

Annex F

Bounded partial attribution

What kind of mathematical object this is — and what it is not.

A self-contained technical note, for readers coming from mathematics or the quantitative sciences.

No part of the monograph is required to read it.

Read as a model, this annex will disappoint

The recurring reaction from mathematical readers

That no clear model is presented, and that the text reads as a succession of ideas criticising a system rather than as a construction one could test.

That reaction is correct, and the fault is in the presentation.

Annex F contains no state equation, no dynamics, no solution concept and no theorem. Read as a model, it is bound to look like a list of caveats. It is not a model: it is a measurement protocol, and the annex never says so.

It does have a mathematical object. The annex names it once — in a subordinate clause on page 255 — and then never returns to it. This note puts that object first, and everything else after.

If you read four slides: 3, the object · 9, the whole construction in one view · 23, the four open questions · 24, three propositions to establish or refute.

The object: a partial identification problem

Not a dynamical system. A set-valued inverse problem.

From an observed path, and a counterfactual that is declared but not unique, recover the SET of attributable shares consistent with the data and with pre-registered structural restrictions — and report the set, not a point.

The lineage is Manski (1990): partial identification. The annex cites it once, at §F.4, and the citation carries the whole design — the bounds are bounds of documented judgement or of partial identification, and explicitly not confidence intervals.

Once the object is named, the standard questions become tractable: is the feasible set non-empty, convex, compact? Is the solution correspondence continuous in the counterfactual? Is the reported bound a bound in the strict sense?

Those four questions are gathered at the end. Everything between here and there is the construction.

The question, before any notation

Jurisdiction j tightens monetary policy. Country i 's sovereign pays more on its external debt. How much of that additional cost is attributable to j ?

Four reasons the arithmetic answer is wrong

- **The counterfactual is not unique.** “What j would otherwise have done” is a modelling choice, and different admissible choices give different answers. There is no observation to appeal to.
- **The exposure is endogenous.** Debt measured after the shock already reflects refinancing, reserve depletion and policy responses caused in part by the shock itself.
- **The channels overlap.** One exchange-rate movement moves debt service, remittances and the balance sheet. Subtracting pairwise overlaps is unstable beyond two channels.
- **Gains and losses are borne by different agents.** A gain to the central bank does not offset a loss borne by a household unless a transfer is documented.

A worked line — the exposure

Mechanism A2, the floating-rate surcharge, for the sovereign of country i against source jurisdiction j , in year t .

What is verified before anything is computed

- external floating-rate principal, pre-shock
- indexed to a single reference rate
- reset dates verified
- no caps, no floors, no hedges
- payment status confirmed

$$E^0 = 2.0 \text{ billion}$$

Exposure fixed BEFORE the shock, or before its credible anticipation — not the exposure observed afterwards.

This is the first substantive choice. Using the observed exposure would treat a post-shock adjustment as though it had been predetermined, and would quietly attribute the country's own response to the source jurisdiction.

Two admissible counterfactuals, two answers

The observed reference rate is 5.5%. Both paths below are admissible: each is a policy path j could lawfully and plausibly have followed.

c^o constant-rate benchmark

counterfactual rate **0.5%**

gap to observed **5.0 points**

gross incidence $\Delta\tilde{X}$ **100 million**

c^N gradual normalisation

counterfactual rate **2.0%**

gap to observed **3.5 points**

gross incidence $\Delta\tilde{X}$ **70 million**

$$\Delta\tilde{X} = E^0 \times (\text{observed rate} - \text{counterfactual rate}) \times \text{accrual fraction}$$

$$\omega(A2, k) = 1, \quad \text{so} \quad \Delta X = \Delta\tilde{X}$$

Neither number is more true than the other. They answer two different, equally legitimate questions.

Attribution, and what comes out

The attribution coefficient α is not 1. Table F2 gives the verified-reset range $\alpha \in [0.90, 1.00]$; the register records why the lower bound is not one — a minor tranche whose reset clause is disputed.

The A2 line for the sovereign

90 – 100 million under c_0 **63 – 70 million** under c^N

reported by counterfactual, and never merged into a single figure

This is the design, in one line.

The output is not a number with an error bar. It is a family of intervals indexed by the counterfactual — and the annex forbids collapsing that family into a point. Everything that follows exists to make that family computable, inspectable and refutable line by line.

Sets and indices

i, j	recipient country; source monetary jurisdiction	<i>attribution is bilateral before any consolidation</i>
$h \in H$	affected agent or institutional sector	<i>the distribution comes first, the national aggregate second</i>
$m \in M$	transmission mechanism	<i>reporting channels — not necessarily independent causal primitives</i>
$k \in K$	primitive economic component	<i>each component is counted once, and once only</i>
$c \in \Gamma$	admissible counterfactual	<i>no incidence is defined without a declared, internally consistent c</i>
$\theta \in \Theta$	measurement and behavioural parameters	<i>uncertainty is not confined to the attribution coefficients</i>
$\alpha \in \Phi(c)$	jointly feasible attribution vector	<i>coefficients are jointly constrained, never chosen channel by channel</i>
d_{iht}	currency, price base, discounting factor	<i>one unit, one price base, one reference date</i>

The whole object, as a chain of maps

The annex defines each map in its own section and never displays the composition. Every arrow below is a place where a claim can be lost.

Γ	admissible counterfactuals	$c \in \Gamma \mapsto z_t(c)$	<i>a vector of jointly coherent paths, not a policy rate</i>
g_m	mechanism incidence	$(E^0, z_{\text{obs}}, z_t(c), \theta) \mapsto \Delta X_m^{\sim}$	<i>exposure fixed before the shock, or the timing is unusable</i>
Ω	primitive components	$\Delta X_m^{\sim} \mapsto D_k, \omega_{mk}, \sum_m \omega_{mk} = 1$	<i>overlap resolved before aggregation, never by ex post discount</i>
α	attribution	$D_k \times \alpha, \alpha \in \Phi(c) \subseteq [0, 1]^{ M }$	<i>jointly constrained — never one channel at a time</i>
Σ_h	aggregation by agent	$\mapsto C_{ijh}, O_{ijh}, N_{ijh}$	<i>distribution first; four registers that never sum</i>
min/max	bounds	$[N^-, N^+] = \text{extrema over } \Phi(c) \times \Theta(c)$	<i>extrema over a joint set, not an interval per coefficient</i>

The stopping rules of §F.15 are not a seventh step: they restrict the domain of every map above. A rule that fails returns “not admissible” — not a weaker number.

The counterfactual is a vector of paths, not a rate

$$z_t(c) = \{ i_{j,t}(c), y^{rf}_{j,t}(c), e_{i,t}(c), risk_t(c), y_{j,t}(c), p^{com}_t(c), q_t(c), \dots \} \quad (F-2)$$

Each mechanism depends on a jointly determined set of prices, quantities and expectations. Reducing c to a policy-rate path would make the channels mutually inconsistent: the exchange rate implied by one would contradict the yield curve implied by another.

Admissibility of Γ

A path is admissible only if the source jurisdiction could lawfully and plausibly have followed it, given its mandate and the information available at the time.

A counterfactual price is a model output. Never an observation.

The polar benchmark c_0

Holding rates at their end-2021 level is admissible as a polar benchmark.

It is NOT an upper bound by definition. Its position within the range is an empirical question — and the annex supersedes the body of the text where the latter called it an upper bound.

Measurement, and the timing of exposure

$$\Delta \tilde{X}(\text{fix})_{ijhmt}(c) = g_m(E^0_{ijhm}, z_t(\text{obs}), z_t(c), \theta_m) \quad (\text{F-3})$$

$$\Delta \tilde{X}(\text{dyn})_{ijhmt}(c) = g_m(E(\text{obs})_{ijhm,t-1}, z_t(\text{obs}), z_t(c), \theta_m) \quad (\text{F-4})$$

g_m is the mechanism-specific measurement function. The two lines differ only in which exposure they use — and that difference is the whole point.

Why fixed exposure is the preferred basis for attribution

For $t > 1$, the observed exposure may already reflect refinancing, reserve depletion, fiscal adjustment or capital controls — responses caused in part by the shock. A time lag does not, on its own, make an exposure exogenous.

Where the shock alters the exposure base, the analyst must model the feedback, instrument the exposure, freeze the base at E^0 , or report both with an explicit decomposition of adaptation effects. Dynamic measurement remains informative, but it describes incidence realised after adaptation and is reported separately.

The feasible set $\Phi(c)$ — attribution is joint, not per-channel

$$\alpha_{ij}(c) = (\alpha_{ijm}(c) : m \in M) \in \Phi_{ij}(c) \subseteq [0,1] \times \dots \times [0,1] \quad (F-5)$$

Φ contains the per-mechanism bounds, the cross-mechanism restrictions, common-shock coherence, sign restrictions, and any constraint induced by the empirical design. It is pre-registered before results are computed.

Structural restrictions (Table F2)

A2 verified contractual reset $0.90 \leq \alpha \leq 1.00$

A3 identified policy component $\alpha = 1$

A3 term, expectations, liquidity $0 \leq \alpha < 1$

A5 pre-existing idiosyncratic premium $\alpha = 0$

A7 domestic crowding-out $0 < \alpha^- \leq \alpha \leq \alpha^+ < 1$

B1–B2 originating supply shocks $\alpha = 0$

Why $\alpha = 1$ is almost never admissible

Before any coefficient is set at unity, the risk-free channel is decomposed:

$$\Delta y^{\text{rf}} = \Delta y(\text{pol}) + \Delta y(\text{term}) + \Delta y(\text{exp}) \\ + \Delta y(\text{liq}) + \Delta y(\text{res})$$

Only the policy component carries $\alpha = 1$. Treating the whole domestic rate response as externally caused is inadmissible.

Overlap is removed before aggregation, not after

One exchange-rate movement can move debt service, remittances and the balance sheet at once. Pairwise overlap deductions become unstable as soon as three channels share a determinant. So components are allocated once, before they are summed.

$$0 \leq \omega_{mk} \leq 1 \quad \text{and} \quad \sum_{\{m \in M(k)\}} \omega_{mk} = 1 \quad (\text{F-8})$$

$$\Delta X_{ijhmt}(c) = \sum_{\{k \in K_m\}} \omega_{mk} \cdot D_{ijhkt}(c) \quad (\text{F-9})$$

D is the monetary value of primitive component k; ω is its published allocation share to reporting mechanism m. $\Delta \tilde{X}$ stays published as the gross per-channel figure; ΔX is the overlap-net amount that enters the aggregate.

Every primitive component carries a unique identifier, a source, a sign, a unit, an owner, an allocation rule and a version. The weights may place a component entirely in one mechanism or spread it across several — but they must sum to one. An unpublished overlap adjustment renders the identity unverifiable.

A gain to one agent does not cancel a loss to another

The primary object is the per-agent vector, not the national total. An offset accruing to agent ℓ reduces a loss borne by agent h only to the extent that a transfer is actually documented.

$T(\ell \rightarrow h, t) \in [0, 1]$ share of a gain accruing to ℓ that reaches h in period t

In the absence of a documented or legally committed transfer, $T = 0$ and no netting occurs. An inflow from an earlier easing phase is not set against a later reversal — least of all when it increased leverage or foreign-currency exposure.

The consequence is arithmetical, and deliberate

Refusing to net across agents produces a larger figure than the legacy national net. That difference is not a revision of the facts. It is the price of refusing to offset a loss with a gain that never reached the party bearing it.

Both may be reported, labelled “distributive consolidation” and “legacy national net (robustness)”.

What the framework outputs

$$C_{ijh}(c, \alpha, \theta) = \sum_t d_{iht} \cdot \sum_{\{m: s_m=+1\}} \alpha_{ijm}(c) \cdot \Delta X_{ijhmt}(c, \theta) \quad (F-12)$$

$$O_{ijh}(c, \alpha, \theta) = \sum_t d_{iht} \cdot \sum_{\{m: s_m=-1\}} \alpha_{ijm}(c) \cdot O^*_{ijhmt}(c, \theta) \quad (F-13)$$

$$N_{ijh}(c, \alpha, \theta) = C_{ijh} - O_{ijh} \quad (F-14)$$

The primary result is the distributive vector $N_{ij}(c) = (N_{ijh}(c) : h \in H)$. A consolidated total may be reported, but only labelled an accounting consolidation — never a welfare measure.

$$\kappa_{ijh}(c) = C_{ijh}(c) / G_{ijh}(c) \quad \iota_{ijh}(c) = N_{ijh}(c) / G_{ijh}(c) \quad (F-19)$$

None of these ratios is a causal probability, a percentage of legal liability, or a share of welfare loss. ι may be negative if documented offsets exceed attributable adverse incidence. Nothing may be published until the shock, the population, the window, the price base, the component allocation and the denominator are all fixed and reproducible.

What is never added to what

$$S_{ijht}(c) = D^{FX}_{ih,t-1} \cdot [e_{i,t}(obs) - e_{i,t}(c)] \quad (F-20)$$

S is the exchange-rate revaluation of the external debt stock: a balance-sheet exposure, not a cash flow of the period. It is reported separately and is never added to C, O, N or G.

Four registers coexist and never sum

the operation and its settlement

a cash flow of the period

asset revaluation

a balance-sheet position

income

a flow accruing to an identified agent

exposures

a stock that conditions future flows

This is stated as a discipline of prudence. Whether it can instead be derived — as a failure of commutation between the aggregation operator and the sum — is the fourth of the open questions at the end of this note.

The identity is designed to stop as much as to compute

These are constitutive rules, not optional caveats. Each blocks publication.

No counterfactual, no figure. every line must identify c and the model generating $z_t(c)$

No pre-shock or explicitly modelled exposure, no causal interpretation. dynamic realised incidence may still be reported, descriptively

No allocation into primitive components, no aggregate. an unresolved overlap does not disappear through residual deduction

No identified beneficiary or documented transfer, no distributive offset. gains are not presumed to reach those who bear the losses

No pre-registered Φ , no envelope of extremes. an unregistered Φ is deemed absent, and the envelope is blocked

The annex disqualifies the book's own headline numbers

The magnitudes in the body of the monograph — the Tunisian monetary core of €0.76–1 billion net, and the $\approx$ €4.5 billion gross, all causes — were computed under the earlier version of this annex.

Measured against the constitutive rules above, they do not constitute verifiable aggregates.

What they lacked: a single dominant counterfactual described as an upper bound; a pairwise seam correction rather than component allocation; no register of primitive components; no agent breakdown; no pre-registered Φ .

They are requalified as disciplined illustrations under the legacy framework, retained to make the anatomy legible, and flagged wherever they appear. Their recomputation is the first deliverable of the empirical programme.

I put this slide in deliberately. The silent coexistence of two standards would hand critics of the architecture precisely the weapon this annex exists to take from them.

The annex is not only prose — it runs

A workbook applies the framework to the case the monograph disqualified. Ten sheets, eighty-seven live formulas, and a declaration of its own status on the first sheet.

“Pre-Configuration M. What is missing for Configuration M is listed here.”

sheet 1, Déclaration — the workbook does not claim to be what it is not

15 channel lines, one column headed “Additive?”	Yes · No · Without deduction — non-additivity made operational, line by line
Channel 5b, cereal and energy supply shock	carries $\alpha = 0$, as Table F2 requires
Stock revaluation S	memo item only, never added to C, O, N or G (§F.12)
Offsets O1–O3	reported by beneficiary, never subtracted (§F.7)
Two counterfactuals	co and c^N , each carried through to the end
Reconciliation with the legacy figures	closed, and line by line

One row computes the attractive figure — extended consolidation minus offsets — and labels it: not admissible as a principal result.

The same exposure, under two counterfactuals

Direct external cash-flow core, cumulative 2022–2025, one recipient country. Same data, same channels, same rules.

co constant-rate benchmark

298 – **682** – 1 430

million euros — low, base, high

c^N neutral-rate path

199 – **455** – 953

million euros — low, base, high

A 33% swing in the headline figure, produced by nothing but the choice of counterfactual — both of them admissible.

This is why §F.10 requires bounds reported by counterfactual, and why a single number would be a claim the data cannot carry.

The workbook states the approximation openly: c^N is obtained by a documented two-thirds factor, and is to be replaced by a modelled $z_t(c^N)$. It is a transitional artefact, and says so.

What this annex does not, and cannot, establish

Equations (F-1)–(F-20) discipline the accounting and the bounds. They do not validate the inputs.

The values of α

their joint restrictions require independent estimation, or a defensible partial-identification design. The annex constrains them; it does not estimate them.

The inputs g_m , D_k , ω_{mk} , T

these require published data and an audit. The framework is only as good as the register behind it.

Causal identification

requires an identified shock, exposures fixed or instrumented, multiple counterfactuals, cross-channel interactions, placebo and pre-trend tests, and heterogeneity by regime. That is Annex G, and it is not done.

Until that work is done and replicated, the output is an accounting incidence with bounded partial attribution. It can inform measurement and scenario design. It cannot establish fault, liability, compensation or price.

The identification was attempted — and refuted by its own test

A companion pipeline implements the Annex G design. It is pre-registered, reproducible, and it reached a negative conclusion about itself.

What was built

- pre-registration deposited before estimation
- Fed shocks: Jarociński–Karadi, pure and information
- ECB shocks: EA-MPD surprises
- exposures predetermined, ten debtors
- 36 tests: 2 banks $\times$ 3 exposures $\times$ 2 outcomes $\times$ 3 horizons
- Benjamini–Hochberg FDR at $q = 0.1$, plus placebo

What it found

Seven of the thirty-six tests survive FDR control.

The placebo then fails: a future Fed shock explains present growth.

The survivors are therefore not credible. Persistent policy cycles and persistent exposure create a spurious correlation that fixed effects do not absorb.

The annual design is invalidated for causal inference — demonstrated, not assumed. High-frequency event identification is not an option but the only credible route.

One further result was withdrawn along the way: the debt-to-GNI finding proved to be an outlier artefact — a single country — and did not survive winsorisation.

Four open questions

These are the reason for the note. None of them is settled here.

1 **Is $\Phi(c)$ non-empty, and is the solution unique?**

If Φ admits several jointly feasible vectors, the annex publishes a selection from a set. Is that set convex and compact, and what is then the honest object to report?

2 **Is the correspondence $c \mapsto \Phi(c)$ hemicontinuous?**

α is not a function of c but a set, so the question is hemicontinuity — Berge — and continuity of the bound functions $N^-(c)$, $N^+(c)$. A modulus there would turn a methodological caution into a result: the single most valuable addition.

3 **Is the reported “bound” a bound in the strict sense?**

The annex reports intervals. They are not confidence intervals; they are constrained ranges of documented judgement. Is that a defensible object, and is there a better one?

4 **Can non-additivity be derived rather than imposed?**

Four registers never sum, and the annex states this as a rule. Homogenisation is the theory of when averaging a heterogeneous medium is legitimate — could this instead be a failure of commutation between the aggregation operator and the sum?

What a proof would look like

The questions opposite ask whether something is true. These state what would have to be proved. None of the three appears in the annex.

1 Existence of the bounds

If Γ , $\Theta(c)$ and $\Phi(c)$ are non-empty and compact and N_{ijh} is continuous in (α, θ) , the extrema defining $[N^-, N^+]$ are attained.

That is Weierstrass. The work is in the three conditions, stated nowhere in the annex.

2 Single counting, derived rather than imposed

If K partitions the incidence and $\sum_m \omega_{mk} = 1$ for every k , then $\sum_m \Delta X_m = \sum_k D_k$: each component is counted once — none twice, none lost. It may still enter several mechanisms.

The annex imposes the unit sum as a rule; stated this way it follows from the construction.

3 Stopping rules as domain restrictions

Read the rules of §F.15 as restrictions on the domain of the output map rather than as warnings attached to its image.

A failed rule then returns “not admissible”, and no degraded number can be produced.

If any of the three is false, the annex is claiming something it cannot support — which is more useful to know than a confirmation.

What a model would be — and why this annex is not one

Annex F measures. It does not propagate. A separate part of the work does try to propagate, and that is where an outside mathematical judgement would matter most.

Annex F — what this note covered

- a static, set-valued measurement problem
- one period at a time, indexed by t
- no state equation, no propagation
- output: a family of intervals
- epistemic rungs 1 and 2

§9.8 and the circuit topology — not covered here

- cascades: protection, and accumulation-reversal
- an impedance vector, and network topology
- a linear-response regime against a threshold regime
- this is where an analogy could be doing no work
- and where physics may be imported unearned

If Annex F survives scrutiny, that is the second discussion — and the harder one.

The materials behind this note

Everything below exists, is versioned, and can be sent on request. In reading order.

1 The Configuration M workbook

Annex F applied end to end: ten sheets, live formulas, two counterfactuals, and a closed reconciliation with the legacy figures.

2 The empirical pipeline

Pre-registration, data dictionary, provenance register for eighteen series, six scripts, shock files, and the synthesis reporting the placebo failure.

3 The evidence register

Three corpora with dated changelogs: the financing mapping, the register of primary instruments, and the fifteen-mechanism schema.

4 The monograph itself

Annex F occupies pages 253–262; Annex G is the empirical bridge; Annex J extends the protocol internationally.