, and each mechanism carries a signed prediction and rank conditions that make it testable (§9.1).

Before any computation, one distinction governs everything that follows: accounting incidence is not causal attribution, and causal attribution is not legal responsibility. Bounded partial attribution replaces a single figure with ranges derived under joint restrictions and an explicit counterfactual (§9.2).

Three registers, never added. Register I is cash-flow incidence: fifteen mechanisms in four families, answering the question how much leaves the country and what share of it is attributable to a decision. It is the register of attribution: it supplies the bands. Register II is balance-sheet revaluation — the restatement, in local currency, of foreign-currency debt outstanding. It is a memo item, never added to the others, and the work does not claim it. Register III is welfare loss: the account that was missing, which the work does not produce but borrows from an independent source — the output foregone and the investment not made.

The rule that governs the base. A transfer has a beneficiary; a destroyed investment has none. Additional interest paid to a creditor is, until shown otherwise, a transfer: it leaves one balance sheet for another, and the global loss, if any, lies in the distortions it triggers rather than in its amount. Output not produced, and investment not made and still not made ten quarters later, are nobody's gain. That, and only that, is where inefficiency resides. The claim of the correction window therefore does not bear on Register I: it bears on the attributable share of Register III. Register I supplies the attribution bands; Register III supplies the base. To conflate them would be to claim a cash outflow in respect of a loss, or a loss in respect of a cash outflow (§9.3).

The epistemic rungs. No figure rises above its rung — accounting incidence, plausible attribution, identified causation. The published Tunisian figures belong to the first rung and are presented as disciplined illustrations, not as verifiable aggregates (§9.2, Annex F).

10.1 From spillover to consequence: what must be established

A negative spillover is not, by itself, an inefficient externality calling for financial correction. Four qualifications must be kept apart — accounting incidence, welfare loss, distributive injustice, internationally wrongful act — and none follows automatically from another. This section draws the institutional consequence: what the chain must establish before attaching a financial consequence to a spillover, and to whom that consequence is owed (§9.6-9.7).

The benchmark, and the objection it carries. The normative literature on spillovers supplies the most serious objection to this work's thesis, and it must be stated without weakening it. Korinek (2016) shows that spillovers arising from unilaterally optimal national policies can be Pareto-efficient — so that further correction would redistribute without creating efficiency — provided three conditions hold: policymakers exercise no market power; they have instruments sufficient to act on their external transactions; and international markets are free of imperfections. The mere fact that a harm is suffered rather than chosen

therefore does not ground a correction. An architecture that priced every measured spillover without clearing this step would be refuted on its most visible case.

What the theorem establishes, and what it is silent on. Korinek's result concerns the efficiency of an allocation given the objectives pursued. It says nothing about the perimeter of those objectives. A central bank pursuing a domestic statutory mandate — maximum employment, price stability — decides optimally with respect to an objective function from which a measured cost, borne by identifiable third parties, is excluded by construction. That exclusion is not a market imperfection: it is a property of the mandate. And this is exactly where Pigou (1920) locates the externality — the decision-maker's private cost diverges from the social cost, so that the decision is optimal only with respect to the wrong magnitude. The thesis defended here is therefore not that the issuing central bank decides badly. It is that it decides on an incomplete ledger.

This thesis is not new, and it is not mine. Ocampo formulated it a quarter of a century ago, in the very terms this work adopts: “Through the impact of monetary policy on economic activity and exchange rates, the center economies generate externalities to the rest of the world that are not internalized by policy makers. These externalities are strongly felt in the developing world, which must adjust to them but lack the degrees of freedom that the ability to supply international currencies provides” (Ocampo, 2001, p. 10). The vocabulary is Pigou's; the diagnosis is this work's; and it is twenty-five years old. The same text names the resulting asymmetry: center economies are “business-cycle makers” and “policy making”; the periphery is made of “business-cycle takers” and is “policy taking”.

This work therefore does not add the diagnosis. It adds the institution that internalises.

The three conditions, applied to the case at hand. For an economy characterised by dominant-currency invoicing, a foreign-currency balance-sheet asymmetry, a shallow foreign-exchange market and imperfectly anchored expectations, the second condition — sufficient instruments — does not fail case by case: it fails by construction. Rey (2015) and Miranda-Agrippino and Rey (2020) establish that in the presence of a global financial cycle the trilemma becomes a dilemma: even a floating exchange rate does not insulate. Gopinath et al. (2020) show that where trade is invoiced in a dominant currency, exchange-rate flexibility loses most of its expenditure-switching power. The receiving country's shortage of instruments is then not a national failing for which it could be blamed: it is a property of the international monetary system. Nor is the third condition met — no complete market allows an exposed jurisdiction to insure against a foreign monetary decision, and it is precisely that missing negotiating forum which §12 analyses.

The burden of proof, and it is the heart of the section. Where the structural characteristics set out above are established, the presumption is one of inefficiency, and it falls to the issuing jurisdiction to rebut it. To place on the receiving country the burden of demonstrating, mechanism by mechanism, that no instrument was available to it would be to impose on it the cost of a design failure it did not author. This inversion is not academic: in July 2023, the governing body of the International Monetary Fund, having found large negative spillovers disproportionately affecting emerging markets, concluded that those markets should float more and better anchor their expectations (§9.7).

What follows. Only where these steps are cleared does the S0-S6 routing authorise a consequence — and even then, only two classes open efficiency correction. The remaining classes route the case towards closure, domestic mitigation, capacity building, consented solidarity or legal remedy (§9.7).

10.2 The monetary programme in detail

Fifteen mechanisms, four families. When major central banks tighten, the cost reaches the receiving country through at least fifteen cash-flow mechanisms grouped in four families — external debt and refinancing; capital flows and exchange rate; trade and terms of trade; domestic financial conditions — plus one balance-sheet memo item that must never be added to the others. Each mechanism carries a signed prediction and rank conditions that make it testable: the decomposition is disciplined, but it is not exhaustive, and it is conditional on the counterfactual adopted (§9.1).

The attribution discipline. Before any computation, one distinction governs everything: accounting incidence is not causal attribution, and causal attribution is not legal responsibility. Bounded partial attribution replaces a single figure with ranges derived under joint restrictions and an explicit counterfactual. Seven attribution rules follow — among them the prohibition on double counting between flows, stocks, valuations, incomes and exposures; the requirement that any figure declare the rung it occupies; and the stopping rule that suspends the inference where the counterfactual cannot be identified (§9.2).

Non-additivity. Three registers never add up: cash-flow incidence, balance-sheet revaluation — a memo item the work does not claim — and welfare loss. Each answers a different question and calls for a different treatment; adding them would produce a number that means nothing, and the temptation to do so is precisely what the rule forbids. Their composition, and the rule that determines which of them supplies the base of a correction, are set out at §10.2 below.

What the figures are, and are not. The published Tunisian figures belong to the first rung — accounting incidence — and are presented as disciplined illustrations under an inherited counterfactual, not as verifiable aggregates. No world-scale aggregate is reported: it was withdrawn in July 2026 and will not return until country coverage, weights, exchange-rate-regime heterogeneity, missing observations and the overlap with domestic drivers are published. This is a stopping rule applied to the author's own most striking figure (§9.3, Annex F).

11. Testing: four grounds, four evidential statuses

Ground	What it tests	Evidential status	§
Integrated demonstrator — an AI data centre	The full path of a case file through the seven functions, and three simulated challenges	Full simulation, not a real case	§8
Monetary policy spillovers	Fifteen signed mechanisms, bounded partial attribution, full cycle, buffer exhaustion	Accounting incidence; causation not established	§9, Annexes F-M
Development finance (Sevilla)	The execution and accountability layer of an existing global framework	Mapping of 129 initiatives, internal coding	§10
Preventive peace	Extension of the preventive loop to conflict prevention	Exploratory extension programme	§11

What they do not demonstrate. None of these grounds demonstrates the superiority of the architecture: they test its coherence and feasibility. The demonstration proper belongs to the three-arm comparative pilot — status quo, light federated solution, full instantiation — on the same case files and against pre-registered indicators (§12.5).

11.1 Development finance: what the Sevilla mapping shows

The Fourth International Conference on Financing for Development (Sevilla, June-July 2025) produced the Compromiso de Sevilla and a Platform for Action — the largest gathering of the decade on development finance. Mapped one by one onto the functional chain, its 130 initiatives display a striking distribution: seventy-one — more than half — bear on financing, but only three on rating, none on adjudication, and none creates a binding instrument. This is a count of initiatives — of institutional attention — not of funds mobilised, which remain far below needs. The system masses its initiatives on financing without closing the gap; in this voluntary corpus rating remains rare, no initiative is primarily adjudicative, and none by itself creates a binding obligation. This distribution is not a failure of the conference: it is a revealed preference of the present system — and the precise specification of the functional and interface deficit the architecture is designed to close (§3.3).

Sevilla did not lack advice. The conference was prepared by an International Expert Commission on Financing for Development, convened for it, whose report — Financing a Sustainable Future — was delivered on 11 February 2025 to inform the negotiation of the outcome document. It was chaired by José Antonio Ocampo, a member, sixteen years earlier, of the 2009 Expert Commission; and among its fourteen members was Joseph Stiglitz, who had chaired that earlier commission. The authors of the 2009 proposal thus advised the 2025 conference.

And the 2025 conference produced seventy-one financing initiatives, three rating initiatives, none on adjudication, and no binding instrument.

What is proposed. The proposal is to add a common, light and federated layer — not a new central organisation. It remains a protocol subordinate to the instruments of the 7+3 and to the mandates of the initiative holders, and it standardises eight objects: an initiative passport (identity, purpose, actions covered, responsible parties, countries, beneficiaries, legal basis, duration); a theory of change with its counterfactual; and the interfaces through which evidence, reasons and challenges travel (§10.5).

11.2 Preventive peace, and the comparative pilot

The peace chapter states the conditions under which the Global Accountability Chain may contribute to conflict prevention — and the conditions under which it must not. Prevention here does not mean forecasting conflict: it means ensuring that a foreseeable, material, cross-border risk can enter a decision, that the competent authority is notified, and that a reasoned answer is given and recorded. The chapter is presented as an exploratory extension programme, not as an established result (§11.1-11.2).

The demonstration proper belongs to the comparative pilot. Three arms — improved status quo, light federated solution, full instantiation — run on the same case files, with pre-registered indicators, stopping rules and independent evaluation. The pilot is designed to be able to fail: if the lighter arms deliver equivalent guarantees at lower cost, the architecture is simplified or abandoned. That possibility is not a rhetorical concession; it is the condition under which the proposal remains refutable (§12.5).

12. What the work establishes, and what it does not

Not all propositions carry the same force. This distinction is not a peripheral precaution: it is part of the theory of accountability, because a decision cannot be more certain than the evidence it carries (§1.10, Table 1).

Status	Meaning	Consequence for reading
VERIFIABLE FACT	Documented by an identifiable source, subject to that source's limits	May support the diagnosis, not necessarily causation or obligation
PROGRAMME RESULT	Estimate, coding or simulation produced by the programme, still contestable and sometimes unreplicated	Requires data, code, assumptions, ranges and a publication threshold
DESIGN PROPOSAL	An institutional or procedural choice offered for comparison	Must be compared with lighter alternatives and submitted to a pilot
OBJECTION	A criticism, conflict of principles or anticipated failure mode	Must remain visible even where an answer is offered
STOPPING OR REJECTION CRITERION	A condition forbidding attribution, pricing, activation of a link or maintenance of the architecture	Guards against automaticity and makes the proposal refutable

Maturity. The twenty-five claimed contributions are graded on a maturity scale measuring what the work is currently entitled to assert, not the importance of the contribution. The conceptual families reach the highest level; the empirical families — anatomy of spillovers, attribution and distributive tracing — remain below it, and the work says so (§1.11).

12.1 The declared limitations

The work states its own limits before its critics do. Four are structural (§14).

Empirical validation. The formal model and the empirical protocol are more advanced than the results they have produced. The monetary figures are accounting incidence, not identified causation; the Tunisian case is a single-country illustration under an inherited counterfactual; the regional panel does not yet exist.

Coding reliability. Both classification exercises — the Sevilla mapping and the register of United Nations sources — were coded by a single coder. Independent double coding and an inter-coder agreement coefficient remain to be produced. The selection protocol for the original 136 sources is not documented, which means the corpus is not yet replicable even with the register in hand.

External legal review. The internal legal texts are advanced, but no external opinion has been obtained act by act. Eight distinct opinions are required — public international law, State responsibility, trade, international taxation, immunities, arbitration and recognition, data, human rights and due process.

Costing. No bottom-up costing of the three configurations exists. Until it does, the claim that the architecture costs less than the counterfactual remains an argument, not a result.

These limitations are not caveats appended to a conclusion. They are the reason the work presents itself as a refutable proposal rather than an established result — and the reason its first expected deliverable is external validation, not wider advocacy.

13. Objections and rejection criteria

Objection	Answer, in one sentence	§
"The powerful will never consent"	The objection assumes a veto over an existing proposal; but one does not veto what no one has proposed. The decisive variable is the absence of a duty-holder, not the balance of power alone.	§2.6-2.7, §13

Objection	Answer, in one sentence	§
“One more commission”	The mandate sought concerns execution and interface failures, with an explicit non-duplication clause and an obligation to inventory prior recommendations before proposing anything new.	§12.7.2
“It will cost more than it returns”	The comparator is not zero but the counterfactual — the cost of inaction and exposure to future crises. The bottom-up costing of the three configurations remains to be produced, and the work says so.	§13.3, §14
“It has not been replicated”	Correct, and the work records this as its first limitation: the empirical programme is more advanced than its results. Independent replication is the first expected deliverable.	§14

Rejection criteria. The architecture must be capable of being abandoned. The criteria are explicit: if the existing architecture or a federated solution achieves equivalent results at lower cost; if the arrangement creates a new capacity asymmetry; if data collection disproportionately infringes rights; or if political adherence requires abandoning traceability, contestation or the stopping rules (§12.7.8, §14).

13.1 Four objections, in full force

The work devotes an entire chapter to objections and discards none. Four deserve to be stated at length, because they are the ones a hostile reviewer would raise first (§13.2).

“The powerful will never consent.” The objection assumes a veto over an existing proposal. But one does not veto what no one has proposed. Where a function has no holder, there is nothing to obstruct — and nothing to concede. The decisive variable is the absence of a duty-holder, not the balance of power alone. This does not make the objection vanish: it relocates it. What must be produced is not consent to an architecture, but a mandate to compare architectures — a far lower political threshold, and the one the adoption sequence targets.

“This is one more commission.” The risk is real and the work treats it as a design constraint rather than a rhetorical difficulty. The mandate sought bears on execution and interface failures, not on a fresh general list of reforms; it carries an explicit non-duplication clause; it requires prior recommendations to be inventoried, with their implementation status and the documented reasons for non-execution, before anything new may be proposed; and it is entrusted to an existing mechanism rather than to a new commission.

“Multiple accountabilities disorder.” Multiplying institutions, forums, audits and judicial avenues risks incompatible oversight demands that no institution could satisfy at once. The answer is discipline, not reduction. To the separation of functions is added a separation of accountability modes: each institution receives a distinct accountability profile which ranks the forms of oversight compatible with its function and excludes those that would corrupt it. A demand legitimate for one institution may be pathological for another — the measurement board must answer to scientific standards and to the data of affected parties, not to a State seeking a more favourable measurement; the Court must remain faithful to the law, never directed by opinion or by those who appointed it. Absent such an order, an organisation oscillates between contradictory expectations and lets its leadership choose, after the fact, the conception of accountability that justifies the decision already taken.

“The institutions will be captured, or will drift.” Any new body may be captured by those it is meant to hold to account, or drift towards self-preservation. The architecture answers with bounded design: revision and sunset clauses; external feedback loops — affected parties, peers, judicial review — rather than self-evaluation alone; professionally heterogeneous composition against the homogeneity that breeds blind spots; and a precise mission whose redefinition must be publicly justified. Preventing pathologies is not an added corrective: it is part of the mandate.

14. Adoption: a sequence, not a switch

A global governance architecture does not pass from publication to legal obligation. Its adoption is a succession of decisions of differing scope; the initial object is not the adoption of the 7+3, but the adoption of a procedure allowing the status quo, a light federated architecture and the full instantiation to be compared on the same case files (§12.7).

Phase	Object	Decision gate
0 — External validation	Open the evidence, legal opinions, replication	Do not seek political traction while key evidence remains closed
1 — Incubation	ECOSOC Forum on Financing for Development, transregional coalition	Proceed only if a transregional coalition exists

Phase	Object	Decision gate
2 — Initial mandate	A procedural General Assembly resolution	No permanent body created, no automatic consequence
3 — Comparative pilots	Three arms, independent evaluators, published results	Simplify or abandon absent a demonstrated net gain
4 — Modular incorporation	Boards, contracts, model laws — function by function	Each act passes the legal gate and the twelve-question test
5 — Targeted binding law	Sectoral protocols, consented jurisdiction	No global treaty if lighter mechanisms provide equivalent guarantees

The sequence is cumulative, modular and reversible: at every gate, simplifying, federating or closing remain honourable outcomes. The elements of a model procedural mandate are set out in Annex Q of the full edition.

Political economy. Adoption requires a coalition, not merely incorporable results. It must be transregional — exposed countries, countries with analytical capacity, small States, emerging economies, development partners, and at least one State able to engage the principal issuing jurisdictions. A coalition composed solely of affected countries would make the project look like a unilateral distributive claim; one confined to major issuers would risk neutralising its substance (§12.2-12.3).

14.1 What is claimed, and what is not

The claim of this work is narrower than a new general theory of global power, and more precise. It is that accountability for cross-border and intergenerational effects is a property of a complete chain rather than of any single institution; that the chain can be specified functionally, tested empirically and adopted modularly; and that the proposal must be capable of being simplified, federated or abandoned on evidence.

What is offered is a transportable grammar for evaluating any accountability architecture — not a finished organisation chart. The seven functions and three instruments are one instantiation among several, and the work requires them to be compared with lighter alternatives before being retained. The theory is what must survive if the instantiation does not.

The distance still to travel is stated as plainly as the claim. The empirical programme is more advanced than its results; both classification exercises were coded by a single coder; no external legal opinion has been obtained act by act; no bottom-up costing exists. Until those are produced, this remains a refutable proposal — which is what it is meant to be.

14.2 In one sentence

The claim of this work is narrower than a new general theory of global power, and more precise: that accountability for cross-border and intergenerational effects is a property of a complete chain rather than of any single institution; that the chain can be specified functionally, tested empirically and adopted modularly; and that the proposal must be capable of being simplified, federated or abandoned on evidence. What is offered is a transportable grammar, not a finished organisation chart.

15. Where to find the demonstration

This short version carries the argument; the full edition carries the evidence.

Section here	In the full edition	Evidence pack
2. The landscape	§2.2-2.4, Boxes 2 and 3	Twenty-families workbook
3. The vignette	§2.1; Annex I	Configuration M, source data
4. The diagnosis	§2.5, §3.1, §3.1.1, §3.3	EP1 Sevilla; EP2 UN register
5. The theory	§1.4, §4, §4.1.2, §4.2	—
6. The safeguards	§1.7, §4.2, §9.7, §5.8.1	—
7. The two loops	§1.5, §15, Annex L	—
8. The instantiation	§5, §6, Annexes N-P	—
9. AI tooling	§7, §8	—
10. Monetary spillovers	§9, Annexes F-M	EP3; companion paper econometric chain
11. The four grounds	§8 to §11	44 LDC panel
12. Evidential statuses	§1.10, §1.11, Table 1, Annex A	Evidence pack index
13. Objections	§13, §14	—
14. Adoption	§12, §12.7, Annex Q	—

Traceability. The evidence pack is versioned and indexed: each correction produces a dated copy, the original is left untouched, and a changelog records the change cell by cell. Its index lists the items, their location and, explicitly, what does not yet exist.

Status of this version

Short version extracted from the full edition — 350 pages, 54 tables, 11 figures, seventeen annexes. It contains no claim absent from the full edition and alters no evidential status. Where formulations differ, the full edition governs.

Working version open to comment. Not yet peer reviewed. Licence: Creative Commons Attribution 4.0 International (CC BY 4.0).