Online Resource 1

Article	Delayed PFAS release under market restrictions: a cohort-resolved substance flow model for the European Economic Area
Journal	Journal of Industrial Ecology
Authors	Matthias Falzberger, Anke Bockreis

Corresponding author	Matthias Falzberger, Unit of Environmental Engineering, Department of Infrastructure Engineering, University of Innsbruck, Innsbruck, Austria, matthias.falzberger@uibk.ac.at

This file documents the model boundary, input provenance, simulation algorithm and assessment procedures. All results, checks and precision statements come from the final simulation campaign (engine input digest `891f3b8713626a1a`, 29 sets), which follows the correction of the fluorinated-gas derogation scope and the scope review of the comprehensive restriction (assumptions A551–A554), and from its extension of 2026-09-15 (digest `a8bb356d207eddc9`, 12 sets), which added the variants VE1, VE2, VS1 and VS2 and changed no default row (Section S8.2). The pack count was decided on the preceding campaign's input state and confirmed on the final one (Section S5).

Machine-readable copies of the input tables are supplied separately as Online Resource 2 (`.xlsx`). Where a table here is a summary of a dataset, the dataset itself is in that file and is named in the legend. In that file every identifier is a hyperlink to the row that defines it, so an assumption number leads to its entry in C2 and a reference number to its entry in C1. Online Resource 3 (`.xlsx`) supplies the annual series behind the results, boundary flows per substance class and stocks per node, 1950–2070, for the three default sets, and the pack-count ladder of Figure S1, with the archive files read and their hashes. Every number shown in a figure of this file is in a table of this file or in Online Resource 3.

How to use this file. A number quoted in the main text is in Section S7 (result tables, three significant figures, with its Monte Carlo band) or in Online Resource 3 (annual series). How a parameter was set is in Section S2 (the input tables), its source and grade in Section S3 and in worksheet C2 of Online Resource 2. How the engine moves mass is in Section S4, how the pack and run counts were chosen in Section S5, what was checked in Section S6, and how to rerun everything in Section S8. Three reading rules apply to every number in this file and in the main text (Section S9): every total, difference, share and index is formed within a run and only then summarised as a median with its 5th–95th percentile band; compositions are means and carry no band; the spread between variants is a difference between parameterisations, not a probability interval.

Masses are tonnes of compound mass throughout: not fluorine-equivalent, not toxicity-weighted. The horizon is 1950–2070 and the spatial scope is the European Economic Area.

Contents

S1 Nodes and the process network **3**
- S1.1 The four bands . . . 3
- S1.2 Node definitions . . . 3
- S1.3 The four absorbing sinks . . . 4
- S1.4 Edges and the release boundary . . . 4

S2 The input tables **5**
- S2.1 Substance classes (A2) . . . 5

S2.2 Sectors and subsectors (A3) . . . 5
S2.3 Scenarios (A4) . . . 7
S2.4 Variants (A5) . . . 8
S2.5 Market inflow anchors (B1) . . . 9
S2.6 Growth rates (B2) . . . 9
S2.7 Allocation shares (B3) . . . 9
S2.8 Transfer weights (B4) . . . 9
S2.9 Residence times (B5) . . . 9

S3 Provenance and source-support grading 11
S3.1 Evidence identification and provenance chain . . . 11
S3.2 Source-support grades . . . 12
S3.3 Mass-weighted source support . . . 13
S3.4 Calculation . . . 13

S4 Simulation algorithm and equations 14
S4.1 Packs and state variables . . . 14
S4.2 Market inflow and the annual cycle . . . 14
S4.3 Drawing on arrival . . . 15
S4.4 Routing and residence sampling . . . 15
S4.5 Release and reported outcome definitions . . . 16
S4.6 Resolving a context to rows . . . 17
S4.7 Reachability and numerical safeguards . . . 17
S4.8 What is sampled and what is assumed . . . 17
S4.9 Random numbers and reproduction . . . 18

S5 Numerical resolution and Monte Carlo precision 18
S5.1 Primary outcomes . . . 18
S5.2 Pack resolution and the gate . . . 18
S5.3 Precision across runs . . . 19
S5.4 Interpretation and scope . . . 19

S6 Verification and external consistency 21
S6.1 Verification . . . 21
S6.2 External consistency . . . 22

S7 Result tables and figures 23
S7.1 Market inflow and release under the three scenarios . . . 24
S7.2 Release committed before 2027 . . . 24
S7.3 Environmental compartment stocks . . . 27
S7.4 Fate of cumulative market inflow . . . 28
S7.5 Variant results . . . 28
S7.6 Sector evidence for the derogation result . . . 32

S8 Reproducibility 32
S8.1 The campaign . . . 32
S8.2 The extension campaign . . . 32
S8.3 Input digest . . . 34
S8.4 Random numbers . . . 34
S8.5 Re-running the campaign . . . 35

S9 Reading rules for the archived results 35

S1 Nodes and the process network

The system represents the European Economic Area over 1950–2070 with 23 nodes, 49 directed links and four absorbing accounting nodes. Initial stocks are zero. Historical stocks of excluded substances, external environmental inputs and chemical transformation are not represented.

S1.1 The four bands

Nodes are grouped into four bands, identified by the node prefix. The grouping defines the release boundary (Section S1.4).

Product cycle (PC). Market entry, production and import, product-use stock, waste collection, recycling, export.

Release accounting (RL). Three junctions, one per releasing stage. They hold no mass.

Technical compartments (TC). Biosolids spreading, landfill, incineration, wastewater treatment, elimination, sewage sludge, landfill leachate. Engineered systems; mass here has not been released.

Environment (EN). Five resident compartments (soil, groundwater, surface water, air and sediment) and two accounting sinks for marine and atmospheric outflow beyond the EEA.

S1.2 Node definitions

Table S1 defines every node. Roles: SOURCE introduces mass (one node); STOCK retains mass across years and carries a residence time (Section S2.9); FLOW passes mass on within the year; AUXILIARY is a bookkeeping junction, not a physical place; SINK is never left.

Table S1: The 23 model nodes, their role and definition

ID	Name	Role	Abs.	Definition
PC0	Market entry	Source	—	Exogenous inflow to the EEA market, resolved by subsector and class.
PC1	Production and import	Flow	—	Bookkeeping and routing of placed-on-market mass from EEA production and net imports; upstream synthesis and manufacture are not modelled.
PC2	Product use	Stock	—	Product-use stock. Release leaves through RL2, end-of-life mass to PC3.
PC3	Waste collection	Flow	—	Collection and sorting of end-of-life articles.
PC4	Recycling	Stock	—	No route separates PFAS from the material; mass re-enters PC2 as recyclate or returns to PC3.
PC5	Export	Sink	yes	Articles and waste leaving the EEA. Not release.
RL1	Release, production stage	Auxiliary	—	Release originating at production and import.
RL2	Release, use stage	Auxiliary	—	Release originating at the use stage.
RL3	Release, waste stage	Auxiliary	—	Release originating at waste collection and recycling.
TC1	Biosolids spreading	Flow	—	Application of sludge and compost to agricultural land.
TC2	Landfill	Stock	—	Landfilled waste, held with a slow residence time. Leachate to TC7.

Table S1 continued

ID	Name	Role	Abs.	Definition
TC3	Incineration	Flow	—	Municipal and hazardous waste incineration; splits into mineralised mass, stack release and ash.
TC4	Wastewater treatment	Stock	—	Splits influent into effluent, sludge and a small share to air.
TC5	Elimination	Sink	yes	Accounting mass assigned to mineralisation; treatment effectiveness is represented by the incoming routing shares.
TC6	Sewage sludge	Stock	—	Sludge routed to land application, incineration or landfill.
TC7	Landfill leachate	Auxiliary	—	Leachate between generation and routing.
EN1	Soil	Stock	—	Agricultural, urban and natural soil, aggregated.
EN2	Groundwater	Stock	—	The groundwater body.
EN3	Surface water	Stock	—	Rivers and lakes.
EN4	Air	Stock	—	Atmosphere. Fluorinated gases are retained per their residence time and disperse to EN7; other classes pass through quickly.
EN5	Marine outflow	Sink	yes	Surface water load leaving the EEA to the open ocean.
EN6	Sediment	Stock	—	Freshwater and coastal sediment.
EN7	Atmospheric outflow	Sink	yes	Long-range atmospheric transport out of the EEA.

Source: `A1_Nodes.csv`, in full in Online Resource 2. *Abs.* marks an absorbing node. Nodes are never split by sector or substance: each pack carries its subsector and class as immutable keys, so a node is partitioned by context at reporting time rather than duplicated in the topology.

S1.3 The four absorbing sinks

Mass arriving at PC5, TC5, EN5 or EN7 stays there. PC5 and EN5/EN7 are terminal because the mass has left the region; TC5 because the assigned mass is treated as mineralised. EN5 and EN7 hold released mass; PC5 does not, and exported mass is not counted as avoided release.

Absorption is declared, not emergent: a node with no outgoing route would accumulate mass without violating the mass balance. A load-time reachability assertion therefore requires every reachable state to have an outgoing route or be declared absorbing (Section S4.7).

S1.4 Edges and the release boundary

A link exists where the routing table carries at least one row with that origin and destination, so the topology follows the parameterisation and cannot drift from it. Table S2 lists all 49.

A link crosses the release boundary exactly when its destination is in the EN band and its origin is not. Fifteen links satisfy this; the model derives the set from the node definitions at load time rather than from a column. EN3 → EN5 and EN4 → EN7 are inside the EN band and are therefore not release: that mass was released on entering the band.

Three further boundary classes are derived the same way: market entry (PC0 → PC1), export (PC3 → PC5) and elimination (PC3 → TC5, TC3 → TC5). 19 of the 49 links carry a boundary class.

Table S2: The 49 directed links of the network

From	To	Release links
PC0	PC1	
PC1	PC2, RL1	
PC2	PC3, RL2	
PC3	PC4, PC5, RL3, TC2, TC3, TC5	
PC4	PC2, PC3, RL3, TC2	
RL1	EN3, EN4	2
RL2	TC4	
RL2	EN1, EN3, EN4	3
RL3	EN1, EN3, EN4	3
TC1	EN1	1
TC2	TC7	
TC2	EN1, EN4	2
TC3	TC2, TC5	
TC3	EN4	1
TC4	TC6	
TC4	EN3, EN4	2
TC6	TC1, TC2, TC3	
TC7	TC4	
TC7	EN2	1
EN1	EN2, EN3, EN4	
EN2	EN3	
EN3	EN5, EN6	
EN4	EN1, EN3, EN7	
EN6	EN3	

Derived from the distinct origin-and-destination pairs in `B4_Transfer_Coefficients.csv`. Destinations are grouped by origin; each listed destination is one link. The third column counts the release links on that row.

S2 The input tables

The model uses twelve semicolon-delimited tables. Ten tables (A1–A5 and B1–B5) are parsed by the engine; two (C1–C2) document references and assumptions. The six default value tables B1–B5 and A4 contain 1056 parameter rows in the default parameterisation: 27 anchors (B1), 234 growth rates (B2), 250 allocation shares (B3), 258 routing weights (B4), 140 residence times (B5) and 147 scenario rows (A4, excluding its three scenario header rows). Node definitions determine the release boundary and the substance table determines each class's first market year, so these definition tables can also affect results.

All twelve are reproduced in Online Resource 2, one worksheet per table. Tables printed here are summaries; the dataset is the record.

S2.1 Substance classes (A2)

Seven classes in three groups (Table S3). Class is an immutable pack attribute; no transformation between classes is modelled (Section S4.8).

S2.2 Sectors and subsectors (A3)

Fifteen sectors, 72 subsectors (Table S4). The subsector, with the substance class, is the context that selects a pack's routing and residence rows. Sector totals are aggregations formed at

Table S3: The seven substance classes and their first year of market inflow

ID	Class	Group	First inflow year
G1A	Terminal PFAAs	PFAAs and their precursors	1950
G1B	Non-polymeric precursors	PFAAs and their precursors	1950
G1C	Side-chain fluorinated polymers	PFAAs and their precursors	1950
G2A	HFCs	Fluorinated gases	1990
G2C	HFOs and HCFOs	Fluorinated gases	2010
G3A	Fluoropolymers	Fluoropolymers and perfluoropolyethers	1950
G3B	Perfluoropolyethers	Fluoropolymers and perfluoropolyethers	1950

Source: `A2_Substances.csv`. Side-chain fluorinated polymers sit with the PFAAs and their precursors, following the restriction dossier (European Chemicals Agency, 2023a). There is no class G2B: the identifier was reserved for CFCs and HCFCs, which are excluded from the model, including their historical stocks. All first inflow years are read from A2 and are model assumptions.

reporting time, never simulated directly.

One subsector is not a product. S07.11, servicing of installed equipment, carries the fluorinated gas sold to top up equipment in operation (assumptions A541–A544, A546, added 2026-09-04). Its anchor share is the dossier's stock-based release of 2018, and its inflow leaves product use one year after placement under the default rounding rule. Its cutoff follows the scenario: 2038 under the partial restriction, the general cutoff under the comprehensive restriction (Section S2.3). It is a proxy for replacement gas associated with escape during use, not a stock-depletion calculation (Section S4.8).

Table S4: The fifteen sectors and all 72 subsectors

ID	Sector	Included subsectors	n
S01	textiles, upholstery, leather, apparel and carpets	Home textiles; consumer apparel; professional apparel and personal protective equipment; technical textiles; leather; home fabric treatments; automotive noise, vibration and harshness textiles	7
S02	food contact materials and packaging	Paper and board packaging; plastic packaging; other packaging; consumer cookware; industrial food production; professional bakeware	6
S03	metal plating and metal product manufacture	Hard-chrome plating; decorative plating; metal product manufacture	3
S04	consumer mixtures and articles	Consumer mixtures	1
S05	cosmetics	Cosmetics	1
S06	ski wax	Ski wax	1
S07	applications of fluorinated gases	Refrigeration; air conditioning and heat pumps; foam-blowing agents; propellants; fire suppressants; mobile air conditioning; servicing of installed equipment	7
S08	medical devices	Implantable medical devices; hernia meshes; wound treatment; tubes and catheters; metered-dose-inhaler coatings; other medical coatings; medical engineered fluids; sterilisation gases; diagnostic and laboratory testing; metered-dose-inhaler propellants; medical venting membranes; medical packaging; rigid gas-permeable contact lenses	13

Table S4 continued

ID	Sector	Included subsectors	n
S09	transport	Vehicle safety applications; hydraulic fluids; mobile air conditioning; transport refrigeration; military mobile air conditioning and refrigeration	5
S10	electronics and semiconductors	Electronics; semiconductors	2
S11	energy	Proton-exchange-membrane fuel cells; other energy applications	2
S12	construction products	Architectural coatings and paints; wind-turbine-blade coatings; coil coatings; composite architectural membranes; pure architectural membranes; ETFE greenhouse film; window frames; bridge and building bearings; PTFE thread-sealing tape; processing aids for non-PFAS polymers; side-chain surface protection; fluorosurfactant wetting agents; non-polymeric processing aids; window-film manufacture	14
S13	lubricants	Lubricants	1
S14	petroleum and mining	Non-polymeric applications; fluoropolymer applications	2
S15	firefighting foams	Municipal and regional fire services; industrial Seveso sites; aviation and airports; offshore oil and gas; marine and military vessels; portable extinguishers; training and testing	7
	Total		72

Source: `A3_Sectors.csv`; subsector names and definitions in Online Resource 2. The classification adapts the universal restriction dossier and the separate firefighting-foam assessment (European Chemicals Agency, 2023d; European Commission, 2025). Subsector resolution and the servicing subsector are modelling choices.

S2.3 Scenarios (A4)

169 rows: three header rows, 147 default rows, 19 variant rows.

What the engine reads. Two columns only: the last year carrying inflow (`last_inflow_year`) and the share of the subsector inflow the row applies to (`share_of_subsector_inflow`). Restriction status and derogation duration are labels for grouping results; the engine ignores them. To check how a scenario is implemented, read the cutoff years, not the labels.

Tranche rows. A subsector-and-group key may carry several rows, each with its own share and its own cutoff year; 26 keys across the three scenarios do. The surviving share in a year is the sum of the shares whose cutoff has not passed, as in Equation (S2). The entry-79 split, where part of a subsector's inflow stops earlier than the rest, is the case a single cutoff year cannot represent.

The anchor. Assumed model restriction date 2027, with no transition period: 2026 is the last year carrying inflow under the general prohibition, 2027 the first year at zero. Time-limited derogations measured from that anchor end in 2031 and 2038; three rows under the partial restriction carry no end date and run to 2070. The scenario cut and the vintage cut coincide at 2027, so two vintage classes suffice.

Specific model rows override the general cutoff. The packaging regulation, the fluorinated gas regulation and REACH Annex XVII entries 79 and 82 apply in every scenario (European Commission, 2024, 2025; European Parliament and Council, 2024, 2025). The most specific input row takes precedence, so dates adapted from adopted measures can be later than the general cutoff. This is a model hierarchy, not a legal rule that a later date in existing law would override a future stricter restriction. Annual cutoffs also approximate within-year application dates and concentration- or use-specific scope. Consequence: the comprehensive restriction carries inflow after 2026 in twelve rows under adopted law (Table S5); the counterfactual has 23 rows ending before 2070. Every one of those twelve rows rests on an adopted measure, so the comprehensive restriction grants no derogation of the proposed entry and all 27 of its detail rows carry `restriction_status` `BANNED`. Servicing of installed equipment, S07.11, is one of the restricted uses: the partial restriction carries it to 2038 under the derogation for maintenance and refilling of existing equipment, and the comprehensive restriction stops it at the anchor, because the restriction option it represents carries no derogations. Its volume follows prescribed growth paths informed by the fluorinated gas regulation; the model does not enforce the HFC quota, which is set in carbon dioxide equivalents, directly. The path is the same in every scenario within a parameterisation, but the cutoff is not, so servicing enters the paired scenario differences.

Table S5: Rows carrying market inflow after 2026 under the comprehensive restriction

Subsector	Name	Group	Last inflow year
S01.1	home textiles	G1	2027
S01.5	leather	G1	2027
S15.7	training and testing	G1	2027
S15.7	training and testing	all	2027
S15.3	aviation and airports	G1	2029
S15.1	municipal and regional fire services	G1	2030
S15.1	municipal and regional fire services	all	2030
S15.3	aviation and airports	G1	2030
S15.3	aviation and airports	all	2030
S15.2	industrial Seveso sites	all	2035
S15.4	offshore oil and gas	all	2035
S15.5	marine and military vessels	all	2035

Source: `A4_Scenarios.csv`, default parameterisation. Firefighting foams sit outside the universal restriction and are governed by REACH Annex XVII entry 82; the textile and leather rows carry the entry-79 date. Subsectors reached by entry 79 carry two rows, one for the share subject to that entry and one for the remainder; rows for G1 take precedence over rows for all groups. An all-group row may therefore not be selected for a G1 context. Counts of declared rows are not counts of active uses; the resolved shares and years determine the actual inflow.

Cutoff years under the partial restriction. Of 91 detail rows: 55 prohibited, 33 subject to a time-limited derogation and three to a derogation without an end date. Of the 55 prohibited rows, 43 stop at the anchor and twelve on a date set by adopted law; of the 33 time-limited rows, 31 end with their derogation tier and two on an adopted-law date that falls earlier than the tier. Cutoffs are therefore 2026 (43 rows), 2031 (10), 2038 (21), 2070 (3), and 14 rows on adopted-law dates — 2027 (4), 2029 (2), 2030 (4), 2032 (1), 2035 (3). The offshore and marine rows on entry 82 paragraph 5(c) count as prohibited from 2026-09-08: a transitional period of adopted law is not a derogation of the proposed entry, and it is present in all three scenarios alike.

S2.4 Variants (A5)

Fifteen rows: the default parameterisation, six source-informed alternatives and eight bounding variants (Table S6); VE1, VE2, VS1 and VS2 were added on 2026-09-15 after the pre-submission review (assumptions A555, A556). Running a variant replaces every default row whose logical

key matches a variant row; other variants' rows are dropped. The bounding variants were generated from the bounds the default rows declare, with the joint direction of change and the positive-weight floor specified by the authors; the generated rows are part of the deposited input tables (Online Resource 2), the generator is not. They are not guaranteed extrema of the outputs.

S2.5 Market inflow anchors (B1)

27 rows: 2020 EEA market placement per sector and group, as a triangular low, mode and high in tonnes per year. The only absolute mass in the model; everything else is a share, a rate or a time. Drawn once per run. Their contribution relative to other sampled parameters has not been isolated by a variance decomposition. The manuscript reports masses in kilotonnes.

S2.6 Growth rates (B2)

234 rows: era growth rates in per cent per year, start year inclusive to end year exclusive. Inflow runs through 2070 inclusive, so the terminal era ends at 2071. Coverage is complete for every live combination except the fluorinated gases before 1990.

Two properties govern how a sector total behaves. Allocation happens before growth, so a sector total is not its anchor grown at one rate: growing first would freeze the substance mix, and the shift from HFCs to HFOs is a change in that mix. And overlapping rows resolve by an explicit precedence rank, higher wins, rising with specificity; rates are not averaged. Section S4 gives the resulting inflow equation.

S2.7 Allocation shares (B3)

250 rows splitting each sector-and-group inflow into subsectors and classes. Values are raw weights, normalised to one within the matching set. The allocation weights define the 2020 composition. They are held fixed within a run, but different subsequent growth factors change the realised composition over time.

Each row carries its own source-support grade: 19 high, 51 intermediate, 180 low. High covers splits with a quantitative source or a structurally degenerate split; intermediate, splits whose sub-use structure is documented but whose magnitudes are partly judgement; low, expert judgement cross-checked for completeness against a published use inventory. The medical device split is low-graded; the sensitivity of results resting on it is reported through the medical variant family, and its tonnage in Table S18 is read with that grade in view.

S2.8 Transfer weights (B4)

258 default rows of routing weights per origin node and context, plus 186 rows each for the two routing variants. Weights out of one node and context normalise to one and become drawing probabilities; all 118 default context sets already sum to one at the mode.

204 default rows declare a triangular range, 54 a point value. The range is not sampled: the modes are normalised within a resolved set to obtain probabilities. The declared bounds enter the results through the routing variants VR1 and VR2 only, as a bracket between runs (Section S4.8). The `flow_class` column is descriptive and unread. The release boundary is not a column here; it is derived from the node definitions (Section S1.4).

S2.9 Residence times (B5)

140 default rows, one residence time per node and context, plus 107 rows each for the two residence variants. Where mass then goes is the routing table's business, so continuous residence models need no exit-specific share. The discrete release-schedule rows are an exception: their parameters are probabilities of departure at specified delays, not destination shares. Ten nodes

Table S6: The default parameterisation and the fourteen declared variants

ID	Family	Overrides	What it changes
V0	base	—	Nothing.
VB1	growth path	B2	Post-2020 growth rates at their declared low bound.
VB2	growth path	B2	Post-2020 growth rates at their declared high bound.
VF1	fluorinated gases	B2	Accelerated substitution; stationary HFO growth turns negative.
VF2	fluorinated gases	B2	Sustained HFO uptake at or above the declared high bound.
VM1	medical scope	A4	Seven subsectors to a cutoff of 2031, one to 2038.
VM2	medical scope	A4	Eleven subsectors renewed once, to a cutoff of 2050.
VR1	routing, containment bound	B4	Every split with a release-directed link: release-directed weights at their declared low, the others at their declared high; 186 rows.
VR2	routing, release bound	B4	The same splits with release-directed weights at their declared high, the others at their declared low; 186 rows.
VL1	residence, short	B5	Every service life in the product-use stock, except the release schedules, with time scale and bounds at 0.65 of the declared value; 107 rows.
VL2	residence, long	B5	The same rows at 1.35 of the declared value; 107 rows.
VE1	residence, technical and environmental, short	B5	Every first-order half-life of landfill, wastewater treatment, sludge, recycling and the environmental compartments at its declared low bound, truncation unchanged; 25 rows.
VE2	residence, technical and environmental, long	B5	The same rows at their declared high bound; 25 rows.
VS1	servicing, half	B1, B3	The servicing term of the fluorinated-gas anchor at 0.5 of the A541 estimate, first fill unchanged, the anchor and its subsector split rescaled; 13 rows.
VS2	servicing, one and a half	B1, B3	The same at 1.5; 13 rows.

Source: `A5_Variants.csv` and the variant rows of `A4_Scenarios.csv`. The override column decides which comparisons are legitimate: the six variants overriding B2 or B1 and B3 change the market realisation a run sees and do not preserve a matched market; the eight overriding A4, B4 or B5 can, because the market draw reads none of those tables (Section S9). The medical variants are keyed to the partial restriction only, so running them against the other two scenarios would reproduce the default exactly; those four sets are not run. A declared routing bound of zero is entered as 10^{-6} so that the reachable state set and the archive layout are identical under every variant. The service-life variants leave the residence times of the technical and environmental compartments as declared; VE1 and VE2 place their half-lives at the declared bounds.

carry rows (Table S7); the default timing models are 74 triangular, 25 first-order, 22 Weibull, 11 normal, six release schedules and two immediate. These counts are asserted at load, so an undeclared edit fails the run.

Residence at the product-use stock is the service life and is drawn per pack, representing spread within a population rather than uncertainty about a parameter. When a row gives lower and upper bounds, the engine truncates individual residence-time draws to those bounds. It does not treat the bounds as uncertain alternatives for the distribution parameters. This implemented truncation applies to continuous residence models in B5.

Bounds that were emptied. An input audit after an earlier campaign found twelve Weibull rows at the product-use stock — heating, ventilation, air-conditioning and refrigeration (HVACR) equipment (S07.1, S07.2), two medical device subsectors (S08.7, S08.9), electronics (S10.1, S10.2) and energy (S11.2) — whose bounds carried the spread of their source means or scales rather than limits on the residence of an individual pack. Because the engine truncates every draw to its bounds, those rows confined every service life to the window: for S08.7 and S08.9 to 11.9–12.0 years, so that every device retired at twelve years. The bounds of the twelve rows were emptied on 2026-09-04 (assumptions A549, A550), shape and scale unchanged; the parameter uncertainty the windows had encoded is carried by the residence variants VL1 and VL2. The notes of the 25 first-order rows describe their bounds as uncertainty in the half-life, whereas the engine applies them as truncation limits of the individual residence time, as stated above. Both readings are kept apart since 2026-09-15: the default rows keep the truncation, and the parameter uncertainty the notes declare is carried by the variants VE1 and VE2, which place the half-life at the row's low or high bound with the truncation unchanged (assumption A555, Table S6). Under VE2 the half-life coincides with the upper truncation limit, so the sampled residence is shorter than an untruncated exponential of that half-life would give; the variant is a bracket within the declared window. Finite truncation also means that these implemented distributions are not memoryless.

Table S7: Residence-time rows by node

Node	Name	Rows
PC2	Product use	113
TC2	Landfill	4
EN1	Soil	4
EN4	Air	4
PC4	Recycling	3
TC4	Wastewater treatment	3
TC6	Sewage sludge	3
EN2	Groundwater	2
EN3	Surface water	2
EN6	Sediment	2
	Total	140

Source: `B5_Service_Life.csv`, default rows. The 113 rows at the product-use stock are the service lives; the remaining 27 are residences in the technical and environmental compartments. The other thirteen nodes carry no row: four absorbing sinks and nine that pass mass on within the year. No matching row means a residence time of zero, not a missing value.

S3 Provenance and source-support grading

S3.1 Evidence identification and provenance chain

Evidence was assembled through an expert-led, parameter-by-parameter search rather than a systematic review. The Annex XV restriction dossier and its annexes supplied the common

2020 EEA frame (European Chemicals Agency, 2023a, 2023b, 2023d). Sources cited there were traced where a more direct value was needed, and targeted searches were used for service lives, waste treatment and environmental fate. Adopted legislation informed the annual cutoff assumptions (European Commission, 2024, 2025; European Parliament and Council, 2024, 2025). The annualised implementation and the hierarchy between model rows are separate modelling decisions. The evidence search closed on 4 September 2026; the input tables were last changed on 8 September 2026, when four scenario rows were corrected from sources already collected (assumptions A551–A554, Section S8.3).

Revision of 2026-09-04. A comparison of an earlier campaign with the restriction dossier (Section S6.2) and an audit of the input tables led to three revisions, each carried by new assumption rows (A541–A550) in the assumption log (Online Resource 2, C2). First, the anchor of the applications of fluorinated gases had been a first-fill quantity, so the gas that escapes during use and is replaced by servicing had no carrier; the servicing subsector S07.11 was added with the dossier's 2018 stock-based release as its central value (European Chemicals Agency, 2023b), and the sector anchor rose from 30.7 to 69.5 kt/a at the mode. Second, the medical-device anchor for fluorinated gases had rested on an assumption identifier that had never been logged; it was re-derived from the dossier's inhaler-propellant figure (0.16/3.08/6.0 kt/a) and the dossier's 30 kt/a placeholder for process gases was documented as excluded. Third, twelve service-life rows carried the spread of their source means in their truncation columns; those bounds were emptied (Section S2.9). The bounding variants VR1, VR2, VL1 and VL2 were added at the same time. The earlier campaign is superseded and no number from it is reported.

Each parameter row follows one traceable chain:

parameter row → `assump_id` → assumption and rationale → reference → source.

The complete chain is supplied in Online Resource 2. Tables C1 and C2 are not read by the simulation, but are hashed with the campaign record. In 104 default residence-time records, the historical `audit_note` refers to assumption A478, the prior on the service-life means; that row was reconstructed in C2 on 2026-09-04 from the citing notes and is the basis of the residence variants. The active `assump_id` links resolve throughout; the historical text should not be used to reconstruct the implemented parameterisation.

Published low, central and high estimates were retained when available. A single published estimate was used either as a point value or with a documented range from a comparable quantity. Conflicting sources, changes of denominator, geographic proxies and adaptations between reference years are recorded in the assumption log. Missing values were not filled automatically; they were declared as assumptions with a rationale.

S3.2 Source-support grades

The grades describe traceability and directness, not parameter accuracy:

High. A source reports the value directly, or the value follows structurally from a directly supported quantity.

Intermediate. A source supports the mechanism or an analogous magnitude, but adaptation is required.

Low. The value is expert judgement or an unvalidated proxy.

Grades were assigned during parameterisation and later audited against the source–assumption–parameter chain. This was a traceability audit, not independent duplicate grading. Allocation table B3 also carries a row-specific source-support grade for each split; this finer grade is used for the disaggregated mass summary below.

S3.3 Mass-weighted source support

Row counts describe the number of parameter decisions; mass-weighted shares describe how much of the 2020 market inflow depends on each grade.

Table S8: Source support by parameter layer

Layer	Basis	High	Interm.	Low
Market-inflow anchors	mass (%)	68.9	1.3	29.8
Market-inflow anchors	rows (count)	24	1	2
Subsector and substance class	mass (%)	14.8	39.8	45.4
Subsector and substance class	rows (count)	19	51	180
Growth rates	rows (%)	22.7	36.4	40.9
Routing weights	rows (%)	9.3	43.8	46.9
Residence times	rows (%)	6.4	41.4	52.1
Scenario rows	rows (%)	12.7	82.7	4.7

Computed from the default input tables as of 2026-09-16, after the re-sourcing described below; before it, routing rows graded 7.0, 32.6 and 60.5 percent and residence rows 6.4, 25.0 and 68.6 percent high, intermediate and low. Mass shares use the 2020 anchor modes (234181 t in total, of which 38806 t is the servicing subsector) and normalised allocation weights. The low-graded anchor mass is the combined anchor of the applications of fluorinated gases (A541), whose servicing part is the dossier's stock-based release estimate; the intermediate anchor is the medical-device propellant (A545). Anchor grades are resolved through C2; allocation grades use the row-specific B3 field. Other layers are counted by active default rows. Percentages may not total 100 because of rounding.

Re-sourcing after the pre-submission review. On 2026-09-16, fifteen assumptions behind the routing and residence rows that the variants identified as carrying the results were checked against published measurements and official statistics that had not been used at parameterisation: the soil, groundwater, surface-water and sediment residence times (A353, A381), the landfill half-lives of foam blowing agents and fluoropolymers (A464), the service lives of building foams, textiles and implants (A162, A382, A169, A170), the sewage-sludge disposal split (A379), the incineration fate (A431), the sludge partitioning of precursors (A377), textile recycling (A104, A108), mobile air-conditioning leakage (A413) and inhaler propellant release (A496); the wastewater, sludge and recycling residence times (A389) were checked and kept low. No value was changed. Each assumption's rationale in C2 now states the value, the corroborating source and the number it reports, replacing the former rationale, and the grade was reassigned by the criteria above: to high where the value follows structurally from a reported quantity (the sludge split from SWD(2023) 157), to intermediate where sources report the magnitude or the mechanism but the value is an adaptation (thirteen assumptions), and left low where a component of the assumption has no source (A389). Twenty-six references were added to C1 (R106–R131). On 2026-09-22 the citation texts of nine of these references were corrected against Crossref and PubMed, the thirteen references cited only in the main text were added to C1 (R132–R144), and the audit notes of A169 and A170 record that the pacemaker medians they quote refer to the devices that reached the replacement interval; no value or grade changed. The tables the engine reads are untouched, so the engine digest of both campaigns stands; the provenance digest changed to `f409da31 265a6dd5 5d665e3d b8b4241c 5d44ca85 b0d6b3cd 96c2141b 9301f09a` on 2026-09-16 and to `01b74519 3e14d184 2b7b75e6 17335072 7162cb55 8573b383 f4126b95 3fa389af` on 2026-09-22 (Section S8.3).

S3.4 Calculation

The 27 anchor modes were weighted by mass and grouped by their C2 grade. Within each sector and substance group, allocation modes were normalised as in the engine, multiplied by the

anchor mass and grouped by the B3 row grade. Growth, routing, residence and scenario rows were counted in the default parameterisation; variant and scenario-header rows were excluded. Superseded C2 records were excluded, and the mode was used throughout. The calculation uses input tables only and can be repeated independently of the simulation outputs.

S4 Simulation algorithm and equations

This section describes the implementation in `loadInput.m` (reading, variant overrides, context resolution, scenario cutoffs) and `simulate.m` with `private/drawResidence.m` (the draws and the annual cycle). Table identifiers refer to the archived inputs. A context is a subsector and substance-class combination. The implementation resolves one set of routing weights and at most one residence row per node and context; the residence row is not selected by destination or flow class.

The equations below expand the same notation used in the manuscript. Market inflow is developed in Section S4.2; routing and residence draws in Sections S4.3 and S4.4; and release, stock index, committed share and derogation share in Section S4.5. Section S4.9 describes the random-number construction used for paired scenario comparisons.

S4.1 Packs and state variables

A cohort consists of the mass entering in one year, subsector and substance class. Every non-empty cohort is divided into $N = 200$ packs of equal mass. Each pack keeps its mass, class, subsector, year of market entry and identifier throughout a run. Its node, next destination, scheduled departure and random-number counter change during movement. Packs do not interact and are never split. Cohort tracking is the implementation chosen here; cohort-resolved compartment models can also retain the year of market entry.

Routing weights determine expected mass fractions over repeated draws. For a link with probability p, the number of packs taking it from N independent arrivals has variance $Np(1-p)$. A rare route can therefore receive no pack in an individual cohort even when its expected mass is positive. Its contribution is therefore assessed in aggregates and through the pack-resolution analysis of Section S5; positive routing probability does not guarantee adequate resolution of each detailed cell.

S4.2 Market inflow and the annual cycle

For subsector s, class c and year t, the main-text inflow equation is $m_{sc}(t) = A_{k(s,c)} w_{sc} G_{sc}(t)$, where $k(s,c)$ identifies the sector-and-group anchor. The growth factor is

$$G_{sc}(t) = \begin{cases} \prod_{u=2021}^{t} (1 + g_{sc}(u)), & t > 2020, \\ 1, & t = 2020, \\ \prod_{u=t+1}^{2020} (1 + g_{sc}(u))^{-1}, & t < 2020. \end{cases} \tag{S1}$$

Here $g_{sc}(u)$ is the annual growth rate as a proportion; values stored in B2 are divided by 100. The whole inflow is set to zero before the class-specific start year in A2. Overlapping growth rows are resolved by precedence rank, with the higher rank taking precedence.

For a scenario with applicable cutoff rows $\mathcal{Q}_{sc}$, the actual inflow is

$$\widetilde{m}_{sc}(t) = m_{sc}(t) \sum_{h \in \mathcal{Q}_{sc}} \varphi_h \, \mathbf{1}(t \le T_h). \tag{S2}$$

The index h denotes a cutoff row, φ_h its share and T_h its last inflow year. These shares sum to one within the applicable set. The run index r is reserved for Monte Carlo runs and is not used

for cutoff rows. Scenario labels and nominal derogation lengths do not replace the explicit shares and years read by the engine.

At the start of each run, the model resolves variant rows and draws the market parameters. All stocks initially equal zero. Each year it creates the new cohorts, processes scheduled departures and repeats zero-delay movements until none remain. It then records annual flows and year-end stocks. Four absorbing accounting nodes retain mass assigned to export, mineralisation, marine outflow and atmospheric outflow. Including these quantities gives the checked identity

$$\sum_{n=1}^{23} S_n(t) = \sum_{u=1950}^{t} M(u), \tag{S3}$$

where $S_n(t)$ is the node inventory and $M(u)$ total actual market inflow. The runtime tolerance is $10^{-6}\max(\sum_u M(u), 1)$ tonnes. Inventories at absorbing nodes are cumulative accounting quantities, not stocks of PFAS physically remaining in the EEA environment.

S4.3 Drawing on arrival

A pack draws its destination and residence time on arrival and retains both until departure. Returning through recycling counts as a new arrival and initiates a new service life. Drawing a complete service life repeatedly while a pack remains in place would change the specified distribution. A correctly implemented age-dependent hazard could give an equivalent alternative, but is not used here.

For a continuous residence distribution with cumulative distribution function F, survival to age a is $1 - F(a)$. The continuous hazard, before annual rounding, is

$$h(a) = -\frac{\mathrm{d}}{\mathrm{d}a}\log[1 - F(a)]. \tag{S4}$$

For an untruncated Weibull distribution with shape β and scale λ, it equals $(\beta/\lambda)(a/\lambda)^{\beta-1}$ and increases with age when $\beta > 1$. An untruncated exponential distribution has a constant hazard. Truncation changes the survival function and generally removes the memoryless property. Thus the bounded first-order rows in B5 must not be described as memoryless residence distributions. Their half-life parameter specifies the underlying exponential shape, not necessarily the median of the truncated and rounded residence times.

S4.4 Routing and residence sampling

Within the resolved routing set, the probability of destination j from node i is

$$p_{ij} = \frac{b_{ij}}{\sum_l b_{il}}, \tag{S5}$$

where b_{ij} is the mode of the declared routing weight. The relevant context is suppressed in this notation. The routing modes are not sampled. A second uniform draw supplies the residence time, independently of the chosen destination conditional on the node and context.

For a continuous residence model with cumulative distribution function F and valid declared bounds $L < U$, the draw is

$$\tau = F^{-1}\{F(L) + v[F(U) - F(L)]\}, \qquad d = \max(\lceil\tau\rceil, 0), \qquad v \sim \mathcal{U}(0,1). \tag{S6}$$

Departure occurs at arrival year plus d. If the bounds are absent or $U \leq L$, the untruncated draw $F^{-1}(v)$ is used. The implementation caps the remapped probability below one to avoid an infinite inverse. Negative normal draws, if allowed by the specified bounds, become zero after the non-negative rounding rule.

Table S9: Residence models in the default B5 input

Model	Rows	Sampling rule
Triangular	74	Minimum, mode and maximum; inverse-CDF draw.
First order	25	Exponential rate $\log(2)/H$, where H is the declared half-life parameter.
Weibull	22	Shape and scale; inverse-CDF draw.
Normal	11	Mean and standard deviation; inverse-CDF draw.
Release schedule	6	Discrete random departure: delay zero with probability 0.9; each delay 1–4 with probability 0.025.
Immediate	2	Delay zero.

Continuous models use declared truncation bounds and upward rounding. A missing residence row also gives delay zero, but does not cause a pack to leave an absorbing node. The release schedule is sampled per pack; exact cohort fractions are not imposed. Service-life variants override 107 product-use rows each and leave the six schedules unchanged.

Under the default servicing row B5-0187, a triangular draw between zero and one becomes a one-year delay almost surely. Thus it does not give same-year release. The short service-life variant also gives one year; the long variant extends the support to 1.35 years and permits two-year delays. The same annual-rounding rule applies to other sub-year rows, including aerosol and inhaler-propellant uses. This is an implementation limitation when assigning release to individual calendar years.

S4.5 Release and reported outcome definitions

Let $q = 0, 1, 2$ denote no further restriction, the comprehensive restriction and the partial restriction, respectively, and let r denote a Monte Carlo run. Let $\mathcal{E}_{\text{rel}}$ be the set of links whose destination is an environmental node and whose origin is not, derived at load time as in Section S1.4, and let $f_{ij,q,r}(t)$ be the mass moving along link $i \to j$ in year t. Annual release is

$$R_{q,r}(t) = \sum_{(i \to j) \in \mathcal{E}_{\text{rel}}} f_{ij,q,r}(t). \tag{S7}$$

Because every pack carries its year of market entry y unchanged, the same sum restricted to packs with $y \leq 2026$ gives $R^{\text{pre}}_{q,r}(t)$ and the remainder gives $R^{\text{post}}_{q,r}(t)$. Thus $R_{q,r}(t) = R^{\text{pre}}_{q,r}(t) + R^{\text{post}}_{q,r}(t)$. The two parts are disjoint and exhaustive. Transfers inside the environmental band, EN3 $\to$ EN5 and EN4 $\to$ EN7, are outside $\mathcal{E}_{\text{rel}}$ and are therefore not counted a second time.

Cumulative release from entry into force to year T is

$$\mathcal{R}_{q,r}(T) = \sum_{t=2027}^{T} R_{q,r}(t). \tag{S8}$$

For environmental compartment e, the stock index used for Q3 is

$$I_{e,q,r}(t) = \frac{S_{e,q,r}(t)}{S_{e,q,r}(2027)}. \tag{S9}$$

It is calculated separately for soil, groundwater, surface water, air and sediment. The index is defined only for a positive 2027 compartment stock, a condition met by every reported run and compartment.

The reported form for Q3 is the aggregate environmental-stock index. Let $\mathcal{A} = \{\text{EN1}, \text{EN2}, \text{EN3}, \text{EN4}, \text{EN6}\}$ be the five resident environmental compartments. Then

$$A_{q,r}(t) = \frac{\sum_{e \in \mathcal{A}} S_{e,q,r}(t)}{\sum_{e \in \mathcal{A}} S_{e,q,r}(2027)}. \tag{S10}$$

The stocks are summed within the run before the ratio is formed, so $A_{q,r}$ is mass-weighted and is not the mean of the five compartment indices. The accounting sinks for marine and atmospheric outflow, EN5 and EN7, are excluded because they hold mass that has left the EEA.

The committed share used for Q2 is

$$C_{q,r} = \frac{\sum_{t=2027}^{2070} R^{\mathrm{pre}}_{q,r}(t)}{\mathcal{R}_{q,r}(2070)}, \qquad q \in \{0, 1, 2\}. \tag{S11}$$

A pre-2027 pack that has not crossed the release boundary by 2070 is outside the numerator. The numerator is therefore a lower bound on release ultimately arising from pre-2027 mass; $C_{q,r}$ itself is explicitly the share within the 2027–2070 reporting horizon.

Avoided release under restriction scenario q is $\Delta_{q,r}(T) = \mathcal{R}_{0,r}(T) - \mathcal{R}_{q,r}(T)$. The derogation share, the second reported form for Q2, is

$$\rho_r(T) = \frac{\mathcal{R}_{2,r}(T) - \mathcal{R}_{1,r}(T)}{\mathcal{R}_{0,r}(T) - \mathcal{R}_{1,r}(T)}. \tag{S12}$$

The numerator is the release left in place by derogations; the denominator is the release avoided by the comprehensive restriction. Both are formed within run r before the distribution is summarised. The share is defined when the denominator is positive, as it was in every reported run.

S4.6 Resolving a context to rows

`loadInput.m` resolves routing and residence rows using the sector, subsector, substance group and substance class. The most specific matching set is used; it is not merged with less specific defaults. Routing weights are normalised only within that set. For example, firefighting training and testing has a specific route to release; combining it with the sector default would introduce a waste route not specified for that subsector. Scenario cutoffs are resolved separately by subsector and substance group in `loadInput.m`; matching tranches within that set are summed.

S4.7 Reachability and numerical safeguards

The allocation table contains 250 rows, of which 239 have positive mode weights and define active contexts. The compiled model checks that every reachable non-absorbing state has an outgoing positive-weight route and that residence resolution is unambiguous. It also checks the zero-delay network. A runtime guard stops a year exceeding 25 movement iterations. This guard detects unexpected cascades; it is not a physical limit on the number of transfers a material can undergo.

S4.8 What is sampled and what is assumed

Market anchors, raw allocation weights and growth-rate rows are sampled once per run. Normalisation makes allocation shares dependent and changes their marginal distributions. A shared parameter row creates dependence across the years or contexts it serves. Route and residence draws occur once per pack arrival. Routing probabilities and residence parameters are fixed within a parameterisation; variants change selected values between parameterisations.

There is no chemical transformation, capacity constraint, interaction between packs or endogenous response of servicing requirements to equipment stock. The servicing inflow is a replacement proxy, not a direct stock-depletion calculation. The class and market-entry year assigned to that proxy must not be interpreted as chemical or physical tracking of the gas that escaped. Excluding transformation preserves compound mass labels but does not establish a general upper bound on real environmental compartment stocks. Input revisions informed by earlier model–dossier comparisons are documented in Section S3.1.

S4.9 Random numbers and reproduction

The master seed and run index determine a run key. Market draws use dedicated counters. Pack keys additionally depend on cohort and pack index; within a pack, counters advance for destination and residence draws. They do not include scenario identifiers. Matching scenarios within one parameterisation therefore share random numbers for matching cohorts. Changing pack count changes the pack-key construction, so pack paths are not nested across the resolution ladder even though market draws remain matched. Exact reproduction requires the code, input bytes, settings and compatible numerical environment, not the seed alone.

S5 Numerical resolution and Monte Carlo precision

The campaign uses 200 packs per cohort and 1000 Monte Carlo runs per set. The criterion file records these settings and the assessment procedure before the production campaigns. Exploratory calculations preceded it, so the criterion hash identifies the version the gate used but does not establish independence from those calculations.

S5.1 Primary outcomes

The resolution criterion names nine outcomes under the default parameterisation:

1. Counterfactual cumulative release, 2027–2050.
2. Counterfactual cumulative release, 2027–2070.
3. Counterfactual committed release, 2027–2070.
4. Counterfactual environmental compartment stock, 2070.
5. Counterfactual product-use stock, 2070.
6. Counterfactual mean annual release, 2046–2050.
7. Counterfactual mean annual release, 2066–2070.
8. Avoided release under the comprehensive restriction, 2027–2070.
9. Avoided release under the partial restriction, 2027–2070.

These aggregate outcomes do not qualify the numerical precision of every individual subsector, substance class or annual cell. Five-year means are used for late-horizon release to reduce finite-pack variation.

S5.2 Pack resolution and the gate

A pilot with 100 packs and 100 runs supplies the p05–p95 band width used to scale each contrast. Forty run keys are replayed at 5, 10, 25, 50, 100, 200, 400, 600, 800 and 1000 packs per cohort. Market draws are identical for a given key across pack counts; pack trajectories are not nested. Figure S1 shows every outcome on this ladder relative to its value at 1000 packs. For outcome Y, run r and pack count N, the convergence tool calculates

$$D_r(N) = Y_r(800) - Y_r(N). \tag{S13}$$

Positive contrasts therefore mean a larger outcome at the reference resolution. This sign convention also applies to avoided-release outcomes, which are first formed by paired scenario subtraction.

The gate summarises the contrasts by their median $\widehat{D}$ and the implemented approximation

$$\widehat{D} \;\pm\; 1.96\,\frac{1.2533\,s_D}{\sqrt{40}}, \tag{S14}$$

where s_D is their sample standard deviation. The factor 1.2533 is the asymptotic normal-distribution approximation for a median's standard error relative to that of a mean, not a distribution-free formula. The software also writes an order-statistic interval for the median and sign-test diagnostics; these permit comparison with the approximation but do not replace the declared gate statistic.

The production count passes when both interval limits fall within ±2% of the pilot band width for every primary outcome. This is a practical agreement criterion relative to a finite 800-pack reference, not a bound on error relative to an exact solution, and the pilot denominator is itself estimated. The intervals are pointwise, so requiring all primary outcomes to pass is a joint decision rule rather than a correction for multiplicity.

Recorded status. The gate was first evaluated before the preceding campaign, on its input state (input digest `02b2f88c1ec5`), and passed; the decision to run at 200 packs rests on that evaluation. That state differs from the final one only in four scenario rows: the derogation scope of fluorinated gases after 2026 (assumptions A551–A553) and the scope of the comprehensive restriction (A554). Because the two avoided-release outcomes involve the restriction scenarios, whose inflow after 2026 changed, the gate was repeated on the final input state (input digest `891f3b871362`), with the same criterion file, seed and ladder; the repeated gate record is the one deposited with the archive. It recorded a pass at 200 packs, seed 42, against an 800-pack reference: all nine declared outcomes passed, and the largest absolute endpoint of a 95% contrast interval was 1.42% of the corresponding pilot Monte Carlo band, below the 2% criterion, for release from pre-2027 mass under the counterfactual. For the two avoided-release outcomes the interval endpoints were within 0.33% (comprehensive restriction) and 0.46% (partial restriction) of the band. The counterfactual outcomes are identical in both evaluations, since no counterfactual flow or stock differs between the two input states.

S5.3 Precision across runs

After the final campaign, 2000 bootstrap resamples of the 1000-run default sets estimate standard errors of the median and the 5th and 95th percentiles. For scenario differences, the bootstrap resamples paired run differences. The declared thresholds are 2% of the Monte Carlo band width for the median and 3% for each band limit. These assess Monte Carlo summary precision, not uncertainty about model structure. Quantiles use MATLAB's percentile definition, linear interpolation between order statistics placed at $(i-0.5)/n$ (Hyndman–Fan type 5). The bootstrap seed is 1. The values reported here are those of the deposited convergence tool run on the final archive.

All nine declared outcomes passed. The largest bootstrap standard error was 1.63% of the Monte Carlo band for a median and 2.91% for a band limit, both for mean annual counterfactual release in 2046–2050, below the respective 2% and 3% criteria. For the two avoided-release outcomes the largest values were 1.30% and 2.58%. Outside the declared set, two band limits of the partial restriction's own levels exceed the 3% criterion marginally: the 95th percentile of cumulative release 2027–2050 (3.04%) and of mean annual release 2066–2070 (3.16%).

S5.4 Interpretation and scope

Pack-discretisation noise contributes to variation across runs. Its effect on a percentile-band width need not be monotonic for every outcome. The ladder separately assesses changes in outcome level

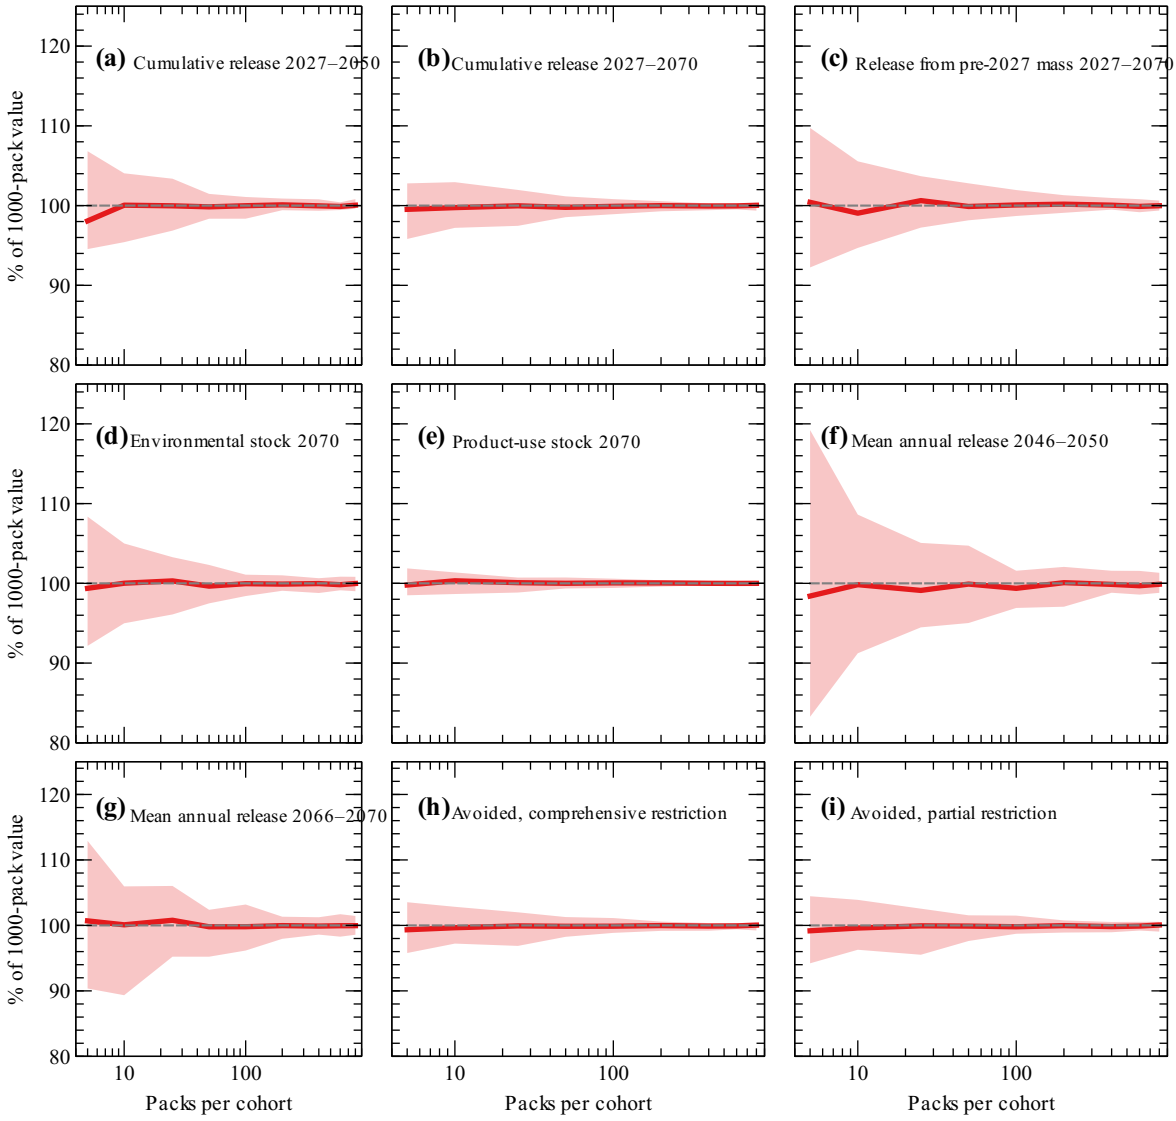


Figure S1: Pack-count ladder for the nine primary outcomes. Each panel shows one outcome of Section S5.1, in the order listed there, as a percentage of the same run key's value at 1000 packs per cohort: the line is the median across the 40 replayed keys, the band their 5th to 95th percentiles, at 5, 10, 25, 50, 100, 200, 400, 600 and 800 packs. The dashed line marks agreement with the 1000-pack value. This is not the gate statistic, which tests the median contrast against the 800-pack rung in units of the pilot band. Values: Online Resource 3, sheet `Convergence`.

with pack count. Neither assessment covers uncertainty in routing weights, residence parameters or omitted processes.

The pack ladder covers the default parameterisation and paired scenario differences only. Retaining the same network in a variant does not guarantee similar pathway probabilities or numerical precision. Consequently, the default gate is not evidence that every variant or detailed result has converged. Run-count precision is assessed for each completed set separately.

S6 Verification and external consistency

Verification checks whether the implementation follows the declared model. External comparisons assess consistency with other estimates. Neither establishes independent empirical validation; Section S6.2 states where the comparison shares its evidence base with the model.

S6.1 Verification

The final campaign contains 29 sets, with seed 42, 200 packs per cohort and 1000 runs per set; the extension campaign of 2026-09-15 adds 12 sets at the same settings (Section S8.2). On the 12 extension sets the five archive checks were run on 2026-09-16 with the same code and passed: no deviation from the archive contract, a largest relative mass-balance residual of 2.34×10^{-8}, a largest aggregation deviation above one tonne of 5.34×10^{-8}, identical pre-2027 histories across the scenarios of each variant, and a largest cutoff deviation of 4.35×10^{-8} (`06_Analysis/phase1/CHECKS_extension.csv`); the largest annual mass-balance residual reported by the engine was 5.2×10^{-8} t. On this archive the engine's annual mass-balance assertion held in every set, with a largest absolute residual of 6.0×10^{-8} t (campaign run log). Pre-2027 histories, all flows and stocks over 1950–2026, were exactly equal across the scenarios of every parameterisation, and between each medical variant and the default counterfactual (check P4). On the default sets, the result-table verification of the deposited code (`verify.m`) passed: avoided release and comprehensive release add to counterfactual release exactly within each run, and a second aggregation path written independently of the result modules reproduced the 28 headline numbers of Tables S12 to S14 to a largest relative deviation of 6.1×10^{-9}.

The five archive checks P1–P5 (Table S10) were run on all 29 sets of the final archive on 2026-09-15 and passed: no deviation from the archive contract, a largest relative mass-balance residual of 2.37×10^{-8}, a largest aggregation deviation above one tonne of 5.45×10^{-8} and a largest cutoff deviation of 8.35×10^{-8}. A first run of P1 on the same day had recorded 24 label-layout deviations, all in the four files of the routing variants VR1 and VR2: a variant that overrides B4 appended its rows after the rows they replace, so the 49 links of those six sets were stored in a different order than under V0, the same links carrying the same numbers. The six sets were brought to the V0 order on 2026-09-15 by `06_Analysis/phase1/repair_P1_edge_order.m`, which permutes every array indexed by link together with its labels and changes nothing numerical; the tidy table behind Table S17 is byte-identical before and after. The engine now places variant rows in the position of the rows they replace, and P1 was repeated with the result above.

Reproduction check. The check reruns the default counterfactual at five runs, on the seed and pack count the archive itself records, and compares the per-run arrays with runs 1–5 of the archived 1000-run set. Because every draw is keyed on the seed and the run index, those five runs must agree bit for bit however many runs the set contains; the summary and cohort reductions are formed over a different run count and are therefore excluded from the comparison. On the final campaign the check passed: `stock`, `flow`, `flowBySector`, `flowBySubgroup` and `flowCumByContext`, and the diagnostic arrays `massBalance`, `maxHops` and `packsCreated`, were bitwise identical, so the archive is regenerable from the input bytes, the seed and the code alone. The rerun was made on 2026-09-14 in the software environment of Table S19 with the

Table S10: Campaign verification protocol

Check	Quantity checked	Criterion
P1	Archive variables, shapes, indices and labels	No deviations
P2	Annual mass balance and release–environment identities	Relative residual $< 10^{-6}$
P3	Agreement across pack, cohort and run aggregates above one tonne	Relative residual $< 10^{-5}$
P4	Pre-2027 histories across scenarios within a parameterisation	Exact equality
P5	Market inflow against declared cutoff years and shares	Relative residual $< 10^{-6}$

Definitions follow the check suite of the deposited code. The values measured are those quoted in the text, all on the final campaign. Zero-inflow conditions are tested separately for exact zero.

reproduction check of the deposited code. Reproduction is assessed on the arrays themselves, alongside the code version, input hashes and software environment.

Variants and software tests. The medical variants change only the partial restriction. Their other scenario combinations are omitted because they duplicate the default. Routing bounds of zero are represented by 10^{-6} in the bounding variants to retain the compiled state layout. Engine tests cover input resolution, reachability, random draws, residence sampling, variant overrides, inflow, mass balance and the convergence gate. These tests support implementation consistency; they do not verify the empirical meaning of a parameter or the legal interpretation of a scenario.

S6.2 External consistency

The comparison uses the Annex XV report's Tables 1 and 3 (printed pp. 40–41 and 54–55) for 2020 market quantities, release estimates and selected technical stocks (European Chemicals Agency, 2023d). Model values are from the default counterfactual of the final campaign; firefighting foams are excluded on both sides. Modelled 2020 market inflow is 236 kt [201–268], against the report's 225 kt (range 138–312 kt); by substance group it is 33.1, 73.7 and 128.5 kt against 32.5, 64.0 and 128.5 kt. Modelled 2020 release, which counts all vintages and includes the waste stage, is 80.4 kt [63.9–95.0]. It lies within the range of the report's total that includes emissions from technical stocks (74.1 kt; 56.0–92.2 kt) and is 2.2 times its total from new-to-market tonnage alone (36.6 kt). By group, release of PFAAs and their precursors (8.85 kt) and of fluoropolymers and perfluoropolyethers (12.1 kt) lies within the report's ranges, whereas release of fluorinated gases, 59.6 kt, exceeds its range of 42.9–49.9 kt. The modelled stock of fluorinated gases in products in use at the end of 2020, 653 kt, is 1.3 times the report's technical stock volume of 493 kt, which refers to 2018; for transport the two agree (164 against 160 kt).

Fluorinated-gas sector, Annex E. Annex E, Table E.86 (printed p. 254) reports mean emissions of the fluorinated-gas sector, excluding mobile air conditioning and transport refrigeration, under a baseline and under restriction option RO1, a full prohibition (European Chemicals Agency, 2023c). For 2025–2055 it gives 1942 kt under the baseline and 92.6 kt under RO1, a reduction of 95%; for 2025–2070, 3407 kt and the same 92.6 kt, a reduction of 97%. The model counterpart is release from applications of fluorinated gases, which likewise excludes mobile refrigeration (assigned to transport) and includes servicing of installed equipment. Under the default parameterisation of the final campaign it is 1180 kt [727–1970] under no further restriction, 719 kt [494–1050] under the partial restriction and 403 kt [292–516] under the comprehensive restriction for 2025–2055, a paired reduction of 66.0% [52.1–78.6] under the comprehensive restriction; for 2025–2070 it is 1770 kt [1040–3150], 758 kt [525–1100] and 428 kt [309–546], a reduction of 76.1% [63.7–85.9]. Two differences are visible. The dossier's RO1 total is identical for both windows, whereas modelled release under the comprehensive restriction keeps accruing after

2055 because pre-2027 mass in installed equipment continues to reach the release boundary. The dossier's baseline was also constructed before Regulation (EU) 2024/573 was adopted, whereas the counterfactual here includes its phase-down, so the dossier's baseline and the counterfactual here are different scenarios. Dossier values are source means; model values are medians with the Monte Carlo band, formed within the run.

Dependence on the parameter sources. Market-inflow anchors use the dossier, so agreement principally checks transcription and allocation. Releases and product-use stocks are derived outputs, but share the dossier's market evidence. An earlier comparison also informed the servicing and medical-device input revisions. Accordingly, the comparison is partly a source-reproduction check and partly a consistency assessment after model development; it is not a held-out validation exercise.

Servicing supply and medical devices. The servicing proxy uses the dossier's stock-based release estimate of 38806 t/a (Annex B, Section B.9.9) (European Chemicals Agency, 2023b). Comparing total fluorinated-gas release with that estimate reuses input information. Annex A, Table A.97 reports EEA bulk supply of 77566 t/a in 2019 (European Chemicals Agency, 2023a); it offers a separate supply-side comparison, subject to reference-year and coverage differences, rather than a like-for-like measurement of the model's 2020 inflow. For medical devices, the model anchor is the dossier's inhaler-propellant figure (0.16/3.08/6.0 kt/a, assumption A545, Section S3.1); the dossier's sector total also contains a 30 kt/a placeholder for process gases, which is outside the quantified anchor. The comparison therefore uses the dossier total without that placeholder, 4.7–21.5 kt/a; the excluded quantity is recorded beside the adjusted basis in the assumption log.

Harmonising the comparison. The report distinguishes releases estimated from annual market quantities from estimates that use technical stocks. Model release includes multiple life-cycle stages and cohorts entering since 1950, so the comparison above pairs it with the report's total that includes emissions from technical stocks. Firefighting foams are outside the two cited dossier tables and are excluded from the model totals used for that comparison. Reported ranges are source ranges, not Monte Carlo bands. No sufficiently matched measured time series was identified for independent evaluation of the complete model.

S7 Result tables and figures

This section gives the numerical record behind the findings in the main text. Unless a table states otherwise, entries are medians with the 5th and 95th percentiles in brackets and use the default parameterisation V0 of the final campaign. Values are rounded to three significant figures. The variant results in Section S7.5 use all 29 sets.

Three aggregation rules apply to every table here and are not repeated in the legends. Totals across years, sectors or classes are summed within a run before the median and band are taken, because adding quantiles does not give the quantile of a total. Every ratio, index, difference and share is likewise formed within the run and only then summarised, so it is not the quotient or difference of the tabulated medians. Compositions by sector, substance class or market-entry vintage instead use component mean masses divided by their summed mean mass, so that the parts add to the whole; these carry no Monte Carlo band, and a set of component medians need not sum to the total.

The tables are the record. Each figure of this section draws the values of the table beside it and adds no number of its own. The annual series from which the cumulative quantities are formed are in Online Resource 3.

S7.1 Market inflow and release under the three scenarios

Market inflow is a scenario input; release is a simulated outcome. Table S11 gives both as annual levels in selected years, together with the cumulative inflow each scenario admits and the years at which a given share of counterfactual release has been avoided. Table S12 gives the cumulative release and the avoided and retained quantities derived from it.

Table S11: Annual market inflow and annual release in selected years, cumulative market inflow, and the years at which a share of counterfactual release has been avoided

Quantity	Year or share	Counterfactual	Partial restriction	Comprehensive restriction
Annual market inflow (kt/a)	2026		268 [221–310]	
	2027	256 [208–302]	119 [101–139]	12.0 [7.29–17.3]
	2031	278 [222–333]	119 [98.4–149]	0.130 [0.0660–0.210]
	2036	310 [248–369]	86.0 [69.9–112]	0
	2040	347 [280–412]	8.73 [3.37–14.3]	0
	2070	564 [464–673]	16.1 [5.90–27.8]	0
Cumulative market inflow (kt)	2027–2070	18400 [15200–21800]	1610 [1240–2060]	13.2 [8.32–18.4]
Annual release (kt/a)	2027	72.0 [57.5–86.4]	71.8 [57.3–85.9]	70.9 [56.7–84.8]
	2040	83.5 [62.9–112]	28.8 [22.3–35.5]	17.4 [13.7–21.2]
	2050	96.3 [72.2–131]	23.9 [15.5–33.2]	8.82 [7.19–10.6]
	2070	127 [94.8–179]	20.9 [11.2–31.8]	5.16 [4.20–6.27]
Avoided annual release, counterfactual minus scenario (kt/a)	2027	—	0.159 [0.0852–0.370]	1.02 [0.632–1.63]
	2040	—	53.8 [37.1–82.2]	65.9 [47.7–94.1]
	2070	—	106 [76.4–154]	122 [89.6–174]
Year at which avoided release reaches	25 %	—	2042 [2040–2043]	2029 [2028–2031]
	50 %	—	2055 [2051–2061]	2037 [2033–2040]
	75 %	—	not reached	2055 [2048–2063]

Medians with the 5th and 95th percentiles in brackets. Market inflow before 2027 is identical in all three scenarios, so the 2026 row is given once; it is a median, whereas the historical series of Figure 2a in the main text is a series of means. A zero entry is an exact zero in every run. The avoided annual release rows are paired differences formed within each run; the single-year 2070 values, like the 2070 release rows, are more sensitive to pack discretisation than the five-year means of Table S13. The crossing years are formed within each run as the first year in which cumulative avoided release reaches the stated share of cumulative counterfactual release; under the partial restriction no run reaches 75 % by 2070. Final campaign, V0. The complete annual series of market inflow and release are in Online Resource 3, sheet `Flows_annual`, boundary classes `market_entry` and `release`, class `ALL`, vintage `all`.

S7.2 Release committed before 2027

For the decomposition of counterfactual release in the main text (Results, second question), release over 2027–2070 under no further restriction splits within each run into four parts: release from mass already placed on the market, 704 kt [585–817]; release from inflow still permitted by adopted law, 4.63 kt [3.03–6.38]; release avoided by the partial restriction, 2680 kt [1930–3820]; and release left in place by its derogations, 842 kt [496–1240]. The four parts add exactly within each run, as the result-table verification of Section S6.1 confirms; their medians do not add. Their means, 702, 4.67, 2770 and 851 kt, add to the mean counterfactual total of 4330 kt.

Table S12: Cumulative and avoided release under the default parameterisation

Quantity	Window	Counterfactual	Partial restriction	Comprehensive restriction
Cumulative release (kt)	2027–2050	1990 [1530–2650]	1110 [865–1430]	584 [476–684]
	2027–2070	4250 [3230–5730]	1540 [1160–2000]	709 [588–822]
Avoided release (kt)	2027–2050	—	876 [620–1220]	1400 [1010–2020]
	2027–2070	—	2680 [1930–3820]	3530 [2580–4990]
Avoided release (% of counterfactual)	2027–2050	—	44.3 [38.2–49.0]	70.8 [64.8–77.3]
	2027–2070	—	64.0 [56.7–70.1]	83.4 [79.2–87.4]
Release retained by derogations (kt)	2027–2050	—	532 [348–817]	—
	2027–2070	—	842 [496–1240]	—
Derogation share (% of comprehensive avoided release)	2027–2050	—	38.0 [30.3–44.5]	—
	2027–2070	—	23.4 [16.5–31.2]	—

Final campaign, V0. Avoided release is the paired counterfactual-minus-restriction difference; release retained by derogations is the paired partial-minus-comprehensive difference. The derogation share is the paired partial-minus-comprehensive difference divided by comprehensive avoided release within the same run. Figure S2 draws the paired rows.

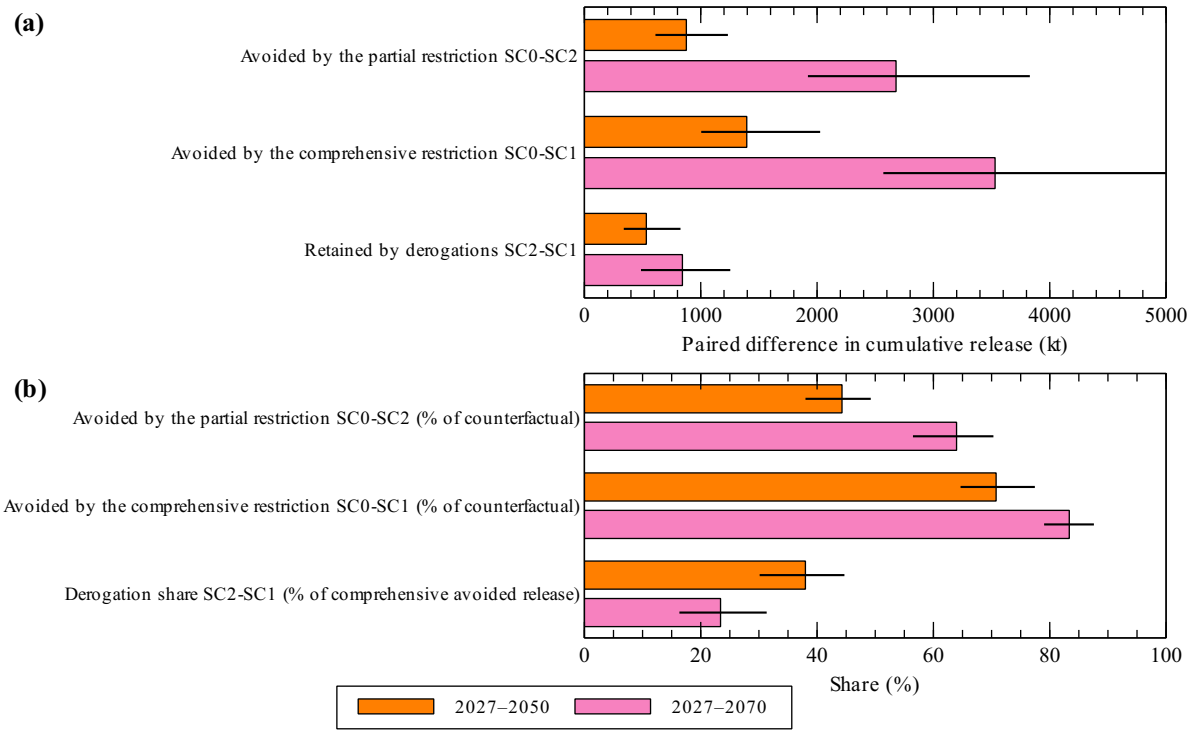


Figure S2: Avoided and retained release under the default parameterisation, 2027–2050 and 2027–2070. Panel a: the paired differences in cumulative release, kt. Panel b: the same quantities as shares, avoided release as a percentage of counterfactual release and the derogation share as a percentage of comprehensive avoided release. Bars are medians, whiskers the 5th to 95th percentiles. The values are those of Table S12.

Table S13: Magnitude and persistence of release from pre-2027 mass

Quantity	Partial restriction	Comprehensive restriction
Committed release, 2027–2070 (kt)	704 [585–817]	704 [585–817]
Committed share of scenario release, 2027–2070 (%)	45.2 [35.8–58.1]	99.3 [99.1–99.5]
Mean annual committed release, 2027–2031 (kt/a)	46.7 [38.3–55.1]	
Mean annual committed release, 2066–2070 (kt/a)	5.41 [4.45–6.38]	
Pre-2027 mass still upstream of release in 2070 (kt)	1250 [1050–1470]	
Share of pre-2027 market mass still upstream in 2070 (%)	15.8 [14.2–17.6]	
Committed release by window (kt)		
2027–2031	233 [191–275]	
2032–2036	144 [117–172]	
2037–2041	94.0 [75.0–113]	
2042–2046	66.4 [52.5–79.4]	
2047–2051	47.2 [38.7–55.6]	
2052–2056	35.2 [29.2–41.2]	
2057–2061	31.0 [25.7–36.2]	
2062–2066	29.1 [24.2–34.5]	
2067–2070	21.4 [17.6–25.3]	
Sector composition of committed release 2027–2070, mean (kt, % of total)		
Applications of fluorinated gases, S07	337 (47.8%)	
Textiles, upholstery, leather, apparel and carpets, S01	248 (35.4%)	
Construction products, S12	57.0 (8.19%)	
Food contact materials and packaging, S02	24.4 (3.50%)	
Medical devices, S08	15.2 (2.19%)	
All other sectors (10)	20.4 (2.94%)	

The committed stream is identical across scenarios within V0. The sector composition rows are compositions of mean masses and carry no band; they are the values of Figure 4b in the main text. The first four rows are calculated from the release-boundary arrays; the next two sum pre-2027 vintage mass in product-use and technical nodes at 2070 and divide it by cumulative pre-2027 market inflow, within run. The remaining mass makes the reported committed release a lower bound on eventual release from pre-2027 mass. Under no further restriction the committed share is 16.5 [12.5–20.7]%. The window rows sum release from pre-2027 mass over each window within the run; the last window has four years because 2027–2070 spans 44. Final campaign, V0; the annual pre-2027 series is in Online Resource 3, sheet `Flows_annual`, vintage `pre-2027`. Figure S3 draws the window rows.

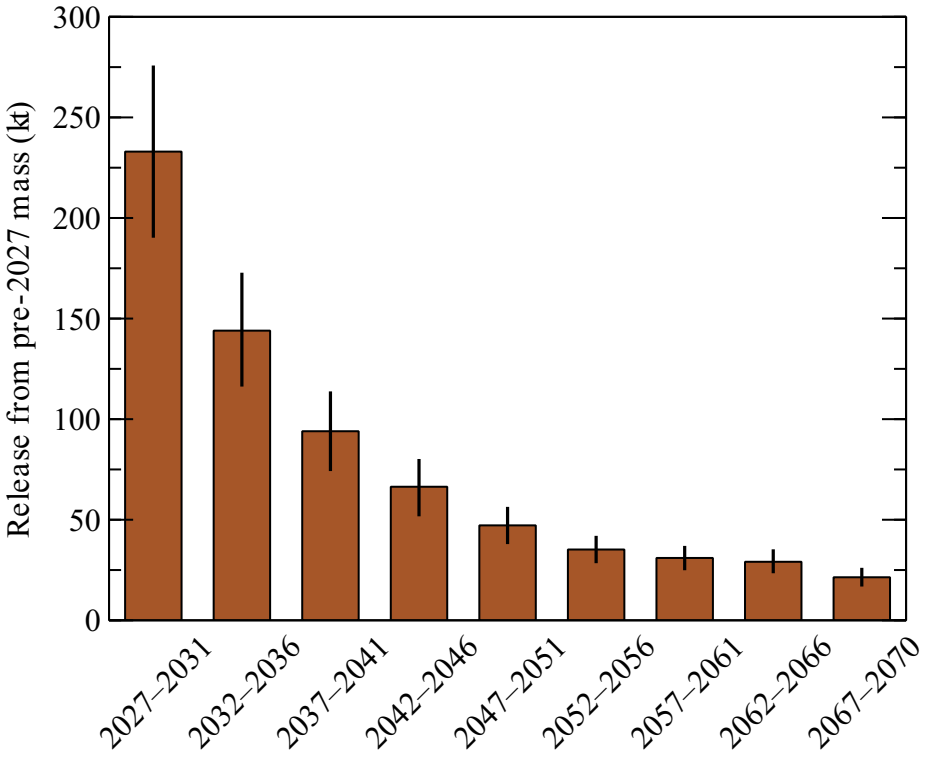


Figure S3: Release from mass placed on the market before 2027, by window, kt. Bars are medians, whiskers the 5th to 95th percentiles, formed within the run; the stream is identical across scenarios. The values are the window rows of Table S13.

S7.3 Environmental compartment stocks

Table S14: Environmental compartment stock and index in 2070

Compartment	Scenario	Stock median (kt)	p05–p95 (kt)	Index median [p05–p95]
Soil	Counterfactual	1740	1190–2290	402 [350–455]
	Partial restriction	587	417–753	137 [129–144]
	Comprehensive restriction	551	392–708	128 [122–135]
Groundwater	Counterfactual	76.1	53.9–99.9	287 [261–325]
	Partial restriction	65.9	45.7–88.2	249 [229–274]
	Comprehensive restriction	63.9	44.2–85.7	241 [223–264]
Surface water	Counterfactual	20.4	15.8–25.8	261 [206–334]
	Partial restriction	3.19	2.42–4.05	40.4 [32.8–48.5]
	Comprehensive restriction	3.06	2.30–3.91	38.8 [31.5–46.8]
Air	Counterfactual	287	154–429	37.1 [21.2–57.0]
	Partial restriction	220	91.7–362	28.5 [12.6–48.1]
	Comprehensive restriction	26.2	18.6–34.6	3.43 [2.82–4.01]
Sediment	Counterfactual	57.7	42.9–74.2	489 [424–565]
	Partial restriction	21.0	16.4–25.9	177 [161–197]
	Comprehensive restriction	19.4	15.1–24.1	164 [150–182]
Aggregate	Counterfactual	—	—	175 [135–217]
	Partial restriction	—	—	72.0 [58.1–88.4]
	Comprehensive restriction	—	—	53.5 [42.3–65.3]

The index is formed within each run as stock divided by its 2027 value and multiplied by 100. The aggregate row is the aggregate environmental-stock index: the five stocks are summed within each run and the sum is indexed to its 2027 value; it is not the mean of the compartment indices, and it excludes marine and atmospheric outflow. Air holds 61.6% [52.2–70.6] of the summed stock in 2027 and accounts for 82.9% [62.7–90.3] of the partial-minus-comprehensive difference in summed stock in 2070. In 2040 the aggregate index is 101 [90.3–113], 86.2 [78.4–94.0] and 75.6 [67.7–83.6] under the three scenarios. Final campaign, V0; the annual stocks of every compartment are in Online Resource 3, sheet `Stocks_annual`. Figure S4 draws the table.

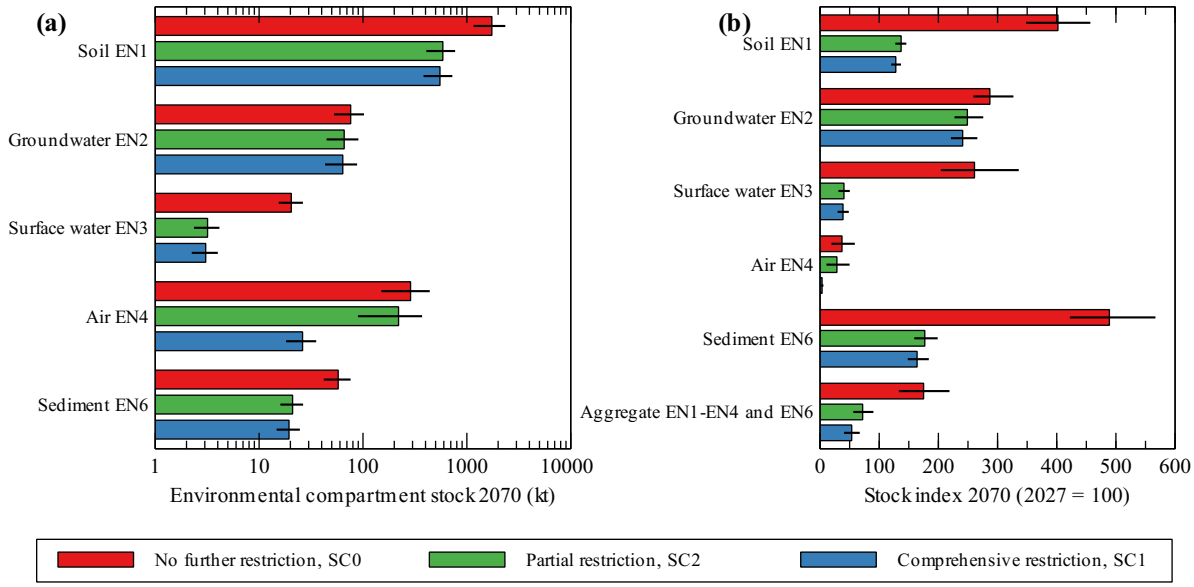


Figure S4: Environmental compartment stocks in 2070 under the three scenarios. Panel a: stock, kt, logarithmic axis. Panel b: stock index with 2027 = 100, including the aggregate over soil, groundwater, surface water, air and sediment. Bars are medians, whiskers the 5th to 95th percentiles. The values are those of Table S14.

S7.4 Fate of cumulative market inflow

Table S16 and Figure S5 give the mean fate of all mass placed on the market over 1950–2070 under each scenario: released, exported, mineralised or still held in products, technical compartments and the environment in 2070.

S7.5 Variant results

Across all variants, the comprehensive restriction avoided 75.4 to 90.3% of counterfactual release over 2027–2070 and the partial restriction 56.5 to 72.7%, against 83.4 and 64.0% under V0. Release from pre-2027 mass supplied 99.2 to 99.5% of release under the comprehensive restriction in every variant (Table S17). Its mean annual rate in 2066–2070 ranged from 3.39 kt/a with shorter service lives (VL1) to 9.25 kt/a with short technical and environmental half-lives (VE1), against 5.41 kt/a under V0; VE1 moves more of the pre-2027 mass through landfill, wastewater treatment and recycling within the horizon, VE2 (4.43 kt/a) less. The derogation share over 2027–2050 ranged from 30.3% (VM1) to 40.1% (VS2), against 38.0% under V0. The servicing variants change the absolute fluorinated-gas quantities, counterfactual release between 3670 and 4800 kt and the release retained by derogations between 697 and 990 kt, but not the derogation share (23.2 to 23.5%) nor its fluorinated-gas composition (90.4 to 93.3%, against 92.2% under V0); with less servicing the medical devices carry 63.8% of the retained release, with more 44.8%, against 52.5% under V0.

Table S15: Environmental compartment index in 2070 under the residence variants of the technical and environmental compartments

	VE1 (short)	V0	VE2 (long)
Comprehensive restriction, index 2070 (2027 = 100)			
Soil, EN1	142 [137–146]	128 [122–135]	126 [119–133]
Groundwater, EN2	196 [180–215]	241 [223–264]	342 [314–379]
Surface water, EN3	61.2 [51.6–70.2]	38.8 [31.5–46.8]	32.5 [25.9–40.3]
Air, EN4	2.51 [1.93–3.11]	3.43 [2.82–4.01]	4.04 [3.36–4.73]
Sediment, EN6	141 [127–159]	164 [150–182]	172 [157–190]
Aggregate, five compartments	75.1 [62.3–87.9]	53.5 [42.3–65.3]	46.5 [36.2–57.7]
Aggregate index 2070 under the other scenarios			
No further restriction	232 [186–278]	175 [135–217]	152 [116–191]
Partial restriction	93.3 [78.7–109]	72.0 [58.1–88.4]	65.4 [51.4–82.0]
Share of air in the 2027 stock (%)	49.0	61.6	67.0

Medians [5th–95th percentiles], each index formed within a run as the 2070 stock divided by the 2027 stock of the same run; the aggregate sums the five live compartments and excludes the marine and atmospheric outflow sinks. VE1 and VE2 place the 25 first-order half-lives at their declared low and high bounds (Table S6). The direction of every compartment is the same in all three parameterisations; the aggregate moves because the share of air in the 2027 stock moves with the atmospheric half-life, and because shorter technical half-lives deliver more of the pre-2027 mass to the environment within the horizon. V0 values agree with Table S14 to rounding. The servicing variants VS1 and VS2 leave every compartment but air unchanged and move the aggregate under the comprehensive restriction to 68.4 [56.2–80.2] and 43.9 [33.8–55.2], because the share of air in the 2027 stock moves from 61.6% to 50.9 and 68.5%.

Table S16: Mean fate of cumulative 1950–2070 market inflow

Scenario	Released	Mineralised	Exported	Still in system	Total
Counterfactual	6780 (25.7%)	9880 (37.4%)	762 (2.89%)	8970 (34.0%)	26400
Partial restriction	4000 (41.8%)	3800 (39.7%)	285 (2.98%)	1490 (15.6%)	9570
Comprehensive restriction	3150 (39.5%)	3310 (41.5%)	249 (3.13%)	1260 (15.8%)	7970

Masses are kt; percentages are shares of the row total. These are compositions of mean masses and carry no Monte Carlo band. Released mass is cumulative release since 1950, not the 2027–2070 quantity in Table S12. The total is mean cumulative 1950–2070 market inflow. Final campaign, V0. Figure S5 draws the table.

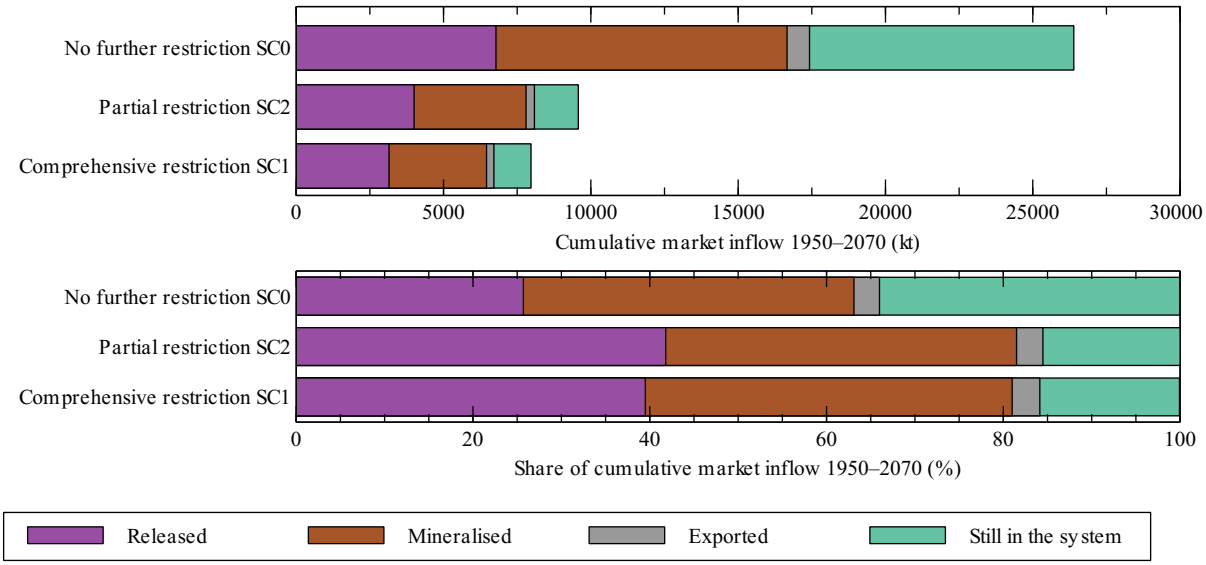


Figure S5: Mean fate of cumulative 1950–2070 market inflow under the three scenarios, as mass (upper panel, kt) and as shares of the row total (lower panel, per cent). Compositions of mean masses; no Monte Carlo band. The values are those of Table S16.

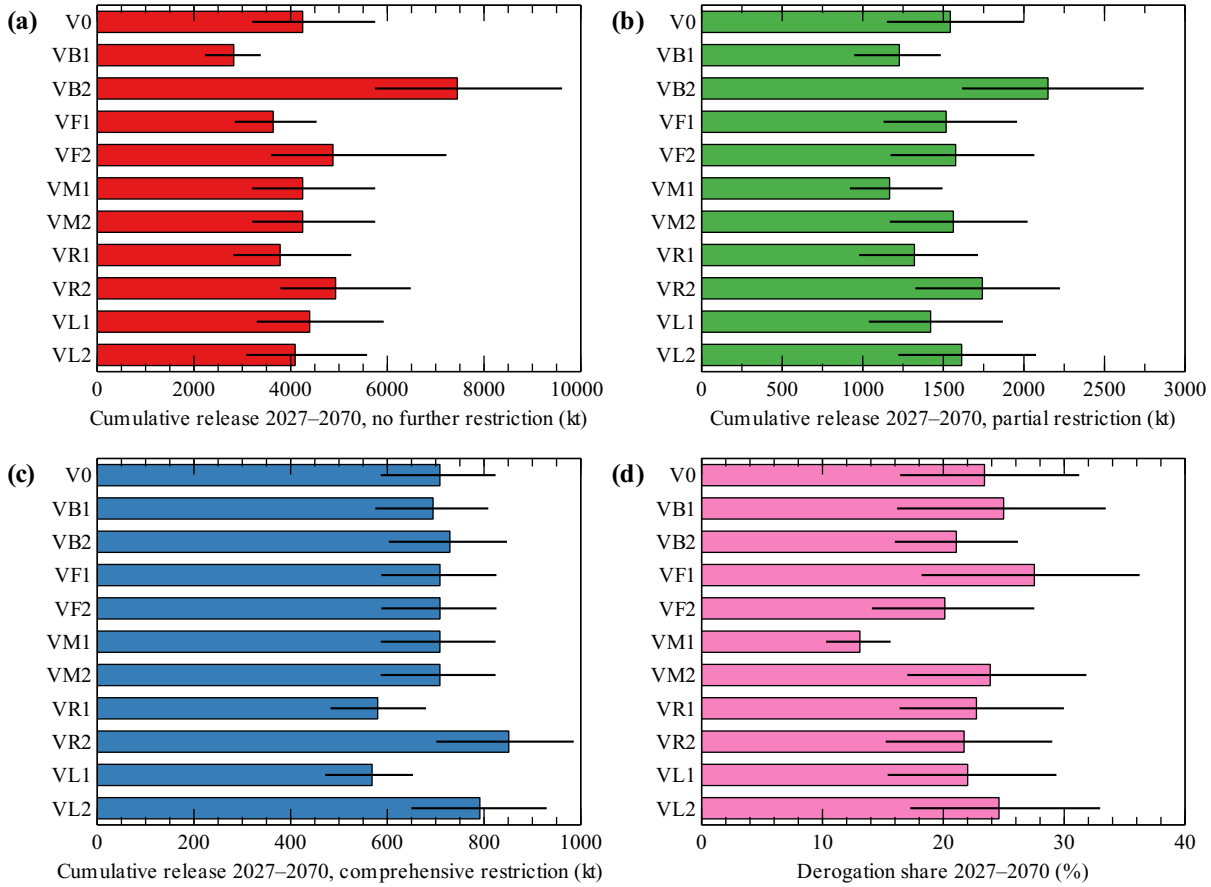


Figure S6: Cumulative 2027–2070 release under each scenario (panels a–c, kt) and the derogation share (panel d, per cent) for the default parameterisation and the declared variants. Bars are medians, whiskers the 5th to 95th percentiles, each formed within the runs of that variant. A difference between variants is a difference between parameterisations, not a Monte Carlo band. The values are those of Table S17.

Table S17: Cumulative 2027–2070 release, derogation share and committed release across variants

	Cumulative release 2027–2070 (kt)			Derogation share
ID	Counterfactual	Partial restriction	Comprehensive restriction	2027–2070 (%)
V0	4250 [3230–5730]	1540 [1160–2000]	709 [588–822]	23.4 [16.5–31.2]
VB1	2820 [2250–3360]	1230 [954–1480]	695 [577–807]	25.0 [16.3–33.3]
VB2	7450 [5770–9590]	2150 [1620–2740]	729 [605–845]	21.1 [16.1–26.1]
VF1	3640 [2870–4520]	1520 [1140–1950]	709 [590–823]	27.5 [18.3–36.1]
VF2	4870 [3620–7200]	1580 [1180–2060]	709 [590–823]	20.1 [14.2–27.4]
VM1	4250 [3230–5730]	1170 [927–1490]	709 [588–822]	13.1 [10.4–15.6]
VM2	4250 [3230–5730]	1560 [1170–2020]	709 [588–822]	23.9 [17.1–31.7]
VR1	3780 [2840–5230]	1320 [984–1710]	580 [485–678]	22.7 [16.5–29.9]
VR2	4930 [3810–6460]	1740 [1330–2220]	851 [704–983]	21.7 [15.3–28.9]
VL1	4390 [3320–5900]	1420 [1040–1860]	568 [474–651]	22.0 [15.5–29.3]
VL2	4090 [3110–5560]	1610 [1230–2070]	791 [652–927]	24.6 [17.4–32.9]
VE1	4390 [3350–5860]	1690 [1310–2150]	850 [706–984]	23.5 [16.6–31.3]
VE2	4210 [3200–5690]	1510 [1120–1960]	676 [559–788]	23.4 [16.5–31.1]
VS1	3670 [2890–4590]	1380 [1030–1760]	696 [581–805]	23.2 [14.8–32.1]
VS2	4800 [3560–6950]	1700 [1260–2240]	721 [597–837]	23.5 [17.7–30.6]

	Derogation share	Release from pre-2027 mass	
ID	2027–2050 (%)	Share of comprehensive release (%)	2066–2070 mean (kt/a)
V0	38.0 [30.3–44.5]	99.3	5.41 [4.45–6.38]
VB1	36.1 [27.5–44.8]	99.4	5.32 [4.36–6.28]
VB2	39.0 [34.0–43.2]	99.3	5.55 [4.58–6.55]
VF1	40.0 [31.1–48.0]	99.3	5.41 [4.44–6.39]
VF2	36.2 [29.4–42.4]	99.3	5.41 [4.44–6.39]
VM1	30.3 [24.3–35.6]	99.3	5.41 [4.45–6.38]
VM2	38.9 [31.6–45.4]	99.3	5.41 [4.45–6.38]
VR1	38.0 [30.5–44.2]	99.5	4.37 [3.59–5.18]
VR2	35.4 [27.9–42.0]	99.4	6.29 [5.16–7.39]
VL1	35.7 [28.4–42.2]	99.2	3.39 [2.63–4.20]
VL2	40.0 [32.2–46.5]	99.4	6.00 [4.98–7.02]
VE1	38.0 [30.3–44.5]	99.4	9.25 [7.36–11.2]
VE2	38.0 [30.3–44.5]	99.3	4.43 [3.66–5.21]
VS1	34.8 [26.1–42.9]	99.3	5.41 [4.44–6.39]
VS2	40.1 [33.2–45.5]	99.4	5.41 [4.46–6.38]

Medians [5th–95th percentiles], formed within each run of the variant; the committed share of release under the comprehensive restriction is given as its median only. The upper block holds the release and derogation-share results, the lower block the shorter-window derogation share and the pre-2027 quantities for the same variants. Variant definitions and families are in Table S6. Medical variants change only the partial restriction; their counterfactual and comprehensive entries repeat V0 by construction, and their derogation share pairs each run with the same run of V0. Final campaign, all 29 sets, and the extension campaign for VE1, VE2, VS1 and VS2 (Section S8.2). Figure S6 draws the four columns of the upper block.

Table S18: Sector contributions to release left in place by derogations

	Means				Within run
Sector	G1	G2	G3	Sum	Median [p05–p95]
Medical devices, S08	8.25	435	2.09	445	443 [158–748]
Applications of fluorinated gases, S07	0	347	0	347	321 [178–594]
Textiles, upholstery, leather, apparel and carpets, S01	8.59	0	30.2	38.8	38.5 [25.6–53.1]
Electronics and semiconductors, S10	7.63	0.364	2.27	10.3	10.0 [6.97–14.1]
Lubricants, S13	0.0188	1.67	4.11	5.80	5.80 [4.82–6.77]
All other sectors (10)	1.33	0	1.72	3.05	2.99 [not reported]

Masses are kt for the paired partial-minus-comprehensive difference over 2027–2070. G1, G2 and G3 are the substance groups of Table S3: PFAAs and their precursors, fluorinated gases, and fluoropolymers and perfluoropolyethers. The four *Means* columns are compositions of mean masses, so the substance-group entries add to the group sum; the last column is the sector total formed within each run and then summarised, which is why it differs from the group sum; it is the value cited in the main text. The bundled remainder carries no band because component quantiles cannot be added. Final campaign, V0.

S7.6 Sector evidence for the derogation result

Fluorinated gases make up 92.2% [86.7–95.0] of the release retained by derogations over 2027–2070, formed within the run from the substance-class arrays.

For the comparison of retained and stopped inflow in the main text, the paired differences isolate the post-2027 market inflow each decision keeps or stops, and pre-2027 mass cancels. The derogations retain 1590 kt [1220–2050] of market inflow over 2027–2070, 8.74% [6.80–11.0] of the inflow the comprehensive restriction stops. By 2070, 52.8% [40.6–61.6] of this retained inflow has been released, against 16.1% [13.0–21.2] of the post-2027 inflow stopped by the partial restriction, whose release is read under no further restriction. For the group comparison in Figure 5b of the main text, which takes whole uses over 1950–2070 and therefore includes their pre-2027 inflow, uses retained by a derogation released 53.8% of their cumulative market inflow (3160 of 5870 kt) and uses stopped by the partial restriction 22.7% (848 of 3730 kt). These are compositions of mean masses and carry no band.

S8 Reproducibility

This section records the final campaign and the input hashes. The code, the input files, the campaign manifest and the simulation output are archived in Zenodo at `https://doi.org/10.5281/zenodo.22751472`. Reproduction requires the code and original input bytes as well as the seed and numerical settings; bitwise agreement must be checked in the specified software environment, and was confirmed for the archived runs in Section S6.1.

S8.1 The campaign

Table S19 records the settings under which the 29 sets of the final campaign were computed, as written by the engine into the manifest and run log.

S8.2 The extension campaign

After the pre-submission review, four variants were added on 2026-09-15 (VE1, VE2, VS1, VS2; assumptions A555, A556; Table S6) and run as an extension of 12 sets, the three scenarios under each, at the same seed, pack count and run count, into its own folder of the archive

Table S19: Settings of the completed campaign

Setting	Value
Master seed	42
Packs per cohort (`n_packs`)	200
Runs per set (`n_runs`)	1000
Campaign plan	`minimal`: 29 sets
Sets	three scenarios under V0, VB1, VB2, VF1, VF2, VR1, VR2, VL1, VL2; the partial restriction under VM1 and VM2
Pack-count criterion	evaluated separately (Section S5)
Horizon	1950–2070, 121 annual steps
MATLAB	25.2.0.3123386 (R2025b) Update 3
Campaign started	9 September 2026, 00:00:31 CEST
Wall time	60 h 57 min for the 29 sets
Largest mass-balance residual	6.0×10^{-8} t (absolute, annual, any set)
Engine commit	`20a4eb48db2eff0b52ba72ca59aa326497c59e3b`, branch `figures-2026-09-08`, with uncommitted changes; the input tree was not pinned to a commit, and the digests of Table S20 identify the data

Sources: the campaign manifest and run log written by the engine, which record the seed, the two counts, the plan, the sets, the per-table digests, the git state, the MATLAB version and the wall time. The pack count and the run count are arguments of the campaign; the engine chooses neither.

(`05_Simulation_Output/extension`) with its own manifest. The extension read an input state that differs from the final campaign's only in the added variant rows; every default (ALL) row is byte-identical, so the 29 sets of the final campaign remain valid and were not recomputed. The engine refuses to resume into a folder whose manifest carries another input digest, which is why the extension has a folder of its own. Its per-table hashes differ from Table S20 for the five tables that gained rows: A5 `c1fc848e 66c0720a`, B1 `b30952bc 4970e717`, B3 `06873b97 783554fd`, B5 `d6d7f578 05aa8fdc`, C2 `2d0b671a 0de2405e` (first 16 characters); the combined digests are `a8bb356d 207eddc9 4f48b988 339665d7 6ab938a2 b158b279 de977d56 0ba64e7c` over A1–B5 and `675caad5 3e23b938 51d10990 4c2afe94 25400a91 1ad6a4fd 350c4670 de07363b` over C1–C2. Online Resource 2 reproduces the extended state, and its index sheet lists these hashes. After the extension, on 2026-09-16, the provenance tables were revised once more without touching any engine-read table: fifteen assumptions were re-sourced and regraded and twenty-six references added (Section S3.2); on 2026-09-22 the reference metadata were corrected, thirteen references added and two audit notes extended, so the deposited C1 and C2 carry the hashes `276baf5b 1e051cb7 407297e2 7f68f76a d3d52046 10c63044 192d4c47 2614cb1a` and `08f5785f 7fe999a0 f4c4f7f5 76729760 3a8dadd6 547ae510 60630c5d ae0a9693` and the provenance digest `01b74519 3e14d184 2b7b75e6 17335072 7162cb55 8573b383 f4126b95 3fa389af`; the engine digests of both campaigns are unaffected by construction. The extension started on 15 September 2026 at 13:17 CEST, was interrupted by the host session after eight sets, and was resumed on 16 September at 09:20 with the finished sets skipped by the engine's resume guard (Section S8.1); the partly computed ninth set was recomputed from scratch, and the twelfth set completed at 18:24 after 26 h 18 min of set time in total. The largest annual mass-balance residual in the 12 sets was 5.2×10^{-8} t; the archive checks P1–P5 on the 12 sets are reported in Section S6.1. Also on 2026-09-15 the six sets of the routing variants were relabelled to the link order of V0 (Section S6.1), and the engine now places variant rows in the position of the rows they replace; a rerun of VR1 or VR2 with the deposited code therefore reproduces the archived numbers cell for cell but not the archived random realisations, whereas every other set, V0 included, is unaffected.

S8.3 Input digest

The engine hashes every input table at the start of a campaign and writes the digests to the manifest (Table S20). Ten tables are parsed; the two provenance tables are hashed for the record and never read, so they cannot affect a result. Two combined digests are recorded: one over the ten parsed tables, which is the digest a result depends on, and one over the two provenance tables. Each combined digest is the SHA-256 of the lines `<table>:<hash>` in table order, so it can be recomputed outside MATLAB from the per-table hashes. The convergence criterion file is hashed the same way by the convergence tool.

Table S20: SHA-256 digests of the input tables as read by the campaign

Table	SHA-256
A1 nodes	`38f4e420 fae91fbe f6f90ebc 369fc8ca 57e5ea12 a2a1bbbc 88bf86e7 6b0a6858`
A2 substances	`57565c16 38f6e7f8 2c28453e 7a8592aa fa3428e9 82a2d045 0d081363 a9b6e037`
A3 sectors	`5e2934e4 1d332aff 8c90c5ee c4b2be50 76a10322 62da4723 44da04a4 f5a37b47`
A4 scenarios	`2adf954b 331fbd9e 4b4c97bf 3c165236 90a941f1 d984fd59 39a07061 6bc6feff`
A5 variants	`6b00de68 afad9c6a 2b1b1ff9 67f2b110 ceeef602 59db4ee0 3ba21871 2f88cd54`
B1 market inflow anchors	`3e5394d0 f3f4bf50 5ddd45d2 e0b09aff 0ec2ac3e 8c05bf1a 555b5bb1 bd55cdde`
B2 growth rates	`d57b59d1 20047a9b c5ff85d2 2b5321f5 31468f73 4f373417 164fc882 ddaed09d`
B3 allocation shares	`179b1505 3ac14812 f785ae63 2cb71b23 0367ce40 e5994666 e0b036b3 7589e9b2`
B4 transfer weights	`f613a886 4add6a39 eacd8aff 747aad22 f945dcbf 453dc228 5cfa7f06 df3cbd8e`
B5 residence times	`5a74516b d3078d17 caf351eb 3bc7790a 54df7f9c 5af339ab ef777730 2b5a462f`
C1 references	`50f3fa53 8f72a5ef 58dad156 2e4697ad c9e2014a 4cbb33ab 036b914b ba44fa8d`
C2 assumption log	`305e4357 8c9fff61 8476ac99 888e2b75 63b8cf13 59e057b7 61fa470b d10748d3`
Digest over A1–B5	`891f3b87 13626a1a d319a19e 1834825b 0ca2ceed e7cc0111 0a492311 9e513962`
Digest over C1–C2	`b17fb74b 9812854b e380fb78 be8c0125 6a068f29 11e5beb4 42c19a3c 5489692a`
Convergence criterion file	`3ad7d346 218c71da be25307d 4edb384e fd7aeacd 06335976 1eaf8cb2 ce737b5e`

The per-table hashes are those the final campaign recorded in its manifest; the source files of Online Resource 2 carry the same hashes, and its index sheet lists them. Thin spaces break the 64 hexadecimal characters into groups of eight for reading only.

Matching the deposit to the campaign. The deposited CSV files and the source files of Online Resource 2 were hashed before deposit and carry the hashes of Table S20. Recreating a CSV from a spreadsheet may change quoting, encoding or line endings even when the cell values agree, so the deposited bytes, not the worksheets, are the input of record. An earlier campaign (seed 42, 100 packs per cohort, 21 sets, completed 2026-08-31) read a previous state of the tables; on 2026-09-04 engine-read values were changed (Section S3.1), so that campaign is superseded and none of its numbers is reported. A second earlier campaign (seed 42, 200 packs, 29 sets, completed 2026-09-05, input digest `02b2f88c1ec5`) read the state before the correction of the fluorinated-gas derogation scope (A551–A553) and the scope review of the comprehensive restriction (A554); it is superseded as well; the pack-count gate was first passed on it and repeated on the final state (Section S5), and no other assessment rests on it. The changes themselves are recorded in the assumption log (Online Resource 2, C2).

S8.4 Random numbers

No random-number generator state is carried. Every uniform draw is a function of a 64-bit key built from the master seed, the run index, the cohort and the pack index, and of a counter that advances within the pack (destination first, then residence time, at every arrival). The mixing function is SplitMix64. Because the key never contains the scenario, a pack draws the same numbers under every scenario until a scenario routes it differently, and because no state is shared between packs, results do not depend on the order in which packs are processed or on the number of threads. The inflow parameters of a run are drawn from a key built from the seed and the run

index alone, so every scenario and every variant that does not override the inflow tables sees the same market in run r.

S8.5 Re-running the campaign

1. Retrieve the original twelve CSV files from the final repository deposit into `01_Input/`. Recompute their SHA-256 hashes and compare with Table S20. Use the original bytes for exact reproduction; Online Resource 2 is a tabular companion.

2. In MATLAB R2025b, run the engine's campaign entry point with 200 packs, 1000 runs, seed 42 and the `minimal` plan. It runs the test suite and the 29 sets and writes the archive with a new manifest; the 12 sets of the extension are run with the explicit set list of Section S8.2 into a separate folder. The pack-count ladder and the run-count bootstrap are run separately with the convergence tool.

3. Compare the per-run arrays of any set with the archived ones; they are expected to be identical bit for bit (the reproduction check of the deposited code does this for the first five runs of the counterfactual).

4. Run the archive checks and the result modules of the deposited code. The tables of Section S7 and the figures of the main text are reductions of the archive formed within the run index as Section S9 states; the tables printed here and the series of Online Resource 3 are the record. The convergence statistics are regenerated by the convergence tool, the external comparison by the comparison module.

Step 3 was carried out on the final archive and passed (Section S6.1).

S9 Reading rules for the archived results

The archive contains detailed cells as well as aggregated outcomes. The three aggregation rules stated at the head of Section S7 — sum within a run before summarising, form every ratio within the run, and use mean masses for compositions — apply to the archive as they do to the tables. The following rules cover what those three do not.

1. **Pair only matched markets.** Scenario differences within one parameterisation, and the eight variants that override the scenario table, the routing weights or the residence times (VM1, VM2, VR1, VR2, VL1, VL2, VE1, VE2) against the default parameterisation, can be formed run by run, because the market draw before scenario cutoffs reads none of those tables. The scenario table subsequently changes the surviving inflow. The four B2 growth variants and the two servicing variants, which override the anchor and allocation tables, change the market realisation and are reported as distributions. Shared random numbers could still define a numerical coupling, but it would not hold the realised market fixed.

2. **Read the bounding variants as brackets.** The routing and residence variants place every affected row at a declared bound at once; they describe no system a source proposes and are reported as selected stress tests. They are not guaranteed lower and upper bounds on outcomes, and their spread is not a probability interval. The service-life bracket (VL1, VL2) covers the product-use stock; the residence times of the technical and environmental compartments are bracketed by VE1 and VE2 within the windows the sources declare (Table S15), so a statement about how an environmental stock responds carries that bracket, not a probability band.

3. **Separate structural equality from empirical support.** Under fixed routing and residence rules, pre-2027 cohorts have the same paths across scenarios. Their equality

follows from the model design. Variants test some of the assumptions behind the amount and timing of that release; they do not test every omitted process.

4. **Treat the horizon as censored.** A terminal value does not imply that all mass has completed its pathway by 2070. Release near the horizon is reported as the 2066–2070 mean because the single year 2070 is more sensitive to pack discretisation.

5. **Follow the three research questions.** Q1 connects EEA-wide market inflow as the model input to annual and cumulative environmental release as outcomes. Q2 uses the aggregate pre/post-2027 release split, committed share and paired release retained by derogations to identify what sustains release after entry into force. Q3 uses the aggregate environmental-stock index, supported by the five compartment indices, to describe environmental stock response across the EEA. Sector, substance-group and node detail supports these quantities but does not replace their EEA-wide reported forms.

These rules apply to every table and machine-readable result in the archive unless a table legend states a more specific procedure.

References

European Chemicals Agency (2023a). *Annex A to the Annex XV Restriction Report on PFASs: Manufacture and uses* (Tech. Rep.). European Chemicals Agency. Retrieved 2026-08-31, from https://echa.europa.eu/restrictions-under-consideration/-/substance-rev/72301/term

European Chemicals Agency (2023b). *Annex B to the Annex XV Restriction Report on PFASs: Information on hazard, exposure and emissions* (Tech. Rep.). European Chemicals Agency. Retrieved 2026-08-31, from https://echa.europa.eu/restrictions-under-consideration/-/substance-rev/72301/term

European Chemicals Agency (2023c). *Annex E to the Annex XV Restriction Report on PFASs: Impact assessment for specific uses* (Tech. Rep.). European Chemicals Agency. Retrieved 2026-08-31, from https://echa.europa.eu/restrictions-under-consideration/-/substance-rev/72301/term

European Chemicals Agency (2023d). *Annex XV Restriction Report: Per- and polyfluoroalkyl substances (PFASs)* (Annex XV restriction report). European Chemicals Agency. Retrieved 2026-08-31, from https://echa.europa.eu/restrictions-under-consideration/-/substance-rev/72301/term (Version 2. Dossier submitters: BAuA (DE), RIVM (NL), KEMI (SE), Norwegian Environment Agency, Danish EPA)

European Commission (2024). *Commission Regulation (EU) 2024/2462 of 19 September 2024 amending Annex XVII to Regulation (EC) No 1907/2006 as regards undecafluorohexanoic acid (PFHxA), its salts and PFHxA-related substances.* Official Journal of the European Union, L 2024/2462. Retrieved 2026-09-01, from http://data.europa.eu/eli/reg/2024/2462/oj

European Commission (2025). *Commission Regulation (EU) 2025/1988 of 2 October 2025 amending Annex XVII to Regulation (EC) No 1907/2006 as regards per- and polyfluoroalkyl substances in firefighting foams.* Official Journal of the European Union, L 2025/1988. Retrieved 2026-09-01, from http://data.europa.eu/eli/reg/2025/1988/oj

European Parliament and Council (2024). *Regulation (EU) 2024/573 of 7 February 2024 on fluorinated greenhouse gases, amending Directive (EU) 2019/1937 and repealing Regulation (EU) No 517/2014.* Official Journal of the European Union, L 2024/573. Retrieved 2026-09-01, from http://data.europa.eu/eli/reg/2024/573/oj

European Parliament and Council (2025). *Regulation (EU) 2025/40 of 19 December 2024 on packaging and packaging waste, amending Regulation (EU) 2019/1020 and Directive (EU) 2019/904, and repealing Directive 94/62/EC.* Official Journal of the European Union, L 2025/40. Retrieved 2026-09-01, from http://data.europa.eu/eli/reg/2025/40/oj