Economic Impacts of the IMO Net-Zero Framework on African Economies — Methodology Annex

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

Contents

1. Introduction

1.1 Purpose

This paper presents the methodology and findings of an economic impact assessment designed to support African delegations in understanding and responding to the economic implications of design choices within the IMO Net-Zero Framework (NZF). The analysis estimates the potential GDP, consumer price, and trade cost impacts of specific NZF design parameters on African economies, and identifies which design choices offer the greatest potential to mitigate those impacts.

The broad architecture of the NZF has been agreed. The regulatory text approved at MEPC 83 (April 2025) establishes a greenhouse gas fuel intensity (GFI) standard for international shipping, a two-tier compliance mechanism through which ships exceeding emission limits must purchase remedial units (RUs), and an IMO Net-Zero Fund that collects the resulting revenue and deploys it through two channels: rewards for ship operators using zero and near-zero emission (ZNZ) fuels, and disbursements to developing countries. Formal adoption of the MARPOL Annex VI amendments was deferred at the extraordinary session of MEPC (MEPC 83 ES, October 2025), with adoption expected at MEPC 84 (April 2026) or a subsequent session. The deferral does not alter the regulatory text; the design details within the agreed architecture remain the focus of ongoing negotiations.

It is these design details — the share of Fund revenue directed to rewards versus disbursements, the eligibility rules for disbursement recipients, the instruments through which disbursements are delivered, the calibration of the reward mechanism — that determine the economic impact on African economies. Previous work by the African Future Policies Hub, including *Navigating Climate Action: Assessing the economic impacts and trade-offs of a shipping carbon tax for African states* (APRI/AFPH, 2024), established that maritime decarbonisation measures would have negative but bounded effects on African GDP, trade, and household incomes, with impacts falling disproportionately on the continent. The present analysis extends that work by moving beyond the question of *whether* the NZF will impose costs, to the question of *which design choices* determine how large those costs are and how effectively they are offset.

1.2 The Comprehensive Impact Assessment

The IMO's Comprehensive Impact Assessment (CIA), conducted in 2024 and considered at MEPC 82, provides the most authoritative economy-level impact estimates available for the NZF. Using a computable general equilibrium (CGE) framework based on the GTAP model, the CIA's Task 3 assessed ten policy scenarios across more than 100 economies. These results establish the aggregate magnitude and direction of impacts under specific policy configurations and serve as the baseline reference for this analysis.

The CIA was designed to evaluate pre-defined policy packages — fixed combinations of GFI stringency, levy level, and compliance mechanism — not to decompose the contribution of individual design choices within those packages. The results cannot directly answer the questions that African delegations now face: What is the economic consequence of a higher versus lower share of Fund revenue directed to rewards? What is the marginal value of securing broad versus narrow eligibility for disbursement? How does the choice of disbursement instrument affect the offset to GDP?

These are the questions that arise when the broad policy architecture has been agreed but detailed design parameters remain open.

This paper addresses that gap. It takes the CIA's Scenario 24 results as a structured reference point and decomposes the underlying economic shocks into components corresponding to identifiable NZF design choices. The approach preserves the country-specific impact structure embedded in the CIA's general equilibrium results while allowing transparent, parameter-by-parameter variation.

1.3 Structure

Section 2 presents the analytical approach: scope and country coverage, methodological rationale, the design-cluster framework, and the complete quantification chain including sensitivity analysis and synthesised findings. Section 3 presents the economic implications in a consolidated, policy-readable format. Section 4 summarises robustness, scope, and limitations. Section 5 presents conclusions.

The Technical Annexes (A–J), published as a companion document, provide derivations, data sources, validation evidence, the complete NZF design topic inventory, and the calculator specification.

An interactive online calculator implementing the identical computation chain is available at tools.fourthtack.com/nzf.

2. Analytical Approach

Key Results

> **Approach.** The analysis takes CIA Task 3 Scenario 24 as a structured reference point and decomposes its aggregate cost shock into components corresponding to specific NZF design choices. Each design choice maps to a standardised parameter; the combined effect is expressed as a sensitivity ratio (SR) that scales the S24 baseline impacts for each country. The model is validated against nine CIA scenarios within its operating range (SR $\in [0.85, 1.45]$), with mean prediction error of approximately -5% (conservative: overstates impact magnitude).
>
> **Principal findings.** (1) Aggregate impact is real but bounded: 2030 MLC increases of +7.7% to +8.6% across the design space, producing GDP impacts for the most exposed AF5 economies of -0.10% to -0.19%. (2) The Fund revenue split (σ) is the dominant policy lever, driving the largest variation in both compliance cost and fiscal offset. (3) By 2050, all scenarios converge; design choices are largely irrelevant. (4) Fuel price uncertainty generates variation comparable to the full range of policy effects. (5) Implementation risk (a non-operational Fund at 2030) can negate the benefits of even favourable design architecture.
>
> **Confidence and limitations.** The 2030 results carry the highest confidence. Reporting precision of ± 0.02 percentage points of GDP; results should be interpreted as first-order impact estimates, not point predictions. The model does not capture general equilibrium feedback, exchange rate adjustment, or strategic carrier behaviour. Omitted channels and their directional biases are documented in Annex B.

2.1 Scope and Country Coverage

The analysis covers the period 2028–2050, corresponding to the NZF's implementation trajectory. Results are presented at three time points — 2030, 2040, and 2050 — consistent with the CIA Task 3 snapshot years. Analytical confidence is highest at 2030, where NZF design parameters are most constrained by the adopted regulatory text, and decreases at 2040 and 2050, where all policy scenarios converge as the fleet transitions to lower-emission fuels.

The primary economic indicators are GDP, consumer prices (CPI), and trade flows (imports and exports), consistent with the CIA. Food price exposure is assessed as a component of CPI impact, given the significance of food imports in several African economies and the explicit reference to food security in the draft revised MARPOL Annex VI (regulation 41).

The model supports analysis of all economies and economy groupings with individual breakouts in the CIA Task 3 results — over 100 individual GTAP economies plus composite groupings that together provide global coverage. The framework can be applied to any country or grouping for which the CIA provides baseline results, without structural modification.

The initial analysis focuses on five African economies selected for their diversity of economic structure, trade exposure, and vulnerability profile: Ghana, Kenya, South Africa, Egypt, and Ethiopia (the "five African reference economies," AF5). Ghana and Kenya represent mid-sized, trade-exposed economies with significant maritime import dependence and high food-import shares. South Africa provides a comparator as the continent's most industrialised economy. Egypt combines high import dependence with a unique position as a maritime transit economy. Ethiopia is represented through the "Rest of landlocked economies in Africa" GTAP aggregate, which includes thirteen countries with Ethiopia as the dominant component (~46% of GDP). The aggregate's trade structure reflects blended East African, Southern African, and West African shipping corridor exposure. Results labelled "Ethiopia" should be interpreted as representative of landlocked African economies generally.

The analysis additionally reports results for fourteen African economies (the "fourteen African reference economies," AF14), adding Algeria, Cameroon, Côte d'Ivoire, Gabon, Madagascar, Mauritius, Morocco, Nigeria, and the United Republic of Tanzania. Both AF5 and AF14 are reported as GDP-weighted averages of independently computed country results. For comparative purposes, results are also presented for Brazil, China, India, Viet Nam, Bangladesh, Trinidad and Tobago, and the SIDS aggregate, as well as the CIA's standard economy groupings (World, Developed economies, Developing economies, SIDS + LDCs).

2.2 Methodological Framing and Complementarity

The CIA was designed to evaluate pre-defined policy packages, not to decompose the contribution of individual design choices within those packages. Each scenario represents a fixed combination of GFI stringency, levy level, and compliance mechanism. The results cannot directly answer questions such as: "If the reward setting mechanism produces outcome X rather than outcome Y, how does this change the GDP impact on Kenya?" These are the questions that African delegations face in the current negotiation phase, where the broad policy architecture has been agreed but detailed design parameters remain open.

The present analysis addresses this gap with a reduced-form, design-focused approach. Rather than re-running a global CGE model, it takes the CIA's Scenario 24 results as a structured reference point and decomposes the underlying economic shocks into components corresponding to identifiable NZF design choices. Scenario 24 models a WtW GFI with flexibility mechanism and represents the closest analogue to the NZF's regulatory architecture. It serves as the anchor from which design-driven variations are measured.

Each design choice is translated into a standardised economic parameter that can be varied independently. The effects are expressed as a sensitivity ratio (SR) scaling the CIA S24 reference: SR below 1.0 indicates lower costs than S24; above 1.0 indicates higher costs. This preserves the country-specific impact structure embedded in the CIA's general equilibrium results while allowing transparent decomposition of design-driven variation.

The approach offers three advantages for the negotiation context: transparency (every result traces to a specific design assumption); real-time comparison (alternative configurations can be assessed without CGE simulation); and internal consistency (all calculations are anchored to the CIA reference case).

The reduced-form approach does not capture general equilibrium feedback effects, exchange rate adjustment, terms-of-trade effects on export competitiveness, labour market concentration effects, or long-run structural adjustment. It treats cost pass-through as embedded in the CIA baseline rather than independently estimated. These omissions, their direction of likely bias, and several interpretive considerations are documented in Annex B. The results are best interpreted as first-order impact estimates that capture the direct economic effects of NZF design choices but not the full dynamic response.

2.3 Design-Cluster-Based Assessment Framework

The NZF encompasses design topics spanning incentive calibration, lifecycle accounting, fuel certification, registry architecture, fund governance, compliance mechanics, and revenue allocation. Not all materially influence near-term macroeconomic outcomes in African economies. A structured approach is required to identify the subset warranting quantitative treatment, organise those topics to prevent double-counting, and maintain a transparent audit trail from the full design space to the reduced analytical scope.

2.3.1 Identification and Classification of Economically Material Design Topics Drawing on the draft revised MARPOL Annex VI, the report of ISWG-GHG 20, and pre-session submissions to MEPC 84, 48 design topics were identified across six groupings: the IMO Net-Zero Fund (1.1–1.8), ZNZ definition and reward (2.1–2.9), fuel certification (3.1–3.8), GFI and compliance approaches (4.1–4.8), the IMO GFI Registry (5.1–5.11), and the LCA framework (6.1–6.9).

Each topic was evaluated against two criteria:

Criterion A: Macroeconomic Materiality. Does the topic plausibly affect, within 2028–2050, the magnitude of global compliance costs, the timing of cost realisation, the incidence of costs between carriers, exporters, consumers, and governments, or the magnitude or usability of fiscal transfers to States?

Criterion B: Reduced-Form Parameterisability. Can the topic's effect be represented as a standardised economic shock parameter — a cost level, cost timing adjustment, or transfer magnitude — without requiring a change in model class?

The intersection generates three buckets:

- **Bucket X — Material and Reducible.** Topics satisfying both criteria. These form the quantitative core: 8 topics.
- **Bucket Y — Material but Not Reducible.** Topics satisfying Criterion A but not B. Incorporated through a structured overlay framework: 18 topics.
- **Bucket Z — Not Material to Macroeconomic Outcomes.** Excluded from economic modelling: 22 topics.

2.3.2 Clustering of Design Topics by Economic Transmission Mechanism Among the eight X-topics, several act on the same underlying economic variable. Topics are grouped if they (a) affect the same standardised economic shock parameter, (b) cannot be varied independently without implicit assumptions about each other, and (c) correspond to a coherent negotiation-relevant lever.

Design Cluster 1: Effective Compliance Cost Path. This cluster determines the magnitude and trajectory of the freight cost shock transmitted to the global trading system. It is the primary driver of CPI and trade impacts, and the single most consequential set of design choices for African states. Included topics: reward basis and reference line (2.3), reward setting mechanism (2.4), reward responsiveness (2.5), RU/SU interaction (2.7), SU supply constraints (4.7), and SU multipliers for ZNZ fuels (5.10). The cluster is represented through parameters that jointly determine the sensitivity ratio $\mathrm{SR}(t)$.

Design Cluster 2: Net Fiscal Offset to African States. This cluster determines the extent to which NZF revenue is returned to African economies in forms that mitigate the domestic cost burden. Cluster 1 determines the shock magnitude; Cluster 2 determines the offset magnitude. Included topics: disbursement categories, eligibility, and allocation logic under regulation 41 (1.7); financing instruments and access modalities (1.8).

Treatment of Bucket Y Topics. Eighteen topics are classified as material but not reducible. Analysis reveals heavy clustering around two vulnerability points:

- **Fund delivery risk (topics 1.1–1.4).** Implemented as a scenario toggle setting the effective usability factor to zero ($\omega_{\text{effective}} = 0$) at 2030.
- **Certification and LCA bottleneck risk (topics 3.5, 3.7, 3.8, 6.3, 6.4).** Captured through the effective Tier 2 compliance price parameter (T_2), which determines the two-tier adjustment factor α.

2.4 Quantification: The 7-Step Reduced-Form Model

The quantification proceeds in six analytical steps, implemented as a 7-step computation chain. The NZF Design calculator implements the identical computation chain described here.

2.4.1 Baseline and Impact Envelope CIA Scenario 24 (WtW GFI with flexibility mechanism, no levy, no feebate) is adopted as the reference point. S24 is a policy scenario, not a no-policy baseline — all outputs are relative to the CIA's BAULG scenario. S24 does not model the NZF's two-tier RU structure, reward mechanism, or revenue recycling; these NZF-specific features are what the reduced-form decomposition adds.

The reference MLC values (CIA Task 3, Table 8) are:

Year	MLC Increase (Φ_{S24})	Interpretation
2030	+5.78%	Near-term; dominated by RU payments; highest confidence
2040	+26.06%	Mid-transition; fleet shifting to low-carbon fuels; moderate confidence
2050	+35.52%	Long-term; near-complete decarbonisation; lowest confidence

2.4.2 The Compliance Cost Channel The MLC increase under the NZF decomposes into two additive components: a *base compliance cost* and a *reward distortion*.

Base compliance cost. The NZF's two-tier RU structure creates a fundamentally different cost profile from S24's single-tier flexibility mechanism. The two-tier adjustment factor α captures the ratio of fleet-weighted NZF compliance cost to S24 compliance cost:

$$\alpha(T_2) = \frac{1210 + 4.65 \times T_2}{2301}$$

where 1210 is the fleet-weighted Tier 1 compliance cost ($/tCO_2e$), 4.65 is the fleet-weighted Tier 2 cost coefficient, and 2301 is the S24 fleet-weighted compliance cost ($/tCO_2e$). The full derivation is in Annex C. Key values:

Effective Tier 2 price	α	Interpretation
$380/tCO2e	1.29	SU market at RU2 ceiling; NZF ~29% above S24
$300/tCO2e	1.13	Some SU supply pressure; NZF ~13% above S24
$235/tCO2e	1.00	NZF cost equals S24
$150/tCO2e	0.83	SU at S24 clearing price; NZF ~17% below S24

Base compliance cost is:

$$\text{base\_MLC}(t) = \Phi_{\text{S24}}(t) \times \alpha(T_2) \times \pi$$

where π is a fuel price factor (0.85–1.15) reflecting exogenous uncertainty identified in DNV's sensitivity analysis.

Reward distortion. Rewards are internal transfers from the Net-Zero Fund to ship operators using ZNZ fuels. Under efficient price discovery, each reward dollar approximately covers the cost gap between the ZNZ fuel and the ship's next-best alternative, inducing a switch that would not otherwise occur.

$$\text{reward\_distortion} = \frac{R}{M} \times \sigma \times \eta \times \varphi \times \mu$$

where R is total NZF revenue (~$10 billion/year at 2030); M is total global maritime logistics cost (~$800 billion); σ is the reward budget share (0.20–0.80); η is reward calibration efficiency (0.55–0.88); φ is a fuel scope factor (1.00–1.05); and μ is an SU market factor (1.00–1.06).

$R/M \approx 1.25\%$ establishes the scale: the maximum possible reward distortion (all revenue to rewards) would be ~1.3% of MLC.

The parameters η, φ, and μ capture design choices within the reward and compliance mechanism. Their ranges are narrow because the underlying analysis demonstrates small effects on near-term compliance cost:

- **Fuel scope** ($\varphi = 1.00$–1.05). Under a fixed reward budget, restricting eligibility increases the per-unit cost gap but decreases ZNZ volume purchased per reward dollar — effects that largely offset.
- **Reward calibration** ($\eta = 0.55$–0.88). η captures the fraction of reward budget that bridges actual cost gaps vs. the fraction flowing as windfall to inframarginal producers. The lower bound reflects a flat-rate mechanism; the upper bound reflects a competitive pay-as-bid auction.
- **SU market design** ($\mu = 1.00$–1.06). The 2030 SU market is supply-constrained, so SU prices converge toward the $380 ceiling regardless of design.

Sensitivity ratio. The total MLC increase and sensitivity ratio are:

$$\text{total\_MLC}(t) = \text{base\_MLC}(t) + \text{reward\_distortion}$$

$$\text{SR}(t) = \frac{\text{total\_MLC}(t)}{\Phi_{\text{S24}}(t)}$$

At default values ($\sigma = 0.50$, $\eta = 0.88$) for 2030: total_MLC $= 7.48\% + 0.55\% = 8.03\%$; SR $= 8.03\%/5.78\% = 1.39$.

2.4.3 Complete Parameter Summary

Parameter	Symbol	Value/Range	Source
S24 MLC anchor	$\Phi_{\text{S24}}(t)$	5.78% / 26.06% / 35.52%	CIA Task 3
Effective Tier 2 price	T_2	$380 (default); $100–$380	UCL/UMAS (2025)
Two-tier adjustment	$\alpha(T_2)$	1.29 (default); 0.73–1.29	Derived from SU price
Fuel price factor	π	1.00 [0.85–1.15]	DNV sensitivity
Base revenue	R	~$10B/year	Fleet-wide RU payments; UCL/UMAS (2025)

Parameter	Symbol	Value/Range	Source
Total MLC	M	~$800B	CIA Table 13 + logistics/transport ratio
Reward budget share	σ	0.20–0.80	Negotiation space
Allocation scheme	—	Categorical (5 options)	CIA Table 23
Usability factor	$\omega = \delta \times \lambda$	0.10–0.63	GCF/UNEP data; Batini et al. (2014)
Fuel scope factor	φ	1.00–1.05	Analytical derivation (Annex D)
Reward calibration	η	0.55–0.88	UCL/UMAS; IRENA (Annex D)
SU market factor	μ	1.00–1.06	Analytical derivation (Annex D)
Investment multiplier	μ_{invest}	0.8	Literature estimate

2.4.4 The Fiscal Transfer Channel The NZF generates approximately $10–12 billion per year in revenue during 2028–2030, from fleet-wide Tier 1 and Tier 2 RU payments. A portion is allocated to ZNZ rewards; the remainder is available for disbursement to States. The fiscal transfer to a given country is:

$$\theta_{\text{country}} = \frac{r_{\text{country}}}{\text{GDP}_{\text{country}}} \times \omega$$

where r_{country} is the country's share of the disbursement pool ($R \times (1 - \sigma)$), determined by the CIA Table 23 allocation formula: per-capita disbursements proportional to GDP impact magnitude, with the eligible set varying by scheme (SIDS and LDCs only, all developing, or all countries).

The usability factor ω ranges across five modalities:

Modality	δ	λ	ω
Technical assistance	0.35	0.3	0.10
Earmarked loans	0.55	0.4	0.22
Mixed instruments	0.65	0.5	0.33
Flexible in-sector grants	0.80	0.6	0.48
Unrestricted budget support	0.90	0.7	0.63

The δ values are calibrated from GCF disbursement-to-approval ratios, bilateral climate finance data (UNEP, 2023), and instrument-specific evidence. The λ values draw on IMF fiscal multiplier estimates for EMEs and LICs (Batini et al., 2014; Ilzetzki et al., 2013), with SSA-specific evidence from Kimaro et al. (2017).

2.4.5 Country-Specific Outcomes The two transmission channels combine to estimate GDP impact, consumer price impact, and trade cost impact for each country and time period.

GDP impact:

$$\text{net\_GDP}_i(t) = \text{GDP\_cost}_i(t) + \theta_i$$

$$\text{GDP\_cost}_i(t) = \Phi^{\text{GDP}}_{\text{S24},i}(t) \times \text{SR}(t)$$

This leverages the CIA's GTAP model, which captures country-specific trade structure, import dependence, commodity composition, and bilateral route exposure. Country-specific pass-through rates are implicitly embedded in the CIA results.

Cross-scenario validation. Within the model's operating range of SR $\in [0.85, 1.45]$, the mean prediction error is approximately -5% (conservative: the model slightly overstates impact magnitude). The default operating point (SR $= 1.39$) falls above the highest directly validated CIA scenario (SR $= 1.29$) and is interpolated from adjacent scenarios. An affine robustness specification (Annex F) produces 6–10% less-negative GDP estimates at the default operating point, confirming that the proportional model's overstatement provides a safety margin. Beyond SR $= 1.45$, prediction error exceeds 10% and results should be treated with caution.

Consumer price impact:

$$\text{CPI}_i(t) = \Phi^{\text{CPI}}_{\text{S24},i}(t) \times \text{SR}(t)$$

For countries with structural data, CPI is decomposed into food and non-food components using import composition weights and commodity-specific maritime margins.

Trade cost impact:

$$\text{import\_cost}_i(t) = \Phi^{\text{IMP}}_{\text{S24},g}(t) \times \text{SR}(t) \times \rho_i$$

where ρ_i is a country-specific import composition adjustment derived from UN Comtrade data and CIA Table 12 commodity-specific CIF-FOB margins. Values range from 1.01 (India) to 1.43 (Côte d'Ivoire).

2.4.6 Sensitivity, Interaction, and Uncertainty **Sensitivity hierarchy.** The Fund revenue split (σ) is the most consequential policy lever: for Kenya, this single parameter shifts net GDP impact by ~0.04 percentage points — larger than the combined effect of all other design parameters. Disbursement design is second-order. Reward and compliance mechanism design (φ, η, μ) collectively contributes ~0.02 percentage points of GDP variation — an order of magnitude below the Fund revenue split.

Exogenous uncertainty. Fuel price uncertainty ($\pi \in [0.85, 1.15]$) generates approximately ± 1.1 percentage points of MLC variation around the central estimate for 2030 — roughly comparable to the entire range of policy-variable effects.

Temporal confidence gradient. The 2030 results carry the highest confidence. At 2040, the two-tier adjustment converges to $\alpha = 1.0$. At 2050, all scenarios converge to 34.7–36.8% MLC; design choices are largely irrelevant. Fund revenue collapses to ~$1–3 billion at 2050 as fleet decarbonisation eliminates the RU payment base; the 2050 fiscal offset displayed in the calculator is therefore overstated by approximately 3–10×.

Correlated negotiation outcome scenarios. Three scenarios capture the covariance of negotiation outcomes:

Scenario	σ	Allocation	ω	η	φ	μ
Africa-Favourable	0.35	All developing	0.48	0.80	1.00	1.00
Compromise	0.50	All developing	0.33	0.70	1.00	1.00
Africa-Unfavourable	0.65	SIDS/LDC only	0.22	0.60	1.05	1.03

Implementation risk. Fund delivery delay sets $\omega_{\text{effective}} = 0$ at 2030, eliminating the fiscal offset for the near term. The toggle has its largest effect precisely when policy design is most favourable to developing countries.

2.4.7 Correspondence to the NZF Design Calculator

Calculator Control	Paper Reference	Parameter	Range
Fund Revenue Split (Slider 0)	§2.4.2, reward distortion	σ	0.20–0.80
Disbursement Allocation (Slider 1)	§2.4.4, fiscal transfer	allocation_scheme	5 options
Disbursement Modality (Slider 2)	§2.4.4, fiscal transfer	$\omega = \delta \times \lambda$	0.10–0.63
Fuel Eligibility (Slider 3)	§2.4.2, fuel scope	φ	1.00–1.05
Reward Calibration (Slider 4)	§2.4.2, reward calibration	η	0.55–0.88
SU Market Design (Slider 5)	§2.4.2, SU market	μ	1.00–1.06
Fuel price whiskers	§2.4.6, uncertainty	π	0.85–1.15
CIA S24 Reference overlay	§2.4.1	$\text{SR} = 1.0$	Toggle

Calculator Control	Paper Reference	Parameter	Range
Fund delivery delay	§2.4.6, implementation risk	$\omega_{\text{eff}} = 0$ at 2030	On/Off
SU / Tier 2 price	§2.4.2, two-tier adjustment	$T_2 \to \alpha$	\$100–\$380

3. Economic Implications of NZF Design Choices

3.1 Aggregate Impact

At the 2030 horizon, the NZF increases global maritime logistics costs by between 7.7% and 8.6% above the CIA's business-as-usual growth scenario, depending on the design configuration adopted. For the five African reference economies (AF5), the resulting GDP impacts at 2030 range from approximately -0.05% to -0.19% depending on the design configuration and country.

These are small in percentage terms but meaningful in absolute value. The GDP-weighted average impact across AF5 is roughly 2.3 times larger than the impact on developed economies and approximately 2.6 times larger than the impact on India — a consequence of higher trade openness, greater import dependence on seaborne goods, and more limited capacity to absorb freight cost increases.

Consumer price effects follow a similar pattern, with disproportionate food price exposure in countries with high seaborne food import dependence. Morocco (6.1% of GDP in seaborne food imports), Madagascar (4.8%), and Egypt (4.3%) face the highest food price pressure.

3.2 The Dominant Lever: Fund Revenue Split

Among the NZF design parameters, one dominates: the Fund revenue split (σ). This single parameter has the largest effect on African economic outcomes because it simultaneously affects both sides of the equation: a higher reward share increases the compliance cost borne by the global trading system, while reducing the pool of funds available to offset that cost for developing countries.

For Kenya, the Fund revenue split alone shifts net GDP impact by approximately 0.04 percentage points — larger than the combined effect of all other design parameters. This dominance is structural: σ affects both transmission channels, while the reward and compliance mechanism parameters modify only a sub-component of one.

3.3 Negotiation Scenarios

Scenario	Ghana	Kenya	S. Africa	Egypt	Ethiopia*
Africa-Favourable	−0.070	−0.129	−0.086	−0.102	−0.023
Compromise	−0.081	−0.156	−0.091	−0.112	−0.036
Africa-Unfavourable	−0.094	−0.186	−0.097	−0.123	−0.026

**Ethiopia is represented through the "Rest of landlocked economies in Africa" GTAP aggregate.*

Under Africa-Unfavourable, SIDS/LDC-only eligibility excludes Ghana, Kenya, South Africa, and Egypt from the disbursement pool, producing zero fiscal offset for those countries.

3.4 Distributional Patterns

Kenya faces the widest range of outcomes (0.057 percentage points of GDP), driven by high maritime import dependence and vulnerability to exclusion from narrow eligibility rules. South Africa's range is narrower due to its more diversified economy. The Ethiopia result illustrates a structural feature: because the composite includes LDCs, it remains eligible even under restrictive rules, and concentration effects can produce a less negative outcome under unfavourable scenarios than under compromise.

Broad eligibility is as consequential as the revenue split for middle-income African countries. Restricting eligibility to SIDS and LDCs eliminates the fiscal offset for Ghana, Kenya, South Africa, and Egypt entirely.

3.5 Temporal Dynamics

NZF design choices primarily affect the transition period (2028–2040). At 2030, the full range of design configurations produces meaningful variation. At 2040, the two-tier adjustment converges. At 2050, all CIA policy scenarios converge to 34.7–36.8% MLC regardless of design choices. Fleet decarbonisation eliminates the RU payment base, Fund revenue collapses, and the fiscal offset becomes negligible.

3.6 Implementation Risk

Even the most favourable design architecture will deliver no near-term benefit if the Fund is not operational when the NZF takes effect. The difference between a functioning and non-functioning Fund at 2030 is comparable to the entire range of variation across all policy design parameters. This supports treating early Fund operationalisation as an objective of equal importance to the revenue split and eligibility rules.

4. Robustness, Scope, and Limitations

4.1 Sensitivity to Key Assumptions

The sensitivity hierarchy is robust: σ dominates; disbursement design is second-order; reward and compliance mechanism parameters collectively contribute variation an order of magnitude smaller. This holds even with substantially wider parameter ranges (Annex D confirms σ accounts for approximately 75% of total GDP impact variation).

Fuel prices are the primary source of exogenous uncertainty. For any given policy configuration, the fuel price band ($\pi = 0.85$–1.15) produces variation roughly comparable to the entire range of policy effects.

The proportional scaling model is validated against nine CIA cross-scenario results at the 2030 horizon. An affine specification check (Annex F) confirms the proportional model provides a conservative safety margin.

4.2 Analytical Boundaries

The reduced-form approach does not capture several economic channels that a full general equilibrium model would include:

- **Exchange rate and balance-of-payments effects** — particularly relevant for economies with managed exchange rates where persistent import cost increases could trigger adjustment. Direction of bias: understatement.
- **Strategic behaviour in freight markets** — evidence from 2021–2023 container shipping disruptions suggests carriers amplify cost shocks. Direction of bias: likely understatement.
- **Structural changes in the trade environment** — including the AfCFTA. Direction of bias: uncertain, increasing with time horizon.

The omission of these channels means results are best interpreted as first-order impact estimates. The direction of bias is predominantly toward understatement.

4.3 Relationship to the CIA and CGE Analysis

The analysis is designed as a complement to the CIA, not a substitute. The CIA provides the reference case, the country-specific impact structure, and the validation benchmarks. The reduced-form analysis adds the design-decomposition layer that the CIA's fixed-scenario structure cannot provide. Where the two approaches diverge, the divergence signals general equilibrium effects that delegations should consider qualitatively.

5. Conclusions

1. **The NZF's cost burden on African economies is real but bounded.** At the 2030 horizon, GDP impacts for AF5 range from approximately -0.05% to -0.19% depending on the design configuration. These are disproportionate relative to developed economies but economically manageable.

2. **The Fund revenue split is the single most consequential design parameter.** It drives the largest variation in both compliance cost and fiscal offset, and dominates all other parameters by an order of magnitude.

3. **Broad eligibility for disbursement is as consequential as the revenue split for middle-income African countries.** Restricting disbursement to SIDS and LDCs eliminates the fiscal offset for Ghana, Kenya, South Africa, and Egypt.

4. **NZF design choices primarily affect the transition period (2028–2040).** By 2050, all policy scenarios converge regardless of design.

5. **Implementation risk can negate favourable design.** A non-operational Fund at 2030 eliminates the fiscal offset entirely, producing outcomes equivalent to the worst-case design scenario regardless of the policy architecture on paper.

Results should be interpreted as first-order estimates with reporting precision of approximately ± 0.02 percentage points of GDP. The omitted channels predominantly bias results toward understatement.

Conservative Choices Register

Choice	What was chosen	Alternative	Directional bias
Cost pass-through	Embedded in CIA GTAP (full competitive pass-through)	Partial pass-through or carrier amplification	Conservative: understates CPI impact
Reward cost induction	1:1 per reward dollar	>1:1 if ZNZ supply is inelastic	Conservative: understates reward distortion
Default Tier 2 price	$380/tCO2e (RU2 ceiling)	Lower if SU supply develops	Conservative: overstates base compliance cost
Fund revenue at 2050	$10B (same as 2030)	$1–3B (fleet decarbonisation)	Anti-conservative: overstates 2050 fiscal offset
Proportional scaling	Linear SR scaling of CIA baselines	Affine specification (Annex F)	Conservative: overstates impact by ~5%
Exchange rate effects	Not modelled	Would amplify import cost shocks	Conservative: understates impact
Strategic carrier behaviour	Not modelled	Cost amplification (2021–2023 evidence)	Conservative: understates CPI impact
Ethiopia representation	GTAP "Rest of landlocked Africa" aggregate	Ethiopia-specific model	Likely conservative: understates corridor vulnerability

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.).

CIA and Impact Assessment

- MEPC 82/INF.8/Add.1 — Report of the Comprehensive Impact Assessment, Task 2 (DNV). Secretariat, 2024.
- MEPC 82/INF.8/Add.2 — Report of the Comprehensive Impact Assessment, 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.
- ITF-OECD (2022). *Carbon Pricing in Shipping.* International Transport Forum, OECD.

African Economic Impact

- AFPH (2024). *Navigating Climate Action: Assessing the economic impacts and trade-offs of a shipping carbon tax for African States*. Africa Policy Research Institute / African Future Policies 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.
- 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, 2019). *Renewable Energy Auctions* series. International Renewable Energy Agency.
- IRENA (2020, 2022). *Green Hydrogen* series. International Renewable Energy Agency.

Maritime Trade and Transport

- UNCTAD (2024). *Review of Maritime Transport 2024*.
- WTO (2024). *World Trade Statistical Review 2024*.

Trade and Economic Data

- United Nations Comtrade Database. HS 2-digit bilateral trade data. Accessed January 2026.
- World Bank. World Development Indicators (WDI). Accessed January 2026.
- FAO. Food balance sheets and import-dependence indicators. FAOSTAT. Accessed January 2026.
- IIASA. SSP2 GDP projections (January 2024, constant 2017 USD). Accessed January 2026.

Methodological Literature

- Caron, J., Rausch, S., and Winchester, N. (2015). "Comparing Applied General Equilibrium and Econometric Estimates of the Effect of an Environmental Policy Shock." *JAERE*, 2(3), 461–490.
- Fourth Tack (2026). "Shipping, carbon pricing, and inflation pass-through: The Silent Absorber hypothesis." *Research in Transportation Economics*.

Appendix A: Source Files and Reproducibility

Source data

File	Description	Provider
nzf_calculator_data.json	Country data, S24 impacts, parameters	Fourth Tack (compiled from CIA Task 3, UN Comtrade, WDI, IIASA)
engine.js	Computation engine	Fourth Tack
calculator-content-en.json	UI text and slider definitions	Fourth Tack

Verification

The NZF Design calculator at tools.fourthtack.com/nzf/design/ implements the identical computation chain described in this document. Any parameter configuration can be reproduced interactively.