Economic Impacts of the IMO Net-Zero Framework on African Economies — Technical Annexes

Ralph Matlack · Fourth Tack LLC · Last Updated: 09 Apr 2026

Contents

Notation Summary 2
Annex A: Country Coverage Details 3
Ethiopia and the GTAP Landlocked-Africa Aggregate 3
AF14 Composition 4
Comparator Country Selection 4
Annex B: Scope, Boundaries, and Limitations of the Reduced-Form Approach 4
Economic Channels Not Captured 4
Temporal Boundaries 5
Interpretive Considerations 5
Annex C: Derivation of the Two-Tier Adjustment Factor (α) 6
Regulatory Parameters 6
Fleet Composition 6
Derivation 6
Structural Asymmetries 7
Post-2030 Convergence 7
Sensitivity to S24 Clearing Price 7
Sensitivity to Fleet Composition 7
Annex D: Reward Mechanism Parameters — Derivations and Evidence 8
Reward Calibration Efficiency (η): Supply-Curve Framework 8
η by Mechanism Design 8
Market Structure Assumption 9
Fuel Scope Factor (φ): Derivation 9
SU Market Factor (μ): Derivation 9
Robustness to Wider Ranges 9
Usability Factor (ω): Calibration Evidence 10
Annex E: Revenue Estimation and Temporal Trajectory 11
First-Principles Revenue Estimate 11
CIA Task 2 Structural Benchmark 11
Temporal Trajectory 12
Compounding Effects at 2050 12
Maritime Logistics Cost Estimate (M) 12
Annex F: Cross-Scenario Validation and Robustness 13
Full Cross-Scenario Validation Table 13
. 15

Summary Statistics 15
Affine Specification Robustness Check 15
Data Note: CIA Scenario 26 15
AFPH Cross-Check 16
CPI Response Pattern Analysis 16
Food Price Decomposition: Data Sources and Methodology 16
Trade Cost ρ Derivation Methodology 17
Annex G: Calculator Specification and Worked Example 17
Complete Slider Specifications 17
Computation Chain 21
Worked Example: Kenya, 2030 21
Full Range (Kenya, 2030) 22
Fixed Parameters 22
Sources 22
Annex H: NZF Design Topic Inventory 22
Topic Grouping 1: IMO Net-Zero Fund 23
Topic Grouping 2: ZNZ definition and/or reward 26
Topic Grouping 3: Fuel Certification 29
Topic Grouping 4: GFI and GFI compliance approaches 32
Topic Grouping 5: IMO GFI Registry 34
Topic Grouping 6: IMO Life Cycle GHG Assessment (LCA) framework 38
Annex I: Topic Classification and Cluster Mapping 41
Topic-to-Bucket Classification 41
Bucket X Topic-to-Cluster Mapping 42
Bucket Y Topic-to-Risk-Cluster Mapping 43
Annex J: References 43
IMO Instruments and Submissions 43
CIA and Impact Assessment 44
Compliance Cost and Maritime Decarbonisation 44
African Economic Impact 44
Fiscal Multipliers and Climate Finance 44
Renewable Energy and Hydrogen 44
Trade and Economic Data 45
National Statistical Sources (CPI Weights) 45

These annexes provide the detailed derivations, evidence base, reference materials, and robustness analysis supporting the reduced-form framework described in *Economic Impacts of the IMO Net-Zero Framework on African Economies — Methodology Annex*. Each annex is self-contained and referenced from the main text.

Notation Summary

Symb	Name	Range	Source / Reference
α	Two-tier adjustment factor	1.05–1.29	Annex C; §2.4.2
T_2	Effective Tier 2 compliance price	\$100–\$380	Annex C; §2.4.2

Symb	Name	Range	Source / Reference
σ	Reward budget share (Fund revenue split)	0.20–0.80	§2.4.2; Slider 0
η	Reward calibration efficiency	0.55–0.88	Annex D; §2.4.2
φ	Fuel scope factor	1.00–1.05	Annex D; §2.4.2
μ	SU market factor	1.00–1.06	Annex D; §2.4.2
π	Fuel price factor	0.85–1.15	DNV sensitivity; §2.4.2
SR	Sensitivity ratio (total MLC / S24 MLC)	0.85–1.45	§2.4.2; Annex F
ω	Usability factor (= δ × λ)	0.10–0.63	Annex D; §2.4.4
δ	Disbursement effectiveness	0.35–0.90	Annex D; GCF data
λ	Fiscal multiplier	0.30–0.70	Annex D; Batini/Ilzetzki/Kimaro
θ	Fiscal transfer (% GDP, country-level)	Derived	§2.4.4
ρ	Import composition adjustment	Country-specific	Annex F; UN Comtrade
R	Total NZF revenue	~$10B/yr (2030)	Annex E; UCL/UMAS
M	Total global maritime logistics cost	~$800B/yr	Annex E; CIA Table 13
Φ	CIA S24 reference impact value	Country/metric-specific	CIA Task 3 Annexes; Table 3.1

Annex A: Country Coverage Details

Ethiopia and the GTAP Landlocked-Africa Aggregate

Ethiopia is represented in the CIA Task 3 results through the “Rest of landlocked economies in Africa” GTAP aggregate, which includes thirteen countries: Ethiopia, Uganda, Zambia, Mali, Burkina Faso, Niger, Rwanda, Malawi, Chad, South Sudan, Burundi, Central African Republic, and Lesotho. Ethiopia is the dominant component, accounting for approximately 46% of the aggregate’s GDP, with Uganda (15%) and Zambia (9%) as the next-largest components.

The aggregate’s trade structure is diversified across multiple transit corridors: Ethiopia’s imports are routed almost entirely through Djibouti; Uganda’s transit via Mombasa; Zambia’s via Dar es Salaam and Durban; and the Sahelian countries’ via West African ports (Abidjan, Lomé, Cotonou, Douala). This corridor diversification means the aggregate’s effective maritime cost exposure reflects a blend of East African, Southern African, and West African shipping routes, diluting Ethiopia’s specific Djibouti corridor dependency.

Results labelled “Ethiopia” in the analysis and calculator should be interpreted as representative of landlocked African economies generally, with Ethiopia as the dominant but not sole component. Users requiring Ethiopia-specific estimates should note that the aggregate likely understates

Ethiopia's corridor-specific vulnerability, because the blended corridor exposure moderates the concentration risk that Ethiopia faces through its near-exclusive dependence on Djibouti.

AF14 Composition

The African Focus 14 comprises the AF5 (Ghana, Kenya, South Africa, Egypt, Ethiopia) plus Algeria, Cameroon, Côte d'Ivoire, Gabon, Madagascar, Mauritius, Morocco, Nigeria, and the United Republic of Tanzania. Each country's impact is computed independently using country-specific CIA data and structural parameters, then averaged using GDP weights. This preserves country-level differentiation that would be lost in a group-level calculation. South Africa, Nigeria, and Egypt dominate the AF14 average.

Comparator Country Selection

Brazil, China, and India represent large developing economies with significant but structurally different exposure to maritime trade costs. The SIDS aggregate captures economies that feature prominently in the NZF's equity architecture, including revenue disbursement design and beneficiary categories under regulation 41. Viet Nam, Bangladesh, and Trinidad and Tobago provide additional comparators for coalition dynamics. All comparators use the same computation chain applied to their respective CIA S24 baseline values.

Annex B: Scope, Boundaries, and Limitations of the Reduced-Form Approach

Economic Channels Not Captured

The model captures two transmission channels: the compliance cost shock (Cluster 1) and the fiscal offset (Cluster 2). Several additional channels are not modelled. In each case, the direction of likely bias is noted.

Exchange rate and balance-of-payments effects. For small, open economies with managed exchange rates—notably Ghana, Egypt, Kenya, and Ethiopia among the AF5—a persistent increase in import costs constitutes a balance-of-payments shock that may trigger exchange rate adjustment, amplifying the domestic price impact beyond the direct cost-push channel. The GTAP model's standard closure holds the trade balance, so this channel is largely absent from both the CIA results and the reduced-form scaling. For Egypt, where the parallel exchange rate has been a recurring policy concern, and for Ethiopia, where foreign exchange constraints are binding, this omission is particularly relevant. **Direction of bias: understatement of impact.**

Terms-of-trade effects on export competitiveness. Countries whose exports compete with suppliers facing different maritime cost exposure may experience trade diversion effects not captured by the aggregate SR scaling. This is most relevant for South Africa, whose manufacturing exports compete with Asian suppliers. The GTAP model captures terms-of-trade effects internally, but the SR scaling treats them as proportional to the aggregate MLC, which may not hold for countries with distinctive export profiles. **Direction of bias: uncertain, country-specific.**

Labour market and sectoral concentration. Maritime logistics cost increases produce concentrated effects in port cities and logistics hubs—Tema in Ghana, Mombasa in Kenya, Durban in South Africa. The model produces a single GDP impact per country and cannot address who within a country bears the cost, a question that may be more politically salient than the aggregate number. **Direction of bias: not applicable (distributional, not aggregate).**

Strategic behaviour in shipping markets. The model treats the MLC increase as a cost-push transmitted mechanically to freight rates. Freight rates are set in oligopolistic markets where carriers may absorb or amplify cost shocks strategically. Evidence from the 2021–2023 container shipping disruptions suggests carriers amplified cost shocks rather than passing them through proportionally. **Direction of bias: likely understatement of consumer price impact.**

Structural changes in the trade environment. The GTAP baseline predates or only partially reflects the African Continental Free Trade Area (AfCFTA). Over the 2030–2050 horizon, the AfCFTA is expected to shift intra-African trade patterns substantially, altering the exposure profiles underlying the CIA results. The SR scaling implicitly assumes relative country exposure does not change with trade structure—an assumption that weakens over longer horizons. **Direction of bias: uncertain, increasing with time horizon.**

Temporal Boundaries

Several model features become less defensible as the horizon extends from 2030 toward 2050.

Compositional effects. The SR captures aggregate MLC variation but not its composition. In GTAP, route-level and commodity-level cost incidence matters. The NZF's two-tier structure, reward mechanism, and SU market may change route-level incidence in ways the aggregate SR does not capture. This is most defensible at 2030 (where RU payments are uniform across vessel types) and less so at 2040–2050.

Cluster independence. The strict independence of Clusters 1 and 2 is defensible at 2030 but increasingly conservative at later horizons. By 2040–2050, in-sector investments funded through disbursements could reduce compliance costs by lowering the marginal abatement cost curve. The model provides no mechanism for this feedback. **Direction of bias at later horizons: overstatement of net impact for countries receiving substantial in-sector investment.**

Cross-scenario convergence. All CIA scenarios converge to 34.7–36.8% MLC by 2050, implying design choices become progressively less consequential. The model's slider-responsive 2050 results should be interpreted in this context.

Visual interpolation. The graph lines connecting 2030–>2040–>2050 in the calculator are visual interpolations between three modelled snapshot years, not modelled trajectories. Users should not draw conclusions about the time path between years, as the actual trajectory may be non-linear (e.g., compliance costs may spike during the 2033–2037 GFI tightening phase).

Interpretive Considerations

Reward distortion 1:1 assumption. The reward distortion formula assumes each reward dollar induces approximately one additional dollar of system fuel cost. This is conservative at 2030, where supply inelasticity in early-stage e-fuel markets may amplify the per-dollar cost effect, and aggressive at later horizons, where growing inframarginal ZNZ adoption reduces the marginal distortion.

GDP-weighting of group averages. The AF5 and AF14 are GDP-weighted, meaning South Africa dominates AF5, and South Africa, Nigeria, and Egypt dominate AF14. For users interested in smaller, more vulnerable economies, individual country results (always available in the calculator) are more informative. Population-weighted or unweighted averages would produce different summary statistics and might better serve the equity framing central to the NZF debate.

Fund delivery toggle. The binary toggle ($\omega_{effective}$ = 0 at 2030) represents the extreme case of a completely non-operational Fund. In practice, the Fund could be partially operational through interim arrangements. A graduated scale (25%, 50%, 75% of design intent) would better represent the continuum, but the binary toggle was chosen for negotiation clarity: the primary message—that a non-operational Fund eliminates the offset entirely—is more actionable than a continuous distribution.

Annex C: Derivation of the Two-Tier Adjustment Factor (α)

Regulatory Parameters

Verified against the draft revised MARPOL Annex VI (Circular Letter No. 5005).

GFI targets (2030): Base GFI = 85.8 gCO_2e/MJ (8% reduction from 93.3 reference); Direct GFI = 73.7 gCO_2e/MJ (21% reduction).

RU prices (2028–2030): Tier 1: \$100/tCO_2e (deficit between Direct and Base GFI). Tier 2: \$380/tCO_2e (deficit above Base GFI). Post-2030 prices to be determined by the Committee by 1 January 2028.

Fuel GFI values: VLSFO WtW = 91.4 gCO_2e/MJ. Fossil LNG WtW = 77.5 gCO_2e/MJ. ZNZ threshold: ≤19 gCO_2e/MJ through 2034.

Fleet Composition

The derivation assumes approximately 83% conventional (VLSFO-fuelled, attained GFI = 91.4) and 17% LNG-fuelled (attained GFI = 77.5) ships by energy consumption in 2030. This is consistent with CSC fleet projections (ISWG-GHG 20/2/31) and CIA Task 2 fleet modelling.

Derivation

Under the NZF, the compliance cost per million MJ of fuel consumed differs by ship type and tier:

Conventional ships (83% of fleet energy):

- Tier 1 gap: Base GFI − Direct GFI = 85.8 − 73.7 = 12.1 gCO_2e/MJ
- Tier 1 cost: 12.1 × \$100 = \$1,210/M·MJ
- Tier 2 gap: Attained GFI − Base GFI = 91.4 − 85.8 = 5.6 gCO_2e/MJ
- Tier 2 cost: 5.6 × T_2 (variable)
- Total NZF cost: \$1,210 + 5.6 × T_2

LNG ships (17% of fleet energy):

- LNG attained GFI (77.5) is below Base GFI (85.8), so no Tier 2 exposure
- Tier 1 gap: 12.1 gCO_2e/MJ (same for all ships)
- Total NZF cost: \$1,210/M·MJ (Tier 1 only)

Fleet-weighted NZF cost: > $0.83 \times (1210 + 5.6 \times T_2) + 0.17 \times 1210 = 1210 + 0.83 \times 5.6 \times T_2 = 1210 + 4.65 \times T_2$

Under S24, the single-tier clearing price is approximately \$150/tCO_2e (midpoint of DNV's \$140–160 range, CIA Task 2):

S24 fleet-weighted cost:

- Conventional: compliance gap = 91.4 − 73.7 = 17.7 gCO_2e/MJ; cost = 17.7 × \$150 = \$2,655/M·MJ
- LNG: compliance gap = 77.5 − 73.7 = 3.8 gCO_2e/MJ; cost = 3.8 × \$150 = \$570/M·MJ
- Fleet-weighted: 0.83 × 2655 + 0.17 × 570 = 2,204 + 97 = \$2,301/M·MJ

Therefore:

$\alpha(T_2) = (1210 + 4.65 \times T_2) / 2301$

Structural Asymmetries

Two features warrant emphasis.

First, LNG ships face no Tier 2 exposure but their Tier 1 cost (\$1,210/M·MJ) substantially exceeds their S24 cost (\$570/M·MJ). LNG ships individually face higher costs under the NZF than under S24—a structural asymmetry from the universal Tier 1 charge applied to the full Base-to-Direct gap regardless of fuel choice.

Second, for conventional ships, the UCL/UMAS (2025) analysis finds that the near-term least-cost compliance option is paying both Tier 1 and Tier 2, because biofuel abatement costs (~\$484/tCO_2e for bio-MGO) exceed the RU2 price of \$380. Conventional ships therefore face NZF costs of \$3,338/M·MJ versus \$2,655 under S24.

Post-2030 Convergence

The two-tier adjustment converges to α = 1.00 at 2040 and 2050 for three reasons: (a) post-2030 RU prices are undetermined and may be adjusted upward; (b) GFI stringency tightens from 8% to 65% reduction (2040) and net zero (2050), shifting compliance cost from RU payments to fuel switching; (c) fleet decarbonisation eliminates the ship-level GFI deficits that generate the tiered pricing advantage.

Sensitivity to S24 Clearing Price

The S24 clearing price of \$150/tCO_2e is the midpoint of DNV's \$140–160 range. At \$140: α(380) = 1.37; at \$160: α(380) = 1.22. This ±7% variation in α is captured within the model's ±15% fuel price uncertainty band.

Sensitivity to Fleet Composition

The 83/17 conventional-LNG split by energy consumption is sourced from CSC fleet projections and CIA Task 2. Because changes in fleet composition shift both the numerator (fleet-weighted NZF cost) and denominator (fleet-weighted S24 cost) in the same direction, α is relatively insensitive to this assumption. At T_2 = \$380: with an 80/20 split, α = 1.30; at 83/17 (default), α = 1.29; at 85/15, α = 1.29. The ±1% variation is well within the fuel price uncertainty band and does not meaningfully affect downstream results.

Annex D: Reward Mechanism Parameters — Derivations and Evidence

Reward Calibration Efficiency (η): Supply-Curve Framework

The reward calibration efficiency captures the fraction of the reward budget that bridges actual cost gaps between ZNZ fuels and conventional alternatives, as distinct from the fraction flowing as windfall to inframarginal producers whose costs are already below the reward level.

Under a supply-curve framework, individual ZNZ production projects—each with unique delivered cost driven by local electricity prices, electrolyser procurement, financing terms, and logistics—form a merit-order curve when sorted by cost gap above the competitive reference price. The cost-gap distribution's coefficient of variation (CV) determines producer surplus under any single-price mechanism.

Between-pathway dispersion. UCL/UMAS (2025) Table 6 reports ZNZ-qualifying fuel abatement costs at 2030 ranging from $400/tCO$_2$e (bio-LNG) to $628/tCO$_2$e (e-methanol), producing cost gaps above RU2 ($380) ranging from $20 to $248—a 12:1 ratio.

Within-pathway dispersion. IRENA's geospatial green hydrogen supply-cost curves (IRENA 2022, Global Hydrogen Trade Part III) show production costs from ~$0.65/kg at best locations to ~$1.15/kg under less optimistic assumptions, implying within-pathway CVs of 25–35%. IRENA (2020, Green Hydrogen Cost Reduction) documents best-versus-average production cost ratios of approximately 2.5:1, driven by compounding variation in electricity cost, electrolyser CAPEX, efficiency, and financing.

Blended dispersion. For e-fuels-only eligibility (the direction of travel in IMO discussions, per IMarEST 20/2/27 ¶56), the blended CV of the project cost-gap distribution is approximately 35–45%.

η by Mechanism Design

Mechanism	η	Rationale
Flat-rate reward, broad eligibility	0.55	Price set to attract ~90th-percentile supply at CV ~40%; substantial surplus to cheaper projects
Uniform-price auction / calibrated flat rate	0.60	Auction discovers price but creates identical surplus geometry; advantage is price discovery only
Category-specific or flexible ex-post	0.70	Reduced within-category dispersion; some dynamic adjustment
Pay-as-bid auction with JET set-asides	0.80	Discriminatory pricing; set-aside pools reduce competition –> 15–25% bid-shading
Competitive pay-as-bid auction	0.88	Deep bidder pool; 10–15% bid-shading (IRENA 2015, 2019)

The lower bound (0.55) is set by practical plausibility: a mechanism where >45% of the budget is wasted as windfall would be abandoned politically. The upper bound (0.88) reflects the efficiency ceiling of real-world pay-as-bid auctions; no first-generation international maritime mechanism would exceed mature domestic renewable energy procurement performance.

Bid-shading evidence. Empirical evidence from renewable energy auctions (IRENA 2015, Renewable Energy Auctions: A Guide to Design; IRENA 2019, Auctions: Status and Trends Beyond Price) indicates winning bids typically include 10–25% markup over true costs in competitive multi-round auctions, producing $\eta \approx 0.88$. With JET set-aside pools reducing competition, $\eta \approx 0.80$.

Market Structure Assumption

The reward distortion formula applies η on the assumption that freight markets exhibit asymmetric pass-through: operators pass through fuel cost increases but retain reward-generated windfalls as margin. Under this assumption, only the fraction η that induces actual fuel switching enters MLC; the remainder $(1 - \eta)$ is retained by operators and is MLC-neutral.

Under perfectly competitive freight markets, operators would pass through both cost increases and windfall savings symmetrically, and the full reward budget would enter MLC regardless of η. Reality lies between these extremes. The asymmetric assumption is consistent with evidence on oligopolistic carrier behaviour (§3.2, Annex B) and the CIA's implicit treatment of freight rate formation.

The magnitude of uncertainty: at $\sigma = 0.50$, the distortion ranges from 0.344% of MLC ($\eta = 0.55$, asymmetric) to 0.625% (competitive, η irrelevant)—a spread of ~0.28 percentage points, small relative to the base compliance cost of ~7.5%.

Fuel Scope Factor (φ): Derivation

Under a fixed reward budget, restricting fuel eligibility increases the per-unit cost gap between qualifying ZNZ fuel and the conventional alternative, but decreases the volume of ZNZ fuel purchased per reward dollar. Under efficient price discovery, these effects largely offset. The net cost increase from the most restrictive policy (e-fuels only) is bounded by the ratio of cost gaps: using UCL/UMAS 2030 data, the e-ammonia cost gap exceeds the bio-MGO gap by ~5–10%. Under efficient price discovery, the reward mechanism optimises fuel selection against the cost gap, so the scope factor applies only to the incremental cost of excluding cheaper alternatives. Effective range: 1.00–1.05, corresponding to 0.00–0.03 percentage points of MLC.

SU Market Factor (μ): Derivation

The 2030 SU market is supply-constrained: total SU supply from overcompliant ships (~15 $MtCO_2e$, from CIA Task 2) is well below Tier 2 demand from conventional ships (~44.8 $MtCO_2e$), a supply-to-demand ratio of ~0.33. SU prices converge toward the $380 ceiling regardless of design, because demand exceeds supply. Market friction (restrictive design, clearing delays) prevents full convergence, producing $\mu = 1.01$–1.02. SU multipliers for ZNZ ships create additional SU supply but function as unfunded additional rewards—the most expansive design (3× multiplier) adds ~6% to the reward distortion, giving $\mu = 1.06$. The direction is counterintuitive: expansive SU design is cost-increasing in a supply-constrained market.

Robustness to Wider Ranges

At analytically derived ranges, σ accounts for ~87% of total GDP impact variation for Kenya, with Sliders 3–5 contributing ~13% (ratio ~7:1). Under substantially wider ranges (φ extended to 1.30, μ to [0.88, 1.12]), σ still accounts for ~75%, with Sliders 3–5 at ~25% (ratio ~3:1). The dominance of σ is structural: it affects both channels simultaneously, while φ, η, and μ modify only a sub-component of one.

Usability Factor (ω): Calibration Evidence

The usability factor $\omega = \delta \times \lambda$ captures the GDP return per dollar of Fund disbursement to a recipient country. It is decomposed into two components: disbursement effectiveness δ (fraction of allocated revenue that reaches productive deployment) and the fiscal multiplier λ (GDP generated per dollar of effective spending). The five modality values in the model correspond to five instrument types spanning the range of climate finance delivery mechanisms. The calibration of each component draws on the evidence base described below.

Disbursement effectiveness (δ): evidence base. The primary reference is the Green Climate Fund's disbursement-to-approval performance across instrument types, as reported in GCF Annual Reports (2023, 2024) and the GCF-1 Progress Report (2023). Supplementary evidence is drawn from UNEP (2023) bilateral climate finance data and the Climate Analytics (2023) review of GCF's first five years. The key data points are:

Technical assistance ($\delta = 0.35$): GCF readiness and preparatory support grants show cumulative disbursement rates of 30–40% of approved amounts, reflecting high overhead (implementing entity fees, M&E costs, procurement delays) and long implementation timelines. UNEP bilateral climate finance data show similar patterns for TA-intensive modalities.

Earmarked loans ($\delta = 0.55$): Loan-based GCF projects achieve faster disbursement than grant-based equivalents but face reporting, procurement, and co-financing requirements that reduce effective delivery. The GCF-1 portfolio shows loan instruments at ~50–60% disbursement-to-approval.

Mixed instruments ($\delta = 0.65$): Blended finance approaches combining loans, grants, and equity. Set at the midpoint of the earmarked-loan and flexible-grant range.

Flexible in-sector grants ($\delta = 0.80$): Sector-specific grants with fewer conditionalities than TA but more oversight than budget support. GCF simplified approval process projects in the energy and transport sectors show 75–85% disbursement-to-approval for projects past mid-implementation.

Unrestricted budget support ($\delta = 0.90$): Direct transfers to national treasuries with minimal intermediation. Calibrated from bilateral budget support experience (World Bank Development Policy Lending assessments) where disbursement efficiency approaches 90–95% of approved amounts.

Fiscal multiplier (λ): evidence base. The primary references are the IMF fiscal multiplier estimates for emerging market and developing economies (EMDEs): Batini et al. (2014), IMF Technical Notes and Manuals 14/04, which reports spending multipliers of 0.2–0.6 for current expenditure and 0.4–1.0 for public investment in EMDEs; and Ilzetzki, Mendoza, and Végh (2013), who find developing-country output multipliers of 0.2–0.5 for current spending and somewhat higher for capital spending, with wide variation by institutional quality. For Sub-Saharan Africa specifically, Kimaro, Keong, and Sea (2017) estimate fiscal multipliers of 0.5–0.8 for productive government spending, with the higher end reflecting capital investment in infrastructure. The λ values are set as follows:

Technical assistance ($\lambda = 0.30$): TA spending generates GDP primarily through capacity building and institutional strengthening—indirect transmission with long lags. Consistent with the lower end of the Batini et al. current-spending range for LICs.

Earmarked loans ($\lambda = 0.40$): Loan-funded projects typically target specific sectors with moderate multiplier effects. Consistent with the midpoint of the Ilzetzki et al. developing-country range.

Mixed instruments ($\lambda = 0.50$): Blended approaches. Set at the transition between current-spending and capital-spending multiplier ranges in the Batini et al. framework.

Flexible in-sector grants ($\lambda = 0.60$): Sector grants enabling capital investment in energy and transport infrastructure. Consistent with the Kimaro et al. SSA productive-spending estimate and the upper range of Batini et al. for EMDE capital spending.

Unrestricted budget support ($\lambda = 0.70$): General budget support provides maximum fiscal flexibility, enabling governments to allocate to highest-return uses. The higher multiplier reflects both capital and recurrent spending at sovereign discretion. Consistent with the upper end of the Kimaro et al. SSA range and the higher Batini et al. capital-spending estimates.

Multiplicative form. The product $\omega = \delta \times \lambda$ assumes δ and λ are independent. In practice, the instrument choice that determines δ also influences the type of spending (and hence λ): budget support has both high δ and high λ, while technical assistance has both low δ and low λ. The five modality values bundle δ and λ together precisely because they are co-determined by instrument choice. The multiplicative form is retained as a transparent decomposition of two conceptually distinct questions (does the money arrive? does it generate growth?), not as a claim of statistical independence. The resulting ω range (0.10–0.63) is presented as a user-adjustable parameter because the true value depends on institutional choices not yet made.

Annex E: Revenue Estimation and Temporal Trajectory

First-Principles Revenue Estimate

NZF revenue at 2030 is derived from fleet-wide RU payments under the regulatory GFI gaps. At the 2030 horizon, Tier 1 payments are paid by essentially the entire fleet (the Direct-to-Base gap is unavoidable), and Tier 2 payments are paid by the ~83% of the fleet using conventional fuels. Revenue erosion from reward-induced ZNZ adoption is less than 1% at this horizon and is negligible.

The resulting estimate of ~$10 billion per year is corroborated by UCL/UMAS (2025): "initial estimates based on the parameters of the IMO's NZF indicate that $11–$12 billion could be raised annually between 2028 and 2030" (p. 8). The model uses $10 billion as a conservative round figure.

CIA Task 2 Structural Benchmark

CIA Task 2 revenue data (Table D-1, MEPC 82/INF.8/Add.1) provides an independent structural benchmark under S24's flexibility mechanism:

Period	S24 Revenue	Interpretation
2027–2030	$5.8B/year	Early phase; moderate compliance gaps
2031–2040	$8.2B/year	GFI tightening; larger gaps
2041–2050	$2.7B/year	Fleet decarbonisation; RU base collapses

The NZF generates higher revenue than S24 (~1.7× at 2030) because Tier 1 payments are universal—every ship pays for the Direct-to-Base gap regardless of abatement strategy—whereas S24's flexibility mechanism allows compliance trading that reduces total revenue.

Temporal Trajectory

The approximate NZF revenue trajectory: ~$10 billion at 2030; ~$8–12 billion at 2040 (uncertain, depending on fleet transition speed and undetermined post-2030 RU prices); ~$1–3 billion at 2050. The current model applies $10 billion uniformly.

At 2030, this is appropriate. At 2040, the true value is plausibly similar; the model may slightly understate or overstate depending on fleet transition. At 2050, the $10 billion assumption overstates actual revenue by a factor of approximately 3–10×, meaning the fiscal offset is substantially overstated.

Compounding Effects at 2050

The revenue decline at 2050 compounds the α convergence: the time-varying α increases the 2050 cost shock by ~22% (from the fixed-α assumption), while revenue decline reduces the fiscal offset by 70–90%. The net effect: 2050 GDP impacts are likely worse than shown. However, practical significance is limited: by 2050 all scenarios converge to 34.7–36.8% MLC, and the fiscal offset is already small relative to the cost shock (0.02% of GDP vs. 0.4–0.5%). The policy-relevant period for revenue recycling is 2028–2040.

Maritime Logistics Cost Estimate (M)

Total global maritime logistics cost is derived from CIA Task 3 data in two steps. First, maritime transport costs: CIA Table 13 reports the ad-valorem maritime transport margin by commodity and trade route at the GTAP base year (2017). The global trade-weighted average margin is 1.60% of CIF import value. Applied to global merchandise imports of ~$25 trillion (WTO, 2024), this yields approximately $400 billion in maritime transport costs. Second, total MLC: the CIA defines MLC as the sum of maritime transport costs and shipping time costs (CIA Section 4.4, equation 4.1). CIA Table 3 reports total bilateral maritime logistics costs by origin-destination pair; CIA Table 6 reports the breakdown between transport cost and shipping time cost components. The ratio of total MLC to transport cost alone, trade-weighted across all bilateral routes in CIA Table 6, is approximately 2:1. The multiplier reflects the high value-of-time coefficients for containerised manufactured goods, food, and metals, which comprise the majority of seaborne trade by value. Applying this ratio: $400 billion × 2 ≈ $800 billion per year in total MLC. The R/M ratio of ~1.25% ($10B / $800B) is therefore a derived quantity, not a calibrated parameter. Its sensitivity to M is modest: a ±20% variation in M (plausible given GTAP base year vintage) shifts R/M between 1.04% and 1.56%, producing a ±0.06 percentage point variation in the reward distortion—small relative to the fuel price uncertainty band.

Annex F: Cross-Scenario Validation and Robustness

Full Cross-Scenario Validation Table

The CIA provides a natural empirical test of the proportional scaling assumption. Nine of ten S24-adjacent scenarios have distinct MLC values at 2030, spanning 5.78% to 19.09%.

Table F.1: Cross-Scenario Validation — All AF5 Countries, 2030

Panel A: Ghana

Scenario	MLC (%)	SR	Predicted	Actual	Error
Within range					
23	6.32	1.09	−0.074	−0.072	−3%
31	6.38	1.10	−0.075	−0.072	−4%
32	7.26	1.26	−0.085	−0.082	−4%
36	7.46	1.29	−0.088	−0.084	−4%
Outside range					
43	10.83	1.87	−0.127	−0.105	−18%
46	19.09	3.30	−0.225	−0.210	−7%
Interpolated					
NZF default	8.03	1.39	−0.095	−0.088*	−7%

Panel B: Kenya

Scenario	MLC (%)	SR	Predicted	Actual	Error
Within range					
23	6.32	1.09	−0.147	−0.147	0%
31	6.38	1.10	−0.148	−0.139	−6%
32	7.26	1.26	−0.168	−0.150	−11%
36	7.46	1.29	−0.173	−0.167	−3%
Outside range					
43	10.83	1.87	−0.251	−0.195	−22%
46	19.09	3.30	−0.443	−0.361	−18%
Interpolated					
NZF default	8.03	1.39	−0.186	−0.174*	−7%

Panel C: South Africa

Scenario	MLC (%)	SR	Predicted	Actual	Error
Within range					
23	6.32	1.09	−0.077	−0.076	−1%
31	6.38	1.10	−0.077	−0.075	−3%
32	7.26	1.26	−0.088	−0.082	−7%
36	7.46	1.29	−0.090	−0.088	−3%
Outside range					
43	10.83	1.87	−0.131	−0.102	−22%
46	19.09	3.30	−0.231	−0.190	−18%
Interpolated					
NZF default	8.03	1.39	−0.097	−0.092*	−6%

Panel D: Egypt

Scenario	MLC (%)	SR	Predicted	Actual	Error
Within range					
23	6.32	1.09	−0.097	−0.086	−12%
31	6.38	1.10	−0.098	−0.096	−2%
32	7.26	1.26	−0.112	−0.110	−2%
36	7.46	1.29	−0.115	−0.114	−1%
Outside range					
43	10.83	1.87	−0.167	−0.131	−21%
46	19.09	3.30	−0.294	−0.241	−18%
Interpolated					
NZF default	8.03	1.39	−0.124	−0.118*	−4%

Panel E: Ethiopia

Scenario	MLC (%)	SR	Predicted	Actual	Error
Within range					
23	6.32	1.09	−0.039	−0.036	−9%
31	6.38	1.10	−0.040	−0.037	−7%
32	7.26	1.26	−0.045	−0.041	−9%

Scenario	MLC (%)	SR	Predicted	Actual	Error
36	7.46	1.29	−0.046	−0.047	+1%
Outside range					
43	10.83	1.87	−0.067	−0.051	−24%
46	19.09	3.30	−0.119	−0.098	−18%
Interpolated					
NZF default	8.03	1.39	−0.050	−0.048*	−3%

Notes: GDP impacts in percentage points. Ethiopia = "Rest of landlocked economies in Africa." Asterisked values interpolated linearly between Scenarios 36 (SR = 1.29) and 43 (SR = 1.87). Scenario 26 excluded (see data note below). Scenarios 21 and 22 omitted for space; results consistent with pattern shown.

Summary Statistics

Within the core validation range SR in [0.85, 1.29]: mean prediction error −4.4%, 90th-percentile absolute error 9.5%. At the default operating point (SR = 1.39, interpolated): mean error −5.3%, 90th-percentile absolute error 6.6%. Direction consistently conservative (overstatement of impact magnitude).

At larger deviations (SR = 1.87, 3.30): prediction error increases to 18–24%. The relationship between MLC and GDP is concave: substitution effects and GE dampening reduce the marginal GDP response at higher cost levels. The proportional scaling has a well-defined validity boundary; the operating range SR in [0.85, 1.45] is inside it.

Affine Specification Robustness Check

A linear-with-intercept specification (GDP = a + b × MLC) achieves R^2 of 0.97–0.99 for all AF5 countries, compared to 0.74–0.95 for the proportional (zero-intercept) specification. The non-zero intercept reflects GDP effects that do not scale with MLC—structural and terms-of-trade adjustments operating independently of the cost shock.

At SR = 1.39, the affine model produces GDP predictions 6–10% less negative than proportional: Ghana −0.089% (vs. −0.095%), Kenya −0.168% (vs. −0.186%), South Africa −0.089% (vs. −0.097%), Egypt −0.113% (vs. −0.124%), Ethiopia −0.045% (vs. −0.050%).

The proportional specification is retained because: (a) the bias is conservative; (b) it preserves a simpler computation chain; (c) the affine predictions are available here as a robustness check. A quadratic specification captures the concavity more precisely (R^2 = 0.99 for all countries) but adds complexity without material improvement within the operating range.

Data Note: CIA Scenario 26

CIA Task 3 Annex 5 (Scenario 26) is excluded. The country-level GDP values reported are identical to Annex 4 (Scenario 24) despite Scenario 26 having an aggregate MLC of 15.82% versus S24's

5.78%. The group-level data in the main report tables shows the expected differentiation, indicating a transcription error in the country-level annex. This does not affect any result in the analysis.

AFPH Cross-Check

The AFPH (2024) study of maritime carbon pricing impacts on African economies provides an independent cross-check using GTAP-E (the energy-environmental extension of the same GTAP model and GTAP 11 database used in the CIA) with Ghana individually represented as one of 17 regions, applying a \$100/tCO_2 carbon tax to water transport.

The AFPH Ghana GDP impact (−0.085%) and the CIA S24 Ghana impact (−0.068%) are not directly comparable (different policy instruments, cost magnitudes, time horizons, and regional aggregation). However, the result provides directional corroboration: a larger cost shock produces a larger impact; both place Ghana in the moderate-impact range; both rely on the same GTAP trade structure.

Three additional observations from the AFPH analysis are relevant:

1. **Full pass-through on African routes.** AFPH stakeholder consultations found shipping lines implement full cost pass-through on African routes, with importers bearing the entire surcharge incidence. This is consistent with the CIA's implicit pass-through assumptions.
2. **South Africa positive GDP.** The AFPH finds South Africa's GDP increases (+0.009%) under the carbon tax—driven by GE rebalancing where manufacturing exports gain competitiveness. This is the same class of counterintuitive effect producing negative CPI values for Kenya and Egypt in the CIA, providing evidence that these are structural model features, not artefacts.
3. **Ranking divergence.** AFPH places Ghana as more affected than Kenya (−0.085% vs. −0.019%); the CIA reverses this (−0.068% vs. −0.134%). The divergence reflects differences in policy transmission, regional aggregation, and Kenya's GE-driven negative CPI. This underscores that country-level results are indicative of orders of magnitude rather than precise point estimates.

CPI Response Pattern Analysis

Cross-scenario analysis of six CIA no-disbursement scenarios reveals two distinct patterns:

Proportional responders (majority of countries including Ghana, Nigeria, Algeria, Morocco): CPI scales roughly proportionally with MLC across all scenarios, validating SR-based scaling.

Negative-CPI countries (Kenya, Egypt, India, Brazil, Bangladesh): CPI is negative across all scenarios, driven by GE effects (trade rebalancing, terms-of-trade shifts) exceeding the direct cost-push. These values do not scale proportionally with MLC. Negative values are floored at +0.001% for display and scaling to prevent artefactual results (scaling a negative base by SR would reverse the sign for SR < 1). CPI results for these countries reflect the direct cost-push channel only; the GE rebalancing benefit is noted but not quantified.

Food Price Decomposition: Data Sources and Methodology

The food CPI component is derived as:

$food\_share_i$ = (seaborne_food_imports$_i$ × margin_food) / (total_maritime_imports$_i$ × margin_all)

$food\_CPI_i = (CPI_i \times food\_share_i) / food\_weight_i$

where margin_food reflects higher maritime margins on agricultural commodities (9.39% for raw agriculture, 5.32% for processed food) compared to the weighted average (2.50%), from CIA Table 8.

CPI food weights (COICOP Division 01, national statistical office publications): Ghana 43% (GSS 2023 rebasing); Kenya 33% (KNBS Feb 2019 base); South Africa 17% (Stats SA 2022 base); Egypt 33% (CAPMAS 2019/20 base); Ethiopia 43% (CSA, estimated). Base years vary; measurement uncertainty is approximately 3–5 percentage points for countries that have not rebased recently.

Seaborne food import shares derived from WDI series TM.VAL.FOOD.ZS.UN (Food imports as % of merchandise imports, UN Comtrade, World Bank WDI database accessed January 2026) combined with GDP and merchandise import data.

Trade Cost ρ Derivation Methodology

The import composition adjustment ratio ρ_i is derived from each country's HS 2-digit import composition (UN Comtrade, most recent year available per country) mapped to GTAP commodity groups using a standard HS-GTAP concordance. The effective maritime margin is the import-share-weighted average of group-specific margins (CIA Table 12). Comtrade vintage years vary (2017–2025); import commodity composition is structurally stable for the sample economies, with year-to-year ρ variation typically <±0.03.

Annex G: Calculator Specification and Worked Example

This annex provides the complete slider specifications for the NZF Economic Impact Calculator and a worked calculation example. It supersedes Parameter Decision Memo v4.6.

Complete Slider Specifications

Slider 0: Fund Revenue Split

Parameter: σ (reward_budget_share). Range: 0.20–0.80. Most consequential slider—drives both reward distortion and disbursement pool.

Computational role: σ enters both the cost channel and the fiscal offset channel in opposite directions. In the cost channel, reward_distortion = $(R / M) \times 100 \times \sigma \times \eta \times \varphi \times \mu$ — higher σ increases compliance costs. In the fiscal offset channel, disbursement_pool = $R \times (1 - \sigma)$ — higher σ shrinks the pool available for State disbursements. This double effect makes σ the most consequential single parameter. At σ = 0.80, 80% of Fund revenue goes to ZNZ rewards (maximising reward distortion and minimising the fiscal offset). At σ = 0.20, the reverse.

Position	Label	σ
1	Heavily toward rewards	0.80
2	Majority toward rewards	0.65

Position	Label	σ
3	Balanced	0.50
4	Majority toward disbursements	0.35
5	Heavily toward disbursements	0.20

Slider 1: Disbursement Allocation

Parameter: allocation_scheme (categorical, CIA Table 23 formula). Fiscal/GDP channel only.

Computational role: Determines which countries are eligible for Fund disbursements and how the disbursement pool is divided among them. Countries outside the eligible set receive zero fiscal offset ($\theta\_i = 0$).

Allocation formula (positions 1–4): alloc_share_i = (pop_i × |S24_GDP_impact_i|) / Σ_eligible(pop_j × |S24_GDP_impact_j|). Per-capita disbursements are proportional to the magnitude of each country's CIA Scenario 24 GDP impact. Countries with larger populations and larger GDP impacts receive larger shares. Position 5 uses population-only weighting: alloc_share_i = pop_i / Σ_eligible(pop_j).

Eligible set by position: Position 1 (SIDS & LDCs only): country is classified as SIDS or LDC. Non-LDC, non-SIDS developing countries (e.g. Kenya, Ghana, South Africa, Egypt) receive zero disbursement; Ethiopia (LDC) receives a substantial share. Position 2 (Vulnerable states): SIDS and LDCs, plus any developing country whose absolute S24 2030 GDP impact exceeds the median absolute GDP impact among all developing countries in the panel. This is the first position where non-LDC African countries may qualify. Position 3 (All developing states): all countries with status = ”developing”. Position 4 (All GDP-weighted): all countries including developed; developed countries enter the denominator, diluting per-capita shares for developing countries. Position 5 (All equal per capita): all countries, population-only weighting; GDP impact magnitude ignored.

Proxy country handling: Countries represented by GTAP regional aggregates (e.g. Ethiopia via ”Rest of landlocked Africa”) inherit the parent aggregate's allocation share, scaled by the ratio of the country's population to the aggregate's population. Proxy countries are excluded from the denominator to avoid double-counting.

Position	Label	Description
1	SIDS & LDCs only	Concentrated to most vulnerable
2	Vulnerable states	Expanded vulnerability tier
3	All developing states	CIA Table 23, all developing
4	All states (GDP-weighted)	GDP-inverse weighting
5	All states (equal per capita)	Equal per-capita distribution

Slider 2: Disbursement Modality

Parameter: $\omega = \delta \times \lambda$. Range: 0.10–0.63. Fiscal/GDP channel only.

Computational role: Multiplier applied to a country's share of the disbursement pool to determine the effective fiscal offset: $\theta_i = (r_i / GDP_i) \times \omega$, where $r_i = alloc\_share_i \times disbursement\_pool$. Higher ω means each dollar of disbursement translates into more GDP offset — representing more effective delivery instruments with higher fiscal multipliers. ω is decomposed as $\delta \times \lambda$, where δ is disbursement effectiveness (fraction of allocated revenue reaching productive deployment, calibrated from GCF data) and λ is the fiscal multiplier (GDP generated per dollar of effective spending, from IMF estimates for EMEs and LICs).

Position	Label	δ	λ	ω
1	Technical assistance only	0.35	0.3	0.10
2	Earmarked loans	0.55	0.4	0.22
3	Mixed instruments	0.65	0.5	0.33
4	Flexible grants (in-sector)	0.80	0.6	0.48
5	Unrestricted budget support	0.90	0.7	0.63

Slider 3: Fuel Eligibility

Parameter: φ (fuel_scope_factor). Range: 1.00–1.05.

Computational role: Multiplier in the reward distortion formula: $reward\_distortion = (R / M) \times 100 \times \sigma \times \eta \times \varphi \times \mu$. $\varphi > 1.0$ indicates that restricting eligible fuel pathways, or adding multipliers for preferred pathways, increases the effective reward cost by directing funds toward more expensive fuel options. The non-monotonic position ordering (1.05, 1.03, 1.00, 1.01, 1.02) reflects that both restriction (positions 1–2) and differentiation (positions 4–5) increase costs relative to the fuel-agnostic baseline (position 3), but through different mechanisms.

Position	Label	φ	Rationale
1	Highly selective	1.05	E-fuels only; some pool underutilisation
2	Somewhat selective	1.03	E-fuels + advanced biofuels
3	Agnostic (baseline)	1.00	Any fuel meeting ZNZ threshold
4	Somewhat differentiated	1.01	All eligible; mild multipliers
5	Strongly differentiated	1.02	All eligible; strong e-fuel multipliers

Slider 4: Reward Calibration

Parameter: η (reward_calibration_efficiency). Range: 0.55–0.88.

Computational role: Multiplier in the reward distortion formula: $reward\_distortion = (R / M) \times 100 \times \sigma \times \eta \times \varphi \times \mu$. Higher η means more of the reward budget translates into real abatement (which is costly), producing higher freight cost distortion. Lower η means more of the reward budget becomes operator windfall (inefficient but cheaper for consumers). Anchored to UCL/UMAS (2025) Table 6 cost-gap data.

Position	Label	η	Rationale
1	Flat-rate, broad eligibility	0.55	Fixed price; wide dispersion; max surplus
2	Uniform-price auction	0.60	Discovers price; same structural surplus
3	Category-specific / flexible	0.70	Reduced within-category dispersion
4	Pay-as-bid with JET set-asides	0.80	Discriminatory pricing; 15–25% bid-shading
5	Competitive pay-as-bid auction	0.88	Deep bidder pool; 10–15% bid-shading

Slider 5: SU Market Design

Parameter: μ (su_market_factor). Range: 1.00–1.06.

Computational role: Multiplier in the reward distortion formula: reward_distortion = (R / M) × 100 × σ × η × φ × μ. Under supply-constrained SU market conditions, SU prices converge toward the $380 ceiling regardless of design. Market friction (restrictive design) prevents full convergence, adding a small cost premium. Expansive design with ZNZ multipliers creates additional SU supply but functions as unfunded additional rewards — cost-increasing in a supply-constrained market. The non-monotonic position ordering (1.02, 1.01, 1.00, 1.03, 1.06) reflects that both restrictive and expansive designs are cost-increasing relative to baseline, through different mechanisms.

Position	Label	μ	Rationale
1	Restrictive	1.02	Friction prevents efficient clearing
2	Somewhat restrictive	1.01	Some friction
3	Baseline (open, no multiplier)	1.00	Cost-minimising configuration
4	Somewhat expansive	1.03	Mild ZNZ multiplier
5	Expansive (strong multiplier)	1.06	Strong multiplier; significant ZNZ pull-in

SU / Tier 2 Compliance Price

Parameter: T_2 (derives α via the formula in §2.4.2 / Annex C). Range: $100–$380. Default: $380.

Computational role: Determines the two-tier adjustment factor α, which scales the base compliance cost at 2030: $\alpha(T_2) = (1210 + 4.648 \times T_2) / 2300.55$. At T_2 = $380 (UCL/UMAS base case, SU supply-constrained at RU2 ceiling): $\alpha \approx 1.29$. At T_2 = $100 (low SU demand scenario): $\alpha \approx 0.73$. α applies only at 2030; at 2040 and 2050, α = 1.0 (fleet convergence eliminates the two-tier effect).

Fund Delivery Delay Toggle

When active: $\omega_{\text{effective}} = 0$ at 2030; 2040 and 2050 unchanged.

Computational role: At 2030, the fiscal offset term θ_i is set to zero regardless of slider settings. The cost shock (GDP_cost) is unaffected. This represents the scenario where the NZF is not operationally disbursing funds at 2030. The effect is largest when σ is low (high disbursement share) and ω is high (effective instruments) — precisely the configuration most favourable to developing countries.

Computation Chain

Step	Formula
1. Base MLC	$\text{base\_MLC} = \Phi_{S24}(t) \times \alpha(T2) \times \pi$
2. Reward distortion	$\text{reward\_dist} = (R / M) \times \sigma \times \eta \times \varphi \times \mu$
3. Total MLC	$\text{total\_MLC} = \text{base\_MLC} + \text{reward\_dist}$
4. Sensitivity ratio	$SR = \text{total\_MLC} / \Phi_{S24}(t)$
5. Cost shock	$\text{GDP\_cost}_i = \Phi^{GDP}{}_{S24,i}(t) \times SR$
6. Fiscal offset	$\text{pool} = R \times (1 - \sigma)$; $r_i = \text{alloc\_rate}_i \times \text{pool}$; $\theta_i = (r_i / \text{GDP}_i) \times \omega$
7. Net GDP	$\text{net\_GDP}_i = \text{GDP\_cost}_i + \theta_i$

Note: α(T2) is applied at 2030 only; α = 1.00 at 2040 and 2050.

Worked Example: Kenya, 2030

Configuration: All sliders at position 3 (σ = 0.50, all developing, ω = 0.33, η = 0.70, φ = 1.00, μ = 1.00).

Step	Calculation	Result
Base MLC	5.78% × 1.294 × 1.00	7.478%
Reward distortion	(10/800) × 0.50 × 0.70 × 1.00 × 1.00	0.438%
Total MLC	7.478% + 0.438%	7.916%
Sensitivity ratio	7.916% / 5.78%	1.369
Kenya cost shock	−0.134% × 1.369	−0.183%
Disbursement pool	$10B × 0.50	$5.0B
Kenya allocation	2.201% × $5.0B	$110M
Revenue / GDP	$110M / $130B	0.085%
θ (offset)	0.085% × 0.33	+0.028%
Net GDP impact	−0.183% + 0.028%	−0.156%
Offset ratio	0.028 / 0.183	15.3%

Full Range (Kenya, 2030)

Scenario	Cost Shock	Offset	Net	Offset %
Extreme A (max rewards)	−0.187%	+0.000%	−0.187%	0.0%
Balanced (central)	−0.183%	+0.028%	−0.156%	15.3%
Extreme B (max disbursements)	−0.179%	+0.029%	−0.150%	16.2%

Fixed Parameters

Parameter	Value	Source
S24 MLC anchor	5.78% / 26.06% / 35.52%	CIA Task 3, Table 8
Default T2	$380 –> α = 1.294	UCL/UMAS (2025)
Fuel price factor	1.00 (0.85–1.15)	DNV sensitivity
Base revenue	~$10B/yr (2030)	UCL/UMAS corroborated
Total MLC	~$800B/yr	CIA Table 13 + ratio
μ_{invest}	0.8	Reserved for future use

Sources

Draft Revised MARPOL Annex VI (Circular Letter No. 5005). UCL/UMAS (2025) compliance cost analysis. CIA Task 2 (DNV) fleet modelling and fuel price sensitivity. CIA Task 3 (UNCTAD) GTAP modelling and disbursement formula. CSC (ISWG-GHG 20/2/31) fuel emission factors.

Fiscal multiplier (λ): Batini, N., et al. (2014). "Fiscal Multipliers: Size, Determinants, and Use in Macroeconomic Projections." IMF Technical Notes and Manuals 14/04. Ilzetzki, E., Mendoza, E.G., & Végh, C.A. (2013). "How big (small?) are fiscal multipliers?" Journal of Monetary Economics 60(2), 239–254. Kimaro, E.L., Keong, C.C., & Sea, L.L. (2017). "Government Expenditure, Efficiency, and Economic Growth." African Journal of Economic Review 5(2), 34–54. Calderón, C. & Servén, L. (2014). "Infrastructure, Growth, and Inequality." World Bank Policy Research WP 7034.

Disbursement effectiveness (δ): Climate Analytics (2023). "Five years of the Green Climate Fund." Green Climate Fund (2023, 2024). Annual Reports. GCF-1 Progress Report (2023). UNEP (2023). Adaptation Gap Report 2023. ODI (2022). "Fiscal Multipliers." SET Programme.

Reward calibration efficiency (η): UCL/UMAS (2025). IRENA (2022). Global Hydrogen Trade Part III. IRENA (2020). Green Hydrogen Cost Reduction. IRENA (2024). Renewable Power Generation Costs 2023. IRENA (2015). Renewable Energy Auctions: A Guide to Design. IRENA (2019). Renewable Energy Auctions: Status and Trends Beyond Price.

Annex H: NZF Design Topic Inventory

This annex documents the 48 NZF design topics identified from the draft revised MARPOL Annex VI, the report of ISWG-GHG 20, and pre-session submissions to MEPC 84. Five topics are not independently classified in Annex I because their economic effects are fully subsumed by parent topics; they are included here for completeness. For each topic, the current boundary of proposals, the degree of convergence among delegations, and the specific design choices under consideration are summarised. The classification of each topic into Buckets X, Y, or Z is provided in Annex I.

Topic Grouping 1: IMO Net-Zero Fund

Topic 1.1: Scope of "governing provisions" and sequencing of what gets written first
Negotiation issue: whether the Fund's governing provisions should be developed as a complete "package" vs. prioritizing a subset (especially the Board provisions) first.

- **Convergence:** UCL reports broad agreement that a *full* set of governing provisions is ultimately required, but the **Governing Board** provisions are widely viewed as the **first priority** (because the Board will run the Fund).
- **Boundaries:**
  - **Board-first / core-first** approach (write core provisions early; let the Board develop derivatives later).
  - **Package approach** emphasizing interlinkages and adopting the provisions "as a package."

Current status: "Board-first" is the prevalent position, but procedural uncertainty remains regarding whether Board rules can be prioritised without the broader institutional machinery in place.

Topic 1.2: Timeline to operationalize the Fund and appoint the Governing Board
Negotiation issue: *how quickly* to adopt governing provisions and stand up the Board relative to when revenues start (expected from 2029).

- **Convergence:** "Many delegations called for a faster timeline" for adopting core provisions and appointing the Board.
- **Boundaries (range of proposals):**
  - **Accelerated timeline** with core provisions adopted as early as **MEPC 85 (late 2026)**, Board appointed/operational in **2027**, to be ready well before 2029.
  - A similar "prudent" target of adopting governing provisions by **MEPC 85 (2H 2026)** to allow setup time.
  - **Slower sequencing** implicit in the Secretariat workplan referenced by Fiji et al. (risk of provisions only at **MEPC 87 (spring 2028)**), which they argue is too late.

Current status: clear directional push toward earlier adoption, but the implementation pathway is linked to the interim-arrangements debate (Dimension 3).

Topic 1.3: Interim arrangements before revenue inflows **Negotiation issue:** whether to create interim structures/financing to build systems/staffing and do readiness work before 2029, or avoid "temporary measures."

- **Convergence:** none. UCL flags a near 50/50 split: about approximately half favour interim solutions and early establishment; the remainder consider it a lower-priority item.
- **Boundaries:**
  - **Robust interim package** (loans/bridging finance; steering committee; readiness support; early activation steps through Council and MEPC milestones).
  - **Light interim**: "if deemed appropriate," consider a steering committee for preparatory technical work.
  - **No interim arrangements** (reticence toward temporary measures, per UCL synthesis).

Current status: this is one of the clearest *fault lines*—and it directly affects the feasible pace for the "faster timeline" everyone says they want.

Topic 1.4: Governing Board composition and representation model **Negotiation issue:** Board **size**, **seat allocation**, and which equity principles are "hard-wired" (SIDS/LDC, developing/developed balance, gender, region, etc.).

- **Convergence:** strong support for principles of **regional balance** and attention to **SIDS/LDC representation**; UCL notes this was widely supported, with pre-submitted options already on the table.
- **Boundaries:**
  - **Explicit SIDS/LDC set-aside / tiered representation**: a **21-member Board** divided 7/7/7 across (i) SIDS+LDCs, (ii) other developing, (iii) developed; with rotation features and observer options.
  - **More open-ended composition**: EU/EC submission stresses representation, competencies, nomination criteria, and conflict-of-interest safeguards, but does not lock a single numeric model in the excerpted portions—i.e., it frames the *questions* more than a fixed seat split.

Current status: there's meaningful convergence on "SIDS/LDC should be guaranteed," but not yet on the final arithmetic (size/tiers/constituencies), nor on whether observers/industry seats exist formally.

Topic 1.5: Decision-making rules, accountability, and integrity safeguards **Negotiation issue:** how the Board takes decisions (consensus vs voting), term limits/rotation, transparency, audits, and conflict-of-interest controls.

- **Convergence:** recurring "core principles" call: **transparency, accountability, conflict avoidance**.
- **Boundaries:**
  - **Consensus-leaning IMO-style practice** as baseline, with further safeguards on COI and competency-based nominations; explicit attention to term limits/rotation and preventing conflicts.

- **Codified voting fallback**: draft governing instrument language (e.g., "make every effort…by consensus," but quorum + simple majority voting framework).

Current status: broad agreement on the values; divergence on how "legalistic" and prescriptive to be in the initial provisions.

Topic 1.6: Legal/administrative "home" of the Fund and its operational model **Negotiation issue:** how tightly the Fund is embedded in IMO financial/staff rules vs. endowed with more independent legal personality/operational autonomy.

- **Boundaries (as expressed in proposals):**
  - **Operate under IMO Financial Regulations/Rules and Staff Rules**, with explicit provisions describing relationship to MARPOL and a budget separate from IMO's general budget.
  - EU/EC framing explicitly seeks clarity on "how the Fund could be established under the remit of IMO," and leans toward reusing IMO practices where suitable.

Current status: early-stage; many delegations want an expert workshop / learning from existing funds (UCL), which suggests this is not close to closure.

Topic 1.7: Disbursement categories, eligibility, and "who benefits" (Regulation 41 implementation) **Negotiation issue:** how to operationalize Regulation 41 categories into *eligible uses*, *access modalities*, and *allocation shares*, including the relationship to just & equitable transition and DNI/food security.

- **Convergence:** strong signal that disbursement must support **JET**, especially for developing countries incl. SIDS/LDCs, and that Regulation 41 needs refinement into actionable guidance.
- **Boundaries:**
  - **Very targeted JET programming**: capacity development, modernization of small/medium fleets, port/coastal infrastructure, maritime administrations, seafarers, NAPs, etc.
  - **Formal "share-of-proceeds" redistribution concept**: India proposes a binding minimum share of net revenues redistributed to developing countries (with priority to SIDS/LDCs), leaving the numeric **[X%]** open.
  - **Open governance question:** whether/how **non-Parties and the shipping industry** should be involved is explicitly flagged by UCL as still unanswered.

Current status: high salience, low specificity. Convergence is high on principles; specificity on allocation rules, channels, and stakeholder roles remains low.

Topic 1.8: Financing instruments and access modalities **Negotiation issue:** grants vs concessional finance vs guarantees, "direct access" to States, readiness support, and voluntary contributions on top of RU inflows.

- **Boundaries:**

- **Direct access / simplified access** and a "country-driven approach," plus stakeholder involvement; also explicit allowance for voluntary contributions from States/IOs/private entities.
- More general "types of financing mechanisms" as a core element needing definition (UCL framing).

Current status: concept-level; details appear intentionally deferred to future drafting or Board work.

Topic Grouping 2: ZNZ definition and/or reward

Topic 2.1: What "ZNZ" is (definition vs eligibility for reward) Dimension

Whether ZNZ is **only** the MARPOL threshold (starting at **19 gCO2e/MJ**) or whether guidelines add **further constraints** (feedstock/pathway/technology restrictions).

Boundary of proposals

- **Narrow definition**: ZNZ = **threshold only**, no further limits (but could still differentiate reward). UCL reports this is the **majority** view in interventions.
- **"Definition stays broad, but eligibility/reward is differentiated"**: keep 19g definition, then **prioritize "long-term solutions"** via reward design (not via redefining ZNZ). UCL says this is the **direction of travel** if e-fuels are to be favored.
- **Tighter eligibility language**: at least some observer analysis argues for further eligibility definition (e.g., specifying eligible feedstocks/pathways linked to the 2024 LCA Guidelines; explicitly favoring e-fuels and excluding biofuels).

Convergence signal

High convergence on **"don't redefine ZNZ beyond 19g"**; remaining divergence shifts to **reward differentiation** rather than definition.

Topic 2.2: Which solutions get rewarded (fuel-only vs "fuels + energy sources + technologies") Dimension

Whether rewards are strictly for **ZNZ fuels**, or extend to **other energy sources** and **technologies** (notably onboard carbon capture).

Boundary of proposals

- **Fuel-centric in early years**: Some guideline drafts are explicitly framed as "ZNZ fuels used ... 2028–2032" with an option to amend later for "other energy sources and ZNZ technologies, including onboard carbon capture."
- **Explicit pushback on CCS as ZNZ**: UCL reports **more countries against** onboard CCS being considered ZNZ than in favor.

Convergence signal

Moderate convergence around **prioritizing fuels early**, with CCS remaining a **live fault line**.

Topic 2.3: Reward basis: "how much" is a ship rewarded for using ZNZ? Dimension

What the reward is *proportional to* (e.g., tonnes of CO_2e avoided on a WtW basis), and what **reference point** is used to compute "avoided" emissions.

Boundary of proposals

- **CO_2e avoided × reward rate** (simple formula): a draft guideline proposes avoided emissions computed as (reference GFI 93.3 – ZNZ GFI) × LCV × tonnes, then multiplied by a reward rate R (US\$/tCO_2e).
- **Reference GFI choice is itself a design lever**: Korea lays out options where the "reference" could be base target / direct target / the ZNZ threshold, changing the magnitude of "avoided" emissions and therefore the needed \$/tCO_2e reward price.

Convergence signal

Convergence on **WtW avoided-emissions logic**; less convergence on **the reference line** (because it interacts with RU pricing, total fund outflows, and perceived "fairness" of the incentive).

Topic 2.4: How the reward level is set (price discovery / governance / predictability) Dimension

Whether the reward is **fixed ex ante**, **adjusted ex post**, **auctioned**, or structured as **contracts for difference**—and who sets it (MEPC vs Fund) and how often.

Boundary of proposals (explicitly laid out)

One co-sponsored paper frames four "overarching concepts," all paying ships after verified use, but differing in *how/when* the reward level is set: **ex-ante fixed**, **flexible ex-post**, **reward auctioning**, **contract for difference**.

- In parallel, IMarEST argues auctions have advantages vs flat rate, including **limiting Fund liability** and **discovering prices**.

Convergence signal

Low-to-moderate. Many parties like the **simplicity** of a flat/standard method, but the submissions show active exploration of **auction/CfD** variants to manage Fund exposure and bankability.

Topic 2.5: Market responsiveness: fixed \$/tCO_2e vs marginal abatement cost (MAC) indexing Dimension

Whether the reward should be **administratively set** (fixed price for a period) or **market-indexed** using a MAC approach (tying reward to observed fuel prices and abatement).

Boundary of proposals

- **Fixed rate risk**: WSC highlights the risk of "under reward" early and "over reward" later if a price is fixed over a multi-year period.
- **MAC methodology**: WSC proposes Annual Reward = (tonnes CO_2e avoided vs a reference GFI) × MAC, with MAC itself derived from observed market prices and GFIs (including a quantity-weighted approach across ZNZ fuels).

Convergence signal

Moderate divergence: broad recognition that **reward needs to reflect real cost gaps**, but disagreement on whether the "right" implementation is **simple fixed rate** vs **indexing/MAC** (complexity vs accuracy/predictability trade-off).

Topic 2.6: Differentiation tools: multipliers / prioritization of "long-term solutions" Dimension

Whether rewards are **uniform** for all ZNZs or **differentiated** (by fuel type, pathway, TRL, supply-chain complexity, negative emissions, etc.).

Boundary of proposals

- UCL: support exists for prioritizing reward to "long-term solutions," with several countries saying those should be **e-fuels**; direction seems to be **differentiation of reward/eligibility**, not narrowing the ZNZ definition.
- Korea: explicitly proposes **multipliers** applied to avoided emissions to reflect policy priorities (DAC e-fuels, low TRL fuels, very low/negative GFI).
- WSC: suggests a **non-monetized reward component using multipliers** to ease liquidity concerns while maintaining incentives.

Convergence signal

Emerging convergence that **some differentiation mechanism may be needed** (especially to avoid the system simply selecting the cheapest 19g-compliant options), but there is no convergence yet on **how strong** differentiation should be, or **which** attributes it should reward.

Topic 2.7: Interaction with the rest of the NZF (RUs/SUs + consistency constraints) Dimension

How reward design must align with RU pricing and other compliance mechanics to avoid perverse incentives and inconsistent signals.

Boundary of proposals

Korea is explicit that the reward benchmark (reference GFI) and reward price must be designed "in harmony" with RU pricing to avoid excessive/insufficient rewards for the same fuel and conflicts between systems.

Convergence signal

High convergence on the *principle* of needing internal consistency; open divergence on the parameter choices (reference line, price, multiplier logic).

Topic 2.8: Equity / developing-country deployment constraints Dimension

Whether reward guidelines should explicitly incorporate constraints or priorities to ensure ZNZ deployment in developing countries and avoid unintended harms (land use, food security, fund depletion).

Boundary of proposals

Mexico's submission frames the reward design as central to: (i) ensuring sufficient Fund revenue given competing uses; (ii) incentivizing **scalable** long-term ZNZs; (iii) avoiding adverse impacts like emissions shifting, land-use change and food-security risks; and (iv) preserving funds for developing countries.

Convergence signal

Moderate: widespread rhetorical support for "just and equitable transition," but the **degree of hard-wiring equity into reward eligibility/formulas** remains open (i.e., whether this is handled via reward rules, fund disbursement windows, or outside the reward mechanism).

Topic 2.9: Administrative scope and complexity management Dimension

How complex the reward system can be (granularity, number of "reward windows," administrative burden), and whether it can be accessed by ships outside the mandatory size threshold.

Boundary of proposals

- Austria et al. explicitly warn that "carefully balancing" issues can create **very granular, complex reward windows**, and highlight unforeseen consequences in similar systems.
- They also float widening access to ships **400–5,000 GT** on a voluntary basis—while noting this increases demand on limited reward resources.

Convergence signal

Convergence on "avoid unnecessary complexity," but divergence on *where complexity is justified* (e.g., to enable price discovery, differentiation, equity, or bankability).

Topic Grouping 3: Fuel Certification

Topic 3.1: What exactly is being "recognized" by IMO (and at what level)? What's being discussed: A formal IMO process to recognize "Sustainable Fuel Certification Schemes (SFCS)" and define what "recognized" means (scope, conditions, publication). Brazil et al. frame recognition as an IMO process that (i) recognizes schemes, (ii) defines application/assessment procedures, and (iii) defines annual reporting back to IMO.

Boundaries (how wide the proposals vary):

- **Recognition object:** national/multinational frameworks and commercial/voluntary schemes are all contemplated as eligible "SFCS" candidates.
- **Scope of recognition:** explicit suggestion that the recognition decision should state the **scope** for which the SFCS is recognized (implying partial recognition by fuel types/pathways/activities).

Convergence signal: Moderate. UCL notes "some convergence" on "how to define and recognize certification bodies," including openness to a process run via the IMO Secretariat (vs heavy Member State involvement).

Topic 3.2: Institutional design: who does what (Committee vs Secretariat vs "assessment group")? What's being discussed: How centralized the recognition pipeline should be, and whether there is an intermediate assessment mechanism (e.g., assessment group) before the Committee decision.

Boundaries:

- **Member State vs Secretariat involvement:** UCL reports flexibility, with some seeing appointment/recognition as potentially run by the **IMO Secretariat** rather than Member States directly.
- **Appeals/objections architecture:** draft options include (a) scheme owner appeals directly (objection/appeal) vs (b) scheme owner appeals “through supporting Member State” to SG/IMO processes (two explicit options).

Convergence signal: Moderate on “there will be a defined process,” but still open on **how legal/administrative standing is structured** (direct vs via Member State).

Topic 3.3: Scope of certification along the fuel pathway (WtT only vs extending into TtW / ship side) **What’s being discussed:** Whether certification schemes should cover **only upstream Well-to-Tank (WtT)** attributes or also extend to **Tank-to-Wake (TtW)** elements and shipboard verification.

Boundaries:

- **WtT-only model (dominant in submissions):** Brazil et al. define SFCS as certifying the **WtT stage** of the fuel supply chain (explicitly).
- **Concern about “certified WtW FLL” / ship inclusion:** Norway et al. argue that interpreting regulation 34.2 as requiring certification schemes to extend into TtW would drag shipping companies into certification with high burden and unclear benefit, and would blur roles between schemes and Administrations/ROs.

Convergence signal: Growing. Multiple documents and the UCL synthesis point toward **keeping SFCS focused on WtT**, with Administrations/ROs handling ship-side verification, but the legal interpretation edge case remains a live boundary question.

Topic 3.4: Fuel Lifecycle Label (FLL) operationalization: what is reported, when, and in what system? **What’s being discussed:** The FLL as the carrier for certified information (fuel type, feedstock, production process, emissions factors, sustainability themes) and how it links to other IMO reporting infrastructure.

Boundaries:

- **Where the reporting lives:** UCL highlights an open question—integrate FLL reporting into **IMO DCS** vs the future **GFI Registry**.
- **How complete the FLL is at bunkering:** Norway et al. note “may accompany the BDN” is read as allowing cases where some certified info is not finalized at bunkering (notably WtT), and they caution against an operational design that requires shipping companies to “finalize” partially completed labels without robust anti-tamper IT controls.
- **Template standardization:** Norway et al. push for a standardized FLL template that schemes must use, potentially embedded in either SFCS recognition guidelines or FLL circulation guidelines.

Convergence signal: Relatively strong on *what kinds of data fields* should be captured (UCL notes "growing convergence" on reporting mechanisms and data categories), but not yet on **systems integration** (DCS vs Registry) or **handover mechanics** at bunkering.

Topic 3.5: Chain of custody (CoC) model(s): physical traceability vs accounting flexibility **What's being discussed:** How certified attributes are tracked along supply chains—especially whether and how to permit **mass balance** and/or **book-and-claim**.

Boundaries (this is the widest divergence area):

- **Physical restriction end:** Some delegations prefer traceability restricted to the **physical supply chain** (UCL).
- **Mass balance allowed:** Supported by multiple industry submissions; IPIECA/RINA/IBIA define mass balance as tracked quantities with a feasible physical connection, without full segregation, and argue it is essential for pathways like bio-LNG / grid-integrated gases where physical differentiation is impossible.
- **Book-and-claim allowed:** Republic of Korea calls for "policy acceptability" exploration—motivated by uneven fuel availability and practical shipping realities—and notes IMO LCA Guidelines currently appear to prefer physically-linked approaches and do not explicitly recognize book-and-claim.
  > >
  > > CLIA explicitly asks to enable mass balance now and initiate work to develop book-and-claim provisions, referencing use in other sectors and ISO 22095 alignment.

Convergence signal: Low. UCL explicitly flags **diverse opinions** here, with split between "physical-only" and "more flexible accounting approaches (mass balance, etc.)."

Topic 3.6: Alignment / fungibility with other regulatory certification schemes (CORSIA, EU RED, etc.) **What's being discussed:** Avoiding a "siloed" maritime-only scheme that creates duplicative or conflicting certification requirements and slows investment.

Boundaries:

- **Interim "use existing schemes" approach:** IPIECA proposes an interim solution where certifications under existing schemes can be accepted to demonstrate compliance until IMO-recognized SFCS are available, provided they align on criteria, boundaries, and documentation needed for FLL completion.
- **Multiple claims risk management:** IPIECA also suggests allowing multiple claims along the supply chain across overlapping schemes until the final point of regulatory compliance (claimed once), but acknowledges the need for "strong guardrails" against inappropriate dual claims.

Convergence signal: Emerging, but politically sensitive. The practicality arguments (timelines, investment barriers, cross-sector supply chains) are strong; the limiting factor is **integrity / double-claim safeguards** and how that interacts with CoC choices.

Topic 3.7: Assurance rigor and global capacity (accreditation, auditors, trained resources) **What's being discussed:** The requirements for certification bodies, accreditation bodies, auditors, and economic operators, and whether capacity constraints will delay access to "actual emission factors" vs defaults.

Boundaries:

- **High-rigor model:** Brazil et al. emphasize core principles like impartial governance, independent third-party audits, competence and training, and international standards for certification/accreditation bodies.
- **Capacity constraint warning:** IPIECA warns that tight timelines plus limited trained verifiers could delay certification (especially outside major hubs), affecting investment and early deployment.

Convergence signal: High on "need for credibility," but open on **how to scale capacity fast enough** without stalling early fuel deployment.

Topic 3.8: Timeline and phasing (what must be ready by when) **What's being discussed:** Sequencing guidelines adoption vs operational readiness of recognized schemes and certified fuel batches.

Boundaries:

- **Near-term deadlines as a design constraint:** IPIECA points to "obligation for certification by 1 March 2027" and views uncertainties as investment barriers.
- **Process schedule for availability of certified fuels:** Brazil et al. target the ability for schemes to apply right after MEPC 84 to enable certified fuels by early 2028, noting the many steps required (scheme recognition –> auditor training –> economic operator certification –> fuel production/supply).

Convergence signal: Strong on "time is tight" and "pragmatism matters"; weaker on *which interim flexibilities* (e.g., accepting other schemes; CoC flexibility) should be used to make the timeline feasible.

Topic Grouping 4: GFI and GFI compliance approaches

Topic 4.1: "Technology neutrality" vs adjustments (multipliers/correction factors) **What's being decided:** whether the GFI calculation should remain strictly LCA-based and "technology-neutral", or whether *adjustment constructs* (multipliers/correction factors) should be introduced for policy reasons (e.g., to further incentivize particular pathways).

Boundaries of the debate:

- **Center of gravity:** keep the calculation anchored in the formula in draft regulation 33 and the LCA Guidelines as the neutral assessment tool.
- **Caution/constraint position:** multipliers/correction factors should be treated cautiously to avoid preferential treatment and to protect **equity + environmental integrity**; some explicitly opposed multipliers on those grounds.
- **Related but distinct "multiplier" thread:** UCL records that an e-fuels/ZNZ-linked **SU multiplier** concept was discussed (more under "registry/guidelines architecture"), with explicit supporters and explicit opponents citing environmental integrity and just transition concerns.

Convergence level: moderate on "start from neutral LCA" + high sensitivity to anything that looks like preferential treatment; low convergence on whether a multiplier can ever be acceptable.

Topic 4.2: How to count onboard renewables and electricity (wind / shore power / solar) What's being decided: how "energy provided" by non-fuel sources is represented in the attained GFI calculation.

Boundary options observed:

- **Option 1: "Effective energy delivered to the ship"** (Norway approach in 20/2/6) — preferred by many as the technology-neutral baseline.
- **Option 2: "Fuel-equivalent" approach** (advocated by IWSA) — convert wind energy into an equivalent fuel quantity using engine efficiency as a conversion factor (IWSA argues Norway's method mixes input and output energy and can distort equivalency).
  > >
  > > MEPC 84/7/14 notes several delegations saw merit in further exploring IWSA's fuel-equivalent concept, and invited further intersessional work.

Convergence level: emerging convergence on using 20/2/6 as the basis, but clear *open technical controversy* on wind equivalency mechanics.

Topic 4.3: Treatment of special operational contexts (ice-classed ships) What's being decided: whether and how to address specificities of ice-class operations within GFI calculation guidelines.

Boundary: raised as important by some delegations; not yet resolved into a concrete method in the official summary.

Convergence level: low/early-stage (problem recognized; solution space not yet negotiated).

Topic 4.4: Base text for guidelines What's being decided: what draft is used as the working baseline.

Boundary / convergence: broad support to use the **annex to ISWG-GHG 20/2/6** as the basis for GFI calculation guidelines, with further updates to be developed for future sessions.

Topic 4.5: "Guidelines should not go beyond the regulation" What's being decided: how "interpretive" the guidelines can be relative to the agreed draft regulation 36.

Boundary:

- A clear cluster argued that regulation 36 is already clear and that guidelines should **not exceed** it.
  > >
  > > This implicitly constrains the room for adding new compliance instruments via guidelines (as opposed to clarifying procedures).

Convergence level: fairly high as a procedural constraint, but it pushes substantive debates into "is this merely clarification or a new mechanism?"

Topic 4.6: Voluntary cancellation of units outside the sector What's being decided: whether voluntary cancellation by non-sector actors (or in a way that reaches beyond the sector) should be recognized, and under what safeguards.

Boundary:

- **Supported in principle** by some.
- **Opposed/concerned** by others due to **double-counting risk** and the need for complex safeguards.
  > >
  > > UCL also frames "use or cancellation of SUs" as a debated compliance approach area (mixed views).

Convergence level: low—this is a classic integrity vs flexibility fault line.

Topic 4.7: Limiting the use of Surplus Units (SUs) / managing SU supply What's being decided: whether to restrict SU use (or otherwise manage SU market/supply) to avoid outcomes like a very low SU price and weak incentives.

Boundary:

- Several delegations expressed concerns about **limiting** SU use.
- UCL reports a related idea—using the registry/guidelines to apply **constraints on SU supply** to avoid low SU prices and improve predictability—received **little support**, suggesting it is unlikely to progress.

Convergence level: moderate convergence *against* active SU supply constraints (based on the "little support" signal), but not necessarily convergence on what to do instead.

Topic 4.8: Candidate "base documents" for compliance-approach guidelines What's being decided: which past drafts anchor the compliance-approach guideline drafting.

Boundary: MEPC 84/7/14 notes multiple possible base documents, including prior submissions (e.g., annex 2 to ISWG-GHG 17/2/8 (Angola et al.)).

This also intersects with the registry debate because some registry language proposals (e.g., "duly communicated and recorded") were flagged as problematic by EDF.

Convergence level: early-stage—no single agreed base text for compliance approaches emerges from the official summary (unlike the 20/2/6 signal for calculation).

Topic Grouping 5: IMO GFI Registry

Topic 5.1: Registry's role in the architecture: data platform vs "market" function Negotiation issue: what the Registry *is* and *is not*.

- **Narrow view (clear convergence):** Registry as a **data platform to record/track/attest** units (RUs/SUs), *not* a trading venue. This is explicitly recalled in the official report.
  > >
  > > **Convergence:** high.

- **Implication boundary:** even if "not a trading platform," it still must support **transfer of units and proof-of-title** (see dimension 2). The debate is *where to draw the line* between a recordkeeping system and a compliance operations system (dimension 6).

Topic 5.2: Legal status of Registry records and units: "constitutive record" vs unresolved **Negotiation issue:** whether the Registry is the **sole legal proof of title** for units (SUs/RUs), and how disputes are avoided.

- **EDF position (strongly specified):** Registry should be the **constitutive record** of unit ownership/transactions and the *sole legal proof of title* to reduce disputes and provide legal certainty.
- **Boundary / uncertainty:** Canada explicitly **reserved its position** on the legal status of SUs recorded in the Registry (signal that this is not fully settled).

Convergence: medium. There is momentum toward EDF's clarity, but at least one delegation flagged legal uncertainty.

Topic 5.3: Governance & account administration model: centralized vs two-tier / hybrid **Negotiation issue:** who does what—IMO Secretariat vs national administrations.

- **Two-tier structure (dominant direction):** multiple places show support for **national Administrations** handling account opening/closing + participant verification, with the **Secretariat** managing the technical platform (and, per EDF, issuances).
- **EDF adds institutionalization:** propose a **Registry Management Committee** (IMO + national administrators) for oversight of IT updates/operational rules.

Convergence: high on *hybrid/two-tier* as the practical scaling model; details (committee mandate, who "issues," etc.) remain open.

Topic 5.4: Data scope and granularity: Appendix XII minimum vs "fuel-level / automatable" dataset **Negotiation issue:** what data sits in the Registry and at what granularity.

- **Minimum dataset:** some delegations want the data to be determined based on **draft appendix XII** (i.e., avoid over-collection).
- **More granular dataset:** others want greater granularity to enable **automation of GFI calculation and checks** (EDF explicitly calls for fuel-level info and automatic calculation/consistency checks).
- **Operational/IT overlay:** RINA stresses **automation over manual uploads**, built-in validation/error handling, and "single source of truth" if connected to IMO DCS (to manage overlap).

Convergence: medium. There is shared interest in automation, but the "how much data" boundary is still wide.

Topic 5.5: Transparency, access, confidentiality: public information tool vs confidentiality constraints **Negotiation issue:** what is public, what is restricted, and who gets access (including PSC).

- **EDF: strong transparency posture**
  > >
  - Build a **public website** with aggregated non-confidential data; also suggests disclosing compliance status, targets, prior verified emissions, etc.
  - Frames lack of transparency as a **just & equitable transition** issue (burden on smaller/less resourced actors).
- **Official report flags this as a key design topic:** "data access and confidentiality" and "access for Port State Control" are explicitly listed among aspects delegations want covered.

Convergence: medium. Agreement that access/confidentiality must be designed; the boundary is **how "open"** the Registry should be beyond minimum compliance reporting.

Topic 5.6: Integrity controls & operational safeguards: banking-grade security vs lighter touch; verification boundary Negotiation issue: cybersecurity, fraud prevention, and what verification the Registry itself performs.

- **EDF: robust IT/security package** (online-banking equivalent): 2FA, KYC checks, backups, and even a **24-hour transaction delay** to reduce risk.
- **Official report: fraud prevention, data standards/security, interoperability, automated error checks** are repeatedly highlighted as desired guideline content.
- **Verification boundary (important convergence signal):** IACS cautions against discrepancies between Registry balance statements and compliance certificates if records are updated; and emphasizes that under draft regulation 37, **Administrations report verified data**—so the Registry **should not undertake additional verification actions**.

Convergence: high on "needs strong safeguards," but with a clear boundary that **verification authority remains with Administrations** (Registry supports; does not become a verifier).

Topic 5.7: Built-in "compliance functions" inside the Registry: passive record vs active enforcement support Negotiation issue: whether the Registry should merely record units/data, or actively support compliance workflows.

- **EDF: automate compliance-support features** (while keeping oversight with Administrations), e.g.:
  - automatic account restrictions for missing data,
  - public flagging of non-compliance,
  - RU-based adjustments for late submissions / handling late submissions via multipliers or RU requirements.
- **Boundary:** this can feel like "enforcement logic" embedded in IT, versus a neutral ledger.

Convergence: medium. The concept is live and detailed in EDF; the official report lists related themes (account management, fraud prevention, PSC access), but does not show settled consensus on how "active" the platform becomes.

Topic 5.8: Implementation approach & timeline: urgency vs "wait until adopted"; procurement/requirements discipline **Negotiation issue:** how hard to push now, given NZF not yet adopted at MEPC/ES.2, and how to de-risk delivery by 2027/2028 deadlines.

- **Hard deadlines are clear in the text:** ships must have accounts by **1 Oct 2027** and pay first annual admin fee by **30 Jun 2028**.
- **Broad support for Secretariat preparatory work** (mapping requirements/practices), and Singapore pilot welcomed—*but* some delegations argue there is **no urgency** (since amendments not adopted) and caution against allocating substantial resources/budget now.
- **RINA "delivery discipline" package:** need an "**intelligent customer**" at Secretariat, strong requirements/specification before RfP, heavy documentation, stakeholder testing/piloting, and a structured feedback mechanism throughout development and operations.

Convergence: medium-high that **preparatory work and piloting are useful**, but the *pace* boundary remains contested (push now vs wait for adoption certainty).

Topic 5.9: Annual Registry administration fee: cap vs no cap; fixed vs tiered; discounts/installments; timing gap **Negotiation issue:** fee purpose, structure, equity, and timing.

- **Purpose converges strongly:** fee should cover **administrative costs only** and "serve no other purposes."
- **ICS proposal:** fee should be **minimal/nominal or zero**, and **cap** the maximum at ≤ **US$500 per ship**.
- **Republic of Korea emphasis:** highlights operational reality (multi-million US$ annual costs anticipated) and a **timing gap**: ships open accounts by Oct 2027, but first fee not due until Jun 2028—implying bridging/financing and "system finalization" timing issues.
- **Debate boundary in the official report:**
  > >
  - support for defining fee purpose/scope and for early-payment discounts / instalments (linked to RoK paper),
  - support for fixed fee with a cap (linked to ICS),
  - opposition/hesitation: some say it is **premature** to set methodology/caps now, and others propose **tiered fees by tonnage** to accommodate smaller ships / developing states / coastal operators who occasionally go international.

Convergence: medium on *purpose*; low-to-medium on *structure* (cap vs tiering vs "too early").

Topic 5.10: Multiplier applied to SUs (extra incentive for ZNZ / e-fuels) This is a "policy lever" inside Registry guidelines, where the Registry discussion touches broader NZF political economy.

- **UCL readout:** multiplier concept appears in different places; here it is specifically **multiplying the number of surplus units** for ships using ZNZ (often framed as e-fuels), providing stronger incentive than SUs alone.

- **Boundary:** clear split:
  - supporters want additional incentive,
  - opponents cite **environmental integrity** (more SUs can mean higher allowed emissions) and **just & equitable transition** concerns.

Convergence: low (material disagreement).

Topic 5.11: Using the Registry to constrain SU supply (price-stabilization / predictability) This is a "policy lever" inside Registry guidelines, where the Registry discussion touches broader NZF political economy.

- **UCL readout:** proposal to use Registry rules to apply constraints on SU supply (to avoid low SU prices and keep incentives predictable) received **little support**, so it is unlikely to progress.

Convergence: high *against* pursuing SU supply constraints (i.e., the idea seems to be dying).

Topic Grouping 6: IMO Life Cycle GHG Assessment (LCA) framework

Topic 6.1: What "sustainability themes/aspects" mean in practice (and how to implement them consistently) What's being discussed

- Moving from high-level sustainability themes in the 2024 LCA Guidelines to **operational, auditable expectations** for certification schemes and certifiers (illustrated via **water**). The Australia et al. submission points to drawing from existing systems (e.g., ICAO/CORSIA-style non-mandatory guidance) and even references methods like **ISO 14046 / AWARE** as supporting approaches for water-related performance evidence.

Boundary (range of proposals)

- **Light-touch, non-mandatory guidance** aligned with existing certification practice (minimize incremental admin burden) <–> **more prescriptive IMO-level guidance** that reduces interpretive variability across schemes/jurisdictions.
- **Process structure:** a stronger push to **merge/align** the "procedural" work on recognizing schemes with the "substantive" guidance on themes/aspects, to prevent inconsistencies.

Convergence signal

- In the official report, the water submission is framed as a call for **globally consistent application**, leveraging common elements with existing systems, and **establishing a correspondence group** with a target of completing work **by end-2026**.

Topic 6.2: Role and recognition of certification schemes / certification bodies (governance of the "trust layer") What's being discussed

- How IMO recognizes certification bodies/schemes and what role Member States vs IMO plays.

Boundary

- **State-centric appointment/oversight** of certification bodies <–> **IMO Secretariat-led process** to appoint/recognize certification bodies (more centralized / harmonized).

Convergence signal

- UCL notes "**some convergence**" on defining/recognizing certification bodies and that countries "generally showed flexibility" on whether appointment becomes a Secretariat process.

Topic 6.3: Chain of custody and traceability philosophy (physical vs accounting)

(This sits at the LCA framework–certification interface, but it materially affects whether/when "actual values" and sustainability attributes can be claimed.)

Boundary

- **Physical-only traceability** (identity preservation / segregation logic) <–> **more flexible accounting** such as **mass balance** and beyond.

Where debate stood

- UCL characterizes this as **more diverse opinions** than for certifier recognition: some delegations want strict physical traceability; others support mass-balance-type approaches (explicitly noted).

Topic 6.4: Default WtT emission factors: how to build them, how conservative they should be, and how often they change What's being discussed

- The GESAMP-LCA WG pipeline for **submitting/reviewing default factors**, and the statistical philosophy behind "default."

Boundary

- **Upper-bound conservatism** (current "select the upper emission value as default") <–> a **scientifically representative** method (e.g., central tendency/robust representativeness), with explicit **data quality standards** and periodic updates.

Convergence signal

- The official report highlights a concentrated package of methodological proposals: align data with international quality standards; replace "upper value default" with a representative method; implement a **biennial review**; allow temporary credible alternative datasets where fossil pathways lack robust data; clarify crediting incl. negative emissions; and create variants for "cargo used as fuel" cases.
- UCL frames the LCA Guidelines and their default values as central to both **environmental integrity** and whether the policy effectively promotes transition, with ongoing work via **GESAMP** to refine values and methodology.

Net assessment

- This looks like a **high-salience technical convergence area** (many specifics on the table), but still with a real boundary between "conservative by construction" vs "representative + governance/updates."

Topic 6.5: "Actual values" methodology: ISO LCA standards and how to treat contracted electricity (PPAs) What's being discussed

- Whether (and how explicitly) to incorporate **ISO LCA standards** into calculating **actual WtT** emission factors, and how broadly **power purchase agreements (PPAs)** can be used across the fuel life cycle.

Boundary

- **Narrow, cautious use** of ISO/PPAs (avoid loopholes / double counting / weak additionality) <–> **explicitly broadened use** (increase feasibility of demonstrating lower actual values where procurement is contractually evidenced).

Where captured

- The official report explicitly notes proposed amendments to **clarify inclusion of ISO LCA standards** for actual WtT values and to **expand PPAs** for the marine fuel life cycle.

Topic 6.6: How CCS is treated in the LCA framework (pathways + permanence/safety guidance) This appears in *two distinct layers*:

1. **Pathway accounting layer** (e.g., LNG with upstream CCS)

- Norway proposes a **new fuel pathway code** in Appendix 1 to enable verification/certification of **actual WtT** emission factors for LNG with upstream CCS.

2. **Environmental integrity layer** (geological storage quality, monitoring, long-term containment)

- IPIECA & RINA propose incorporating **underground CCS provisions** into the LCA Guidelines to ensure **safe and secure long-term containment** via geological storage, and suggest further IMO guidance as needed.

Boundary

- **Treat CCS primarily as a quant adjustment** in pathway factors <–> **treat CCS as a regulated chain with explicit permanence/leakage/monitoring expectations** embedded in the LCA framework (potentially requiring additional IMO guidance).

Topic 6.7: Carbon source/crediting logic for e-fuels (example: e-methanol and "inherent CO_2") What's being discussed

- Malaysia proposes recognizing **pre-combustion captured inherent CO_2 from natural gas processing** as "carbon neutral feedstock" for low-carbon e-fuels (specifically e-methanol), with crediting linked to using captured CO_2.

Boundary

- **Strict carbon-source hierarchy** (e.g., favor DAC/biogenic, treat fossil-derived CO_2 cautiously) <–> **broader eligibility** where captured "inherent CO_2" can be treated as neutral/creditable in e-fuel production accounting.

Net assessment

- A classic **systems-boundary debate**: whether the framework rewards "captured CO_2 utilization" similarly across fossil/biogenic/DAC sources, and how this interacts with double counting / counterfactual baselines.

Topic 6.8: Avoided emissions / "efecu" recognition (do we credit displacement?) What's being discussed

- Proposals to recognize **avoided emissions** in developing methodological guidance for the WtT emission factor parameter **efecu** (as flagged in the agenda list and submissions set).

Boundary

- **No avoided-emissions crediting** (keep WtT accounting attributional and conservative) <–> **allow avoided/displacement crediting** under specified conditions (requires robust counterfactuals and governance to prevent inflation of claims).

Where the debate stood

- The fact it's framed as a proposal for methodological guidance (rather than already-accepted practice) suggests this is still **open and potentially contentious**, because it changes what the LCA framework is *for* (pure accounting vs also incentivizing system substitution).

Topic 6.9: ILUC risk: whether it's addressed qualitatively or quantitatively and how it gates fuel eligibility What's being discussed

- Whether IMO should incorporate an ILUC-risk approach and how (especially for biofuels). RINA's INF.3 argues IMO's GFS does not yet address ILUC risk adequately and recommends recognizing credible **low-ILUC certification schemes** and restricting ILUC-risk feedstocks unless certified low-ILUC risk.

Boundary

- **Primarily scheme-based safeguards** (recognize credible low-ILUC certifications; restrict high-risk feedstocks) <–> **explicit quantitative risk-based ILUC accounting** (implied by the presence of a "quantitative risk-based approach" submission in the agenda list).

Net assessment

- Debate boundary is wide here: it's not just "how strict," but **what form** the policy takes (certification gating vs quantified ILUC adjustments).

Annex I: Topic Classification and Cluster Mapping

Topic-to-Bucket Classification

Each NZF design topic identified in Annex H was evaluated against two criteria: macroeconomic materiality (Criterion A) and reduced-form parameterisability (Criterion B). The classification methodology is described in §2.3.1.

Bucket X — Material and Reducible (8 topics)

1.7, 1.8, 2.3, 2.4, 2.5, 2.7, 4.7, 5.10

Bucket Y — Material but Not Reducible (18 topics)

1.1, 1.2, 1.3, 1.4, 2.8, 2.9, 3.5, 3.7, 3.8, 4.2, 4.3, 4.6, 5.11, 6.3, 6.4, 6.5, 6.6, 6.7

Note on topics 1.1, 1.2, and 1.4: Individually, these Fund governance and operationalisation topics (scope of governing provisions, timeline to operationalise, and Board composition) relate to institutional process and would satisfy Criterion A only weakly. However, jointly with topic 1.3 (interim arrangements), they determine whether the Net-Zero Fund is operational in time to disburse revenue during the early implementation period. This joint effect creates a material implementation risk for the fiscal offset channel — the Fund delivery risk documented in §2.3.2 and implemented as a scenario toggle in §2.4.6. The reclassification from Z to Y reflects the analytical judgment that Fund operationalisation timing is macroeconomically material when considered as a cluster, even though the individual governance design choices (Board size, decision rules, legal home) are not.

Bucket Z — Not Material to Macroeconomic Outcomes (22 topics)

1.5, 1.6, 3.1, 3.2, 3.3, 3.4, 4.4, 4.5, 4.8, 5.1, 5.2, 5.3, 5.4, 5.5, 5.6, 5.7, 5.8, 5.9, 6.1, 6.2, 6.8, 6.9

Note: Several topics in the Annex H inventory (2.1, 2.2, 2.6, 3.6, 4.1) are not independently classified because they are either definitional (e.g., 2.1: the ZNZ threshold is largely settled at ≤19 gCO_2e/MJ), procedural (e.g., 4.1: technology neutrality as a design principle), or cross-cutting (e.g., 3.6: alignment with other schemes). Their economic effects, to the extent they exist, operate through topics already classified above.

Bucket X Topic-to-Cluster Mapping

The eight X-topics are assigned to two design clusters based on their underlying economic transmission mechanism. The clustering methodology is described in §2.3.2.

Cluster 1: Effective Compliance Cost Path

Determines the magnitude and trajectory of the freight cost shock transmitted to the global trading system. Represented through the sensitivity ratio SR(t).

Topic	Description	Role in Cluster
2.3	Reward basis and reference line	Determines emissions avoidance credited per unit of ZNZ fuel
2.4	Reward setting mechanism and price governance	Determines whether reward tracks market conditions or is fixed
2.5	Reward responsiveness and market indexing	Determines dynamic adjustment of reward to cost-gap changes
2.7	RU/SU interaction and internal consistency	Determines consistency of incentives across compliance and reward
4.7	SU supply constraints and market management	Determines whether flexibility mechanism dampens or amplifies cost
5.10	SU multipliers for ZNZ fuels	Determines whether SU supply expansion reduces effective cost

Cluster 2: Net Fiscal Offset to African States

Determines the extent to which NZF revenue is returned to African economies in forms that mitigate the domestic cost burden. Represented through the fiscal offset θ.

Topic	Description	Role in Cluster
1.7	Disbursement categories, eligibility, and allocation logic	Determines share of Fund revenue reaching African states
1.8	Financing instruments and access modalities	Determines economic effectiveness of disbursement ($\omega = \delta \times \lambda$)

Bucket Y Topic-to-Risk-Cluster Mapping

The 18 Y-topics are grouped by the model parameter they most directly affect. Two risk clusters are implemented quantitatively; the remainder are treated as qualitative analytical boundaries.

Fund Delivery Risk (topics 1.1, 1.2, 1.3, 1.4)

These topics jointly determine whether the fiscal offset channel functions as designed. Implemented as a scenario toggle: ω_effective = 0 at 2030 when activated.

Certification and LCA Bottleneck Risk (topics 3.5, 3.7, 3.8, 6.3, 6.4)

These topics jointly determine whether fuels qualifying under the policy design can be certified and delivered in practice. Captured through the effective Tier 2 compliance price parameter T_2, which determines α.

Qualitative Analytical Boundaries (remaining Y-topics)

Topics 2.8, 2.9, 4.2, 4.3, 4.6, 5.11, 6.5, 6.6, 6.7: these influence cost incidence, strategic behaviour, or environmental integrity in ways that are directionally identifiable but not quantified in the reduced-form framework.

Annex J: References

IMO Instruments and Submissions

Circular Letter No. 5005 — Draft Revised MARPOL Annex VI (Secretariat).

ISWG-GHG 20 — Report of the twentieth meeting of the Intersessional Working Group on Reduction of GHG Emissions from Ships (Secretariat).

ISWG-GHG 20/2/14 — The ZNZs definition and reward mechanism (Fiji, Kiribati, Marshall Islands et al.).

ISWG-GHG 20/2/15 — Draft Guidelines on ZNZ rewards and methodology to determine such rewards (ICS and IBIA).

IMarEST ISWG-GHG 20/2/27 — Provisions, guidance, and other documents for supporting uniform and effective implementation.

CIA and Impact Assessment

MEPC 82/INF.8/Add.1 — Report of the Comprehensive Impact Assessment of the basket of candidate GHG reduction mid-term measures, Task 2 (DNV). Secretariat, 2024.

MEPC 82/INF.8/Add.2 — Report of the Comprehensive Impact Assessment of the basket of candidate GHG reduction mid-term measures, Task 3 (UNCTAD). Secretariat, 2024.

Compliance Cost and Maritime Decarbonisation

UCL/UMAS (2025). IMO Net Zero Framework: Assessing the potential options and costs of compliance. UCL Energy Institute.

CSC — Clean Shipping Coalition fleet projection data (referenced in CIA Task 2).

African Economic Impact

AFPH (2024). Assessing the economic impacts and trade-offs of a shipping carbon tax for African States. Africa Focus Policy Hub.

AFPH (2024). Inequity of EU Climate Regulations for Africa. Africa Focus Policy Hub.

Fiscal Multipliers and Climate Finance

Batini, N., Eyraud, L., Forni, L., and Weber, A. (2014). "Fiscal Multipliers: Size, Determinants, and Use in Macroeconomic Projections." IMF Technical Notes and Manuals 14/04.

Ilzetzki, E., Mendoza, E.G., and Végh, C.A. (2013). "How Big (Small?) Are Fiscal Multipliers?" Journal of Monetary Economics, 60(2), 239–254.

Kimaro, E.L., Keong, C.C., and Sea, L.L. (2017). "Government Expenditure, Efficiency, and Economic Growth: A Panel Analysis of Sub-Saharan African Low Income Countries." African Journal of Economic Review, 5(2), 34–54.

Climate Analytics (2023). "Five years of the Green Climate Fund."

Green Climate Fund (2023, 2024). Annual Reports. GCF-1 Progress Report (2023).

UNEP (2023). Bilateral climate finance data — instrument-level disbursement and effectiveness analysis.

Renewable Energy and Hydrogen

IRENA (2015). Renewable Energy Auctions: A Guide to Design. International Renewable Energy Agency.

IRENA (2019). Renewable Energy Auctions: Status and Trends Beyond Price. International Renewable Energy Agency.

IRENA (2020). Green Hydrogen Cost Reduction: Scaling up Electrolysers to Meet the 1.5°C Climate Goal. International Renewable Energy Agency.

IRENA (2022). Global Hydrogen Trade to Meet the 1.5°C Climate Goal: Part III — Green Hydrogen Supply Cost and Potential. International Renewable Energy Agency.

IRENA (2024). Renewable Power Generation Costs in 2023. International Renewable Energy Agency.

Trade and Economic Data

United Nations Comtrade Database. HS 2-digit bilateral trade data. Accessed January 2026.

World Bank. World Development Indicators (WDI). Series TM.VAL.FOOD.ZS.UN (Food imports as % of merchandise imports). Accessed January 2026.

World Trade Organization (2024). World Trade Statistical Review 2024.

National Statistical Sources (CPI Weights)

Ghana Statistical Service (2023). Consumer Price Index — 2023 rebasing.

Kenya National Bureau of Statistics (2019). Consumer Price Index — February 2019 base.

Statistics South Africa (2022). Consumer Price Index — 2022 basket weights.