MODEL DOCUMENTATION

EnergyPATHWAYS Model

United States

Date	October 26, 2026
Package	EnergyPATHWAYS and RIO Documentation, October 26, 2026

Contents

1 Model structure

- 1.1 Approach
- 1.2 Sectors and subsectors
- 1.3 Demand drivers and projection of inputs
- 1.4 Notation

2 Geography and time

- 2.1 Geographies
- 2.2 Years, units, and currency

3 Service demand

- 3.1 Projection
- 3.2 Service links

4 Stock rollover

- 4.1 Total stock
- 4.2 Survival and retention
- 4.3 Rollover recursion
- 4.4 Sales shares
- 4.5 Specified stocks, linked technologies, and early retirement
- 4.6 Logit technology choice

5 Technology parameters

- 5.1 Efficiency
- 5.2 Service demand modifiers
- 5.3 Costs

6 Energy demand

- 6.1 Subsectors with a technology stock
- 6.2 Subsectors without a technology stock
- 6.3 Calibration to energy benchmarks

7 Measures

8 Load shapes

- 8.1 Shapes
- 8.2 Assignment and aggregation
- 8.3 Reconciliation to system load

9 Outputs and export to RIO

- 9.1 Outputs
- 9.2 Tables passed to RIO
- 9.3 Mode (a): demand optimization off
- 9.4 Mode (b): demand optimization on

References

This document describes the EnergyPATHWAYS (EP) model as applied to the United States energy system. It covers the structure of the demand side, the projection of service demand, the stock rollover, the treatment of technology parameters and measures, the calculation of final energy demand and its hourly profile, the handling of geography, and the tables EP passes to the RIO supply model. The data and settings used in particular runs are documented in the *EnergyPATHWAYS US Database* document, and the supply model is documented in the *RIO Model* document. Williams et al. (2021) give an earlier description of the framework.

1 Model structure

1.1 Approach

EP is an accounting model of energy demand. It does not optimize. For each demand subsector it projects a service demand, the equipment stock that serves it, and the final energy the stock consumes, given drivers, technology characteristics, and specified shares of new equipment sales. EP is solved annually over its projection horizon. Every subsector is calculated independently, except where one subsector supplies a driver or a stock to another (Sections 3.2 and 4.5). No prices pass from RIO to EP.

The calculation proceeds in this order:

1. Drivers are read and mapped to the disaggregation geography (Section 2).
2. For each subsector, service demand or energy demand is projected from its drivers (Section 3).
3. Where the subsector has a technology stock, the total stock is projected and rolled over by technology and vintage (Section 4).
4. Measures modify service demand, sales shares, or energy demand (Section 7).
5. Final energy demand is calculated by final energy type, zone, and year (Section 6), together with demand-side costs (Section 5.3).
6. Hourly load profiles are built bottom up and reconciled to system load (Section 8).
7. Results are written, together with the tables passed to RIO (Section 9).

1.2 Sectors and subsectors

The demand side is divided into sectors (residential, commercial, industrial, transportation) and, within each sector, subsectors. A subsector is the accounting unit: an end use (residential space heating), an industry (paper and allied products), or a transport mode (passenger car). Data in a subsector can carry up to two further indices, such as housing type, building type, or industrial end use.

How a subsector is calculated depends on which input tables it has. EP assigns one of six subsector types:

Type	Inputs	Calculation
Stock and service	Service demand, total stock, technology characterization	Service demand is served by a technology stock rolled over by vintage; energy is service demand times stock-weighted energy intensity
Stock and energy	Energy demand in data years, total stock, technology characterization	Energy is divided by stock-weighted energy intensity to give service demand in the data years; service demand is then projected and served as in the stock and service type

Type	Inputs	Calculation
Service and efficiency	Service demand and an aggregate service efficiency	Energy is service demand times the energy per unit of service
Service and energy	Service demand and energy demand	Both are projected; service demand is a driver of energy demand
Energy	Energy demand by final energy type	Energy demand is projected from its drivers
Link	A technology stock with no energy use	The stock modifies the service demand of other subsectors through service links (for example, a building shell stock modifies heating and cooling service)

Only the first two types, and the link type, carry a technology stock. Technology choice, efficiency standards, and sales share measures therefore apply only to those subsectors. The other types are changed by energy efficiency, fuel-switching, and service demand measures (Section 7).

A stock can be counted in units of equipment (vehicles, water heaters), in units of service capacity, or as a capacity factor. A capacity-factor stock is converted to a service-capacity stock by dividing average service demand per unit time by the capacity factor.

1.3 Demand drivers and projection of inputs

A driver is an exogenous series, such as population, households, floor area, value of shipments, or vehicle miles traveled, that scales a subsector input. A driver can itself be specified relative to a base driver, for example floor area per household. Each projected input (service demand, energy demand, total stock) names up to three drivers and up to two denominator drivers.

Inputs are entered for the years for which data exist and are projected to all model years. EP converts each input into an intensity per unit of its combined driver, extends that intensity over the model years with the interpolation and extrapolation methods the input names (for example linear interpolation, holding the nearest value, logistic, exponential, or regression methods), and multiplies by the driver in each year. For an input X entered as a total for region g of its input geography in data years t, with combined driver $D = \prod_k D^{(k)}$ on the disaggregation geography,

$$X_{d,y} \;=\; \mathcal{T}_y\left[\frac{X_{g,t}}{\sum_{d' \in g} D_{d',t}}\right] D_{d,y}, \qquad d \in g \tag{1}$$

where $\mathcal{T}_y$ denotes interpolation or extrapolation of the intensity to year y. The result is mapped to the primary geography (Section 2.1). An input entered as an intensity per unit of a denominator driver (for example energy per household) is multiplied by that driver before (1) is applied. An input with no driver is interpolated and extrapolated directly.

1.4 Notation

Symbol	Description
y, v, a	Year, vintage, and age $a = y - v$
z, d, g	Region of the primary geography, of the disaggregation geography, and of an input's own geography

Symbol	Description
τ, i, j	Demand technology; i a retiring technology and j a purchased technology
f	Final energy type
h	Hour of the weather year
SD_y	Service demand
$N_{\tau,v,y}$	Stock of technology τ and vintage v in year y; N^0 the stock present in the first year
$Q_{\tau,y}$, $R_{\tau,y}$	Sales and retirements
$S_\tau(a)$, $p_\tau(a)$	Survival function and one-year retention probability
$\sigma_{j,i,y}$	Share of the replacement of retiring technology i purchased as technology j
$e_{\tau,v,y}$	Energy intensity (energy per unit of service)
u_τ	Utility factor: share of service delivered on the main final energy type
$m_{\tau,v,y}$	Service demand modifier
$E_{f,y}$	Final energy demand
L_τ	Book life used to levelize cost
W	Number of weather years

2 Geography and time

2.1 Geographies

EP holds a set of named geographies, each a partition of the modeled area into regions: for example nation, census region and division, state, county, climate zone, and the model zones. A spatial join table lists every county, or county fragment where a county is split between regions, with the region it belongs to in each geography and a set of allocation weights called map keys (households, population, floor area weighted by degree days, industrial activity, and others). Every input row records the geography it is entered on and, optionally, the map key used to move it to another geography. Rows that name no key use a default key.

Two geographies are set for a run:

- The **primary geography** is the geography on which EP reports results and passes data to RIO.
- The **disaggregation geography** is the finer geography on which drivers are held and on which (1) is evaluated. It allows a driver to distribute an input that is entered on a coarse geography, so that, for example, national service demand per household is allocated across zones by where the households are.

A run can also restrict the model to a subset of the regions of the primary geography, and can form a composite geography by replacing selected regions of a geography with the regions of a finer one. Where an input is given both for a large region and for a smaller region inside it, EP subtracts the smaller region's data from the larger region so that the two together reproduce the original total.

Let $k_{s,t}$ be the sum of the map key over the county fragments that lie in source region s and target region t. A total is split in proportion to the key, and an intensity is averaged with the key as weight:

$$X_t = \sum_s X_s \frac{k_{s,t}}{\sum_{t'} k_{s,t'}} \quad \text{(totals)}, \qquad x_t = \sum_s x_s \frac{k_{s,t}}{\sum_{s'} k_{s',t}} \quad \text{(intensities)} \tag{2}$$

2.2 Years, units, and currency

Each subsector is calculated annually from the earliest year in its inputs (and not before 1970) to the end year of the run. Vintages span the same years. A configured current year separates historical data from projection: technology stock data after the current year are discarded unless the input specifies otherwise, so the technology mix of later years comes from sales shares.

Inputs carry their own units. EP converts energy to a single calculation energy unit and mass to a single mass unit at load time. Costs are entered in the currency and currency year of each row and are converted to a single model currency year with an inflation index.

3 Service demand

3.1 Projection

Service demand is projected with (1) from the subsector's service demand input and its drivers. In a stock and energy subsector, service demand is first derived from energy. For each data year, the energy input is divided by the energy intensity of the stock in that year,

$$SD_t \;=\; \frac{E_t}{\bar{e}_t}, \qquad \bar{e}_t = \sum_{\tau,v,f} \frac{N_{\tau,v,t}}{\sum_{\tau',v'} N_{\tau',v',t}}\, e_{\tau,v,t,f} \tag{3}$$

and the resulting series is projected with (1). When the energy input is indexed by technology or by final energy type, the division is made within each technology or final energy type.

The total stock and service demand can be related in three ways, set for each subsector: each is projected from its own drivers; the stock is driven by service demand (for example vehicles from vehicle miles traveled); or service demand is driven by the stock.

3.2 Service links

A technology in one subsector can change the service demand of another subsector. The hot water drawn by clothes washers and dishwashers is a share of water heating service; the building shell changes space heating and air conditioning service. For a linked subsector ℓ with precursor subsectors q, EP calculates the stock-weighted link factor $\bar{\lambda}_{q,y}$ of each precursor's technologies and applies to the service demand of ℓ the multiplier

$$\Lambda_{\ell,y} \;=\; 1 - \sum_q \phi_{q,\ell} \left(1 - \frac{\bar{\lambda}_{q,y}}{\bar{\lambda}_{q,y_c}}\right) \tag{4}$$

where $\phi_{q,\ell}$ is the share of the service demand of ℓ attributable to q and y_c is the link's calibration year. The multiplier enters (1) as an additional driver. Subsectors are calculated in an order that places each precursor before the subsectors that depend on it.

4 Stock rollover

4.1 Total stock

The total stock of a subsector is projected with (1), separately for each combination of zone and the subsector's other indices (a rollover group). The rollover is solved for each group. The annual change in the total stock, $\Delta N_y = N_y - N_{y-1}$, together with retirements, determines total sales.

4.2 Survival and retention

Each technology has a lifetime distribution, entered as a mean and variance or as a minimum and maximum. The survival function is Weibull by default:

$$S_\tau(a) = \exp\left[-\left(\frac{a}{\alpha_\tau}\right)^{\beta_\tau}\right], \qquad \alpha_\tau = \frac{\mu_\tau}{\Gamma(1+1/\beta_\tau)} \tag{5}$$

where μ_τ is the mean lifetime and β_τ is solved so that the distribution has the specified variance. Linear and exponential survival functions are also available. Where only a minimum and maximum lifetime are given, the mean is their midpoint and the standard deviation is a quarter of their range.

Equipment present in the first year has a mix of ages. EP assumes this stock has the age distribution of a steady state, which gives the survival function of the initial stock

$$S^0_\tau(a) = \frac{\sum_{k \geq a} S_\tau(k)}{\sum_{k \geq 0} S_\tau(k)} \tag{6}$$

The retention probability over one year is $p_\tau(a) = S_\tau(a+1)/S_\tau(a)$, and $p^0_\tau(a)$ is defined in the same way from S^0_τ. If one-year survival of any technology in a subsector is below 0.95, the rollover is solved with more than one step per year.

4.3 Rollover recursion

Each vintage retires along its survival curve:

$$N_{\tau,v,y} = N_{\tau,v,y-1}\, p_\tau(y-1-v), \quad v < y; \qquad N^0_{\tau,y} = N^0_{\tau,y-1}\, p^0_\tau(y-1-y_0) \tag{7}$$

where y_0 is the first year of the subsector. Retirements in year y are

$$R_{\tau,y} = \sum_{v<y} N_{\tau,v,y-1}\left(1 - p_\tau(y-1-v)\right) + N^0_{\tau,y-1}\left(1 - p^0_\tau(y-1-y_0)\right) \tag{8}$$

The new vintage replaces retiring stock and supplies stock growth. Replacement of a retiring technology i is allocated to purchased technologies j by the sales share matrix $\sigma_{j,i,y}$, whose columns sum to one (Section 4.4). Growth is first allocated across the retiring technologies in proportion to $R_{i,y}L_i$, which approximates the composition of the existing stock, and then passed through the same matrix:

$$N_{j,y,y} = \sum_i \sigma_{j,i,y}\left(\hat{R}_y \frac{R_{i,y}}{\sum_k R_{k,y}} + G_y \frac{R_{i,y}L_i}{\sum_k R_{k,y}L_k}\right) \tag{9}$$

where $\hat{R}_y = \min\left(\sum_k R_{k,y},\ \sum_k R_{k,y} + \Delta N_y\right)$ is the replacement of retiring stock and $G_y = \max(\Delta N_y, 0)$ is stock growth.

If no stock retires in a year, growth is allocated by the sales shares or by the composition of the prior year's stock. If the total stock falls by more than retirements, the additional reduction is taken as early retirement. Sales are the new vintage grossed up for retirement within the year of purchase:

$$Q_{j,y} = \frac{N_{j,y,y}}{p_j(0)} \tag{10}$$

Sales are reported as replacement sales, up to the natural retirements of each technology, and new sales, the remainder.

4.4 Sales shares

The element $\sigma_{j,i,y}$ is the share of the replacement of a retiring technology i purchased as technology j, for vintage y. A share that does not depend on i is a share of all sales. The matrix is assembled in three steps.

Reference shares. Where a technology has a reference sales share input $s^{\text{ref}}_{j,y}$, it is used for every retiring technology. The remaining share is filled by technologies without an input in proportion to their stock divided by their lifetime, so that equipment is replaced in proportion to the existing stock composition:

$$\sigma^{\text{ref}}_{j,i,y} = \begin{cases} s^{\text{ref}}_{j,y} & j \in \mathcal{J}^{\text{ref}} \\ \left(1 - \sum_{k \in \mathcal{J}^{\text{ref}}} s^{\text{ref}}_{k,y}\right) \dfrac{N_j / L_j}{\sum_{k \notin \mathcal{J}^{\text{ref}}} N_k / L_k} & j \notin \mathcal{J}^{\text{ref}} \end{cases} \tag{11}$$

where N_j is the technology stock in the last year at or before the current year in which technology stocks are specified, or the initial stock. Reference shares that sum to more than one are scaled down. Where the logit choice of Section 4.6 is active, it replaces (11) for the technologies it covers.

Sales share measures without a replaced technology set the share of technology j directly. A technology's share is the sum of its measures, and is never below its reference share. The reference shares of the remaining technologies are scaled to fill the share the measures leave.

Sales share measures with a replaced technology r move a share $s_{j,y}$ of the sales of r to technology j:

$$\Delta\sigma_{j,i,y} = s_{j,y}\,\sigma_{r,i,y}, \qquad \Delta\sigma_{r,i,y} = -s_{j,y}\,\sigma_{r,i,y} \tag{12}$$

If the measures on one replaced technology together move more than all of its sales, they are scaled down to move all of it.

4.5 Specified stocks, linked technologies, and early retirement

A technology stock can be specified in a year, either by historical data up to the current year or by a stock measure. The sales of a specified technology are then set to reach the specified stock, $Q_{\tau,y} = N^{\text{spec}}_{\tau,y} - (N_{\tau,y-1} - R_{\tau,y})$, and the remaining sales are allocated to the other technologies by (9). If a specified stock is below what natural retirement leaves, or specified sales exceed what retirements and growth allow, EP retires additional stock early by advancing the retiring technologies along their survival curves until the required quantity has retired.

The initial stock is taken from the first year at or before the current year in which the specified technology stocks account for between 70% and 143% of the total stock, scaled to the total. If no such year exists, the initial composition is the steady state implied by the reference sales shares. Where there are no reference sales shares either, the initial stock is divided equally among technologies.

A technology can be linked to a technology in another subsector with a stock ratio. The heating and cooling functions of a heat pump are sold in different subsectors, and the cooling stock is set equal to the heating stock times the ratio. The linked stock enters the dependent subsector as a specified stock, and the dependent subsector is calculated after its precursor.

4.6 Logit technology choice

For subsectors with logit parameters, EP can replace the reference sales shares with a mixture of multinomial logits over the total cost of ownership. Purchasers are divided into planning-horizon bins b with horizon H_b and weight w_b. For each bin the cost of technology j is its upfront cost plus annual operating cost over the horizon,

$$C_{j,y,b} = K_{j,y} + O_{j,y} H_b, \qquad O_{j,y} = F_{j,y} + \kappa_{j,y}\left[u_j\, e^{\text{main}}_{j,y}\, \pi_{f^{\text{main}}_j,y} + (1-u_j)\, e^{\text{aux}}_{j,y}\, \pi_{f^{\text{aux}}_j,y}\right] \tag{13}$$

where K is capital plus installation cost less incentives, F is fixed operation and maintenance cost, κ is annual service per unit of stock, and π_f is an exogenous fuel price input. The share of technology j is

$$s_{j,y} = \sum_b w_b \frac{\exp\left(A_{j,y} - \beta_y C_{j,y,b}\right)}{\sum_k \exp\left(A_{k,y} - \beta_y C_{k,y,b}\right)} \tag{14}$$

where β_y is the cost coefficient and $A_{j,y}$ an alternative-specific constant. The constants are calibrated so that (14) reproduces observed sales shares in a calibration year, or are read from a previous calibration. An optional feedback term adds $\gamma \ln(s_{j,y-1}/s_{j,y_c})$ to $A_{j,y}$, which raises the attractiveness of technologies whose share has grown since the calibration year y_c. Shares can be capped by technology-specific maximum shares, with the excess redistributed to the uncapped technologies. Technologies excluded from the logit keep their reference shares, and the logit shares fill the remainder. The logit shares are the same for every retiring technology, and sales share measures still override them. A zero cost coefficient for a zone switches the logit off there.

5 Technology parameters

5.1 Efficiency

Each technology has a main efficiency, by vintage, for its main final energy type. Efficiency is entered as a ratio of a numerator and denominator unit with a flag that states whether the numerator is service or energy (for example Btu of heat per Btu of fuel, or miles per gallon). EP converts every efficiency to energy intensity e, energy per unit of service, in the calculation energy unit. An efficiency can be entered relative to a reference technology, as a multiplier or an addition, and can change with age at a linear or exponential rate.

A technology that uses two final energy types (a plug-in hybrid vehicle, a hybrid heat pump with combustion backup, a solar water heater with electric backup) has an auxiliary efficiency for its second final energy type and a utility factor u_τ, the share of its service delivered with the main type. EP splits the stock of such a technology between the two energy types in proportions u_τ and $1-u_\tau$ and applies the main and auxiliary intensities to the two parts (equation 17). A technology without a utility factor input has $u_\tau = 1$.

A technology can also have parasitic energy, a standing energy use per unit of stock per year that does not depend on service (for example the standby electricity of a gas furnace).

5.2 Service demand modifiers

Technologies in the same subsector can deliver different amounts of service per unit of stock: newer vehicles travel more miles than older ones, and a room air conditioner serves less of a dwelling than a central system. The service demand modifier $m_{\tau,v,y}$ is the ratio of a technology's share of service demand to its share of the stock. Where service or energy demand is entered by technology or by final energy type, EP calculates the modifier as that ratio in the last data year at or before the current year and holds it for later years. A technology-specific modifier input, by vintage and age, replaces the calculated value. Modifiers are normalized so that the stock as a whole delivers the subsector's service demand:

$$\tilde{m}_{\tau,v,y} = \frac{m_{\tau,v,y}}{\sum_{\tau',v'} \nu_{\tau',v',y}\, m_{\tau',v',y}}, \qquad \nu_{\tau,v,y} = \frac{N_{\tau,v,y}}{\sum_{\tau',v'} N_{\tau',v',y}} \tag{15}$$

5.3 Costs

Four cost classes are held for each technology and vintage: capital cost, installation cost, fuel-switching cost, and fixed operation and maintenance (O&M) cost. Capital and installation cost can differ between new purchases (stock growth) and replacement purchases; if only one is entered it is used for both. Costs can be entered as absolute values or relative to a reference technology. Incentives are entered as an absolute value or as a fraction of capital, installation, or fuel-switching cost; where both are entered, the lesser or the greater applies, as the input specifies. Incentives are subtracted from cost.

Fuel-switching cost is a one-time cost paid when a purchase changes the final energy a unit of stock uses, such as a home charger with a first electric vehicle. EP charges it on the sales of each final energy type in excess of retirements of the same type, allocated to technologies in proportion to their sales.

Capital, installation, and fuel-switching costs are levelized over the technology's book life L_τ (its mean lifetime, or the mean of its minimum and maximum) at a real rate r equal to the subsector's cost of capital less the inflation rate:

$$c^{\text{lev}}_{\tau,v} = C_{\tau,v}\, \frac{r}{1-(1+r)^{-L_\tau}} \tag{16}$$

EP reports two cost series. Levelized cost in year y is $c^{\text{lev}}_{\tau,v}$ times the sales of each vintage still within its book life, plus the initial stock, which is charged on a straight-line decline. Annual spending is $C_{\tau,y} Q_{\tau,y}$, paid in the year of sale. Fixed O&M is charged on the surviving stock of each vintage in every year in both series.

6 Energy demand

6.1 Subsectors with a technology stock

Energy demand by final energy type is service demand times the energy intensity of the stock, with each technology and vintage weighted by its share of the stock and its normalized service demand modifier, plus parasitic energy:

$$E_{f,y} = SD_y \sum_{\tau,v} \nu_{\tau,v,y}\, \tilde{m}_{\tau,v,y} \Big[u_\tau\, e^{\text{main}}_{\tau,v,y}\, 1(f = f^{\text{main}}_\tau) + (1-u_\tau)\, e^{\text{aux}}_{\tau,v,y}\, 1(f = f^{\text{aux}}_\tau) \Big] + \sum_{\tau,v} N_{\tau,v,y}\, P_{\tau,v,f} \tag{17}$$

where $P_{\tau,v,f}$ is parasitic energy per unit of stock. Equation (17) is evaluated for each rollover group and summed to the zone.

6.2 Subsectors without a technology stock

In a service and efficiency subsector, energy demand is service demand times the energy per unit of service, by final energy type. In a service and energy subsector and an energy subsector, energy demand by final energy type is projected with (1); in the service and energy type, service demand is one of the drivers. Energy efficiency and fuel-switching measures are then applied (Section 7).

6.3 Calibration to energy benchmarks

Where benchmark final energy demand is supplied for the current year, EP can calculate the ratio of the benchmark to modeled demand by sector, final energy type, and region of the benchmark's geography, map it to the primary geography, and multiply demand in all years by that ratio. The step is a run setting.

7 Measures

A measure is an optional data set that changes the projection of one subsector. A run activates measures by naming them. The calculation without measures is the baseline against which each measure is applied. Five kinds of measure exist:

Measure	Applies to	Effect
Sales share	Subsectors with a technology stock	Changes σ by (11) and (12)
Specified stock	Subsectors with a technology stock	Fixes a technology stock in given years (Section 4.5)
Service demand	Any subsector with service demand	Reduces service demand
Energy efficiency	Subsectors without a technology stock	Reduces energy demand of a final energy type
Fuel switching	Subsectors without a technology stock	Moves energy demand from one final energy type to another

Measure trajectories are projected over the years with the methods of Section 1.3; S-shaped adoption paths use logistic interpolation.

Service demand measures reduce service demand by a fraction of the baseline or by an absolute amount. If the measures in a subsector together save more than the service demand, all are scaled down in proportion.

Energy efficiency measures save a fraction $\theta_{m,y}$ of the energy demand of the final energy types and indices they target, or an absolute amount. Savings are capped at the energy available and scaled down in proportion where they exceed it:

$$\Sigma^{\text{ee}}_{m,f,y} = \theta_{m,y}\, E_{f,y}, \qquad E'_{f,y} = E_{f,y} - \sum_{m} \Sigma^{\text{ee}}_{m,f,y} \tag{18}$$

Fuel-switching measures are applied after efficiency measures. A measure m switches a share $\varphi_{m,y}$ of the demand for final energy type f to type f'. The energy of f' required per unit of f displaced, ι_m, represents the difference in efficiency between the two (for example a gas-fired and an electric process heater):

$$\Sigma^{\text{fs}}_{m,y} = \varphi_{m,y}\, E'_{f,y}, \qquad E''_{f,y} = E'_{f,y} - \sum_{m:f} \Sigma^{\text{fs}}_{m,y}, \qquad E''_{f',y} = E'_{f',y} + \sum_{m:f\to f'} \iota_m\, \Sigma^{\text{fs}}_{m,y} \tag{19}$$

If the measures switching away from one final energy type together exceed its demand, each of them, with its additions, is scaled down by the same factor.

Measure costs. Efficiency, fuel-switching, and service demand measure costs are entered as a levelized cost per unit of energy or service saved. Each measure has a lifetime. EP treats the savings of each measure as a measure stock and rolls it over with (7) to (10) using the measure's lifetime, which gives the new and replacement savings purchased each year. Levelized cost is the levelized unit cost times the measure stock; annual spending is the unit cost times the savings purchased.

8 Load shapes

8.1 Shapes

A load shape is an hourly profile over the weather years of the run. Two kinds exist. A time-slice shape gives a repeating pattern indexed by combinations of month, week, day of year, day of week, day type (workday or non-workday), and hour, and EP expands it to every hour of the weather years. A weather-date shape gives values for each hour of specific calendar years, so that weather-dependent loads such as heating, cooling, and vehicle charging correspond to the weather of the year. A shape can be indexed by final energy type, which lets a hybrid heat pump carry one profile for its electricity and another for its backup fuel, and by distribution feeder.

Each shape is entered on a geography and in local time. EP maps it through time zones to a common time zone and then to the primary geography with (2), using the shape's map key. Shapes are normalized so that each sums to W, the number of weather years, over the hours h; only the pattern of a shape affects results.

8.2 Assignment and aggregation

Each part of a subsector's energy demand is assigned a shape in this order: the technology's own shape, the subsector's shape, the sector's shape, and the default shape of the final energy type. The hourly load of a subsector in zone z is

$$\ell_{z,h,y} = \sum_{\tau \in \mathcal{T}^{\mathrm{s}}} E_{\tau,z,y}\, s^{\tau}_{z,h} + \left(E_{z,y} - \sum_{\tau \in \mathcal{T}^{\mathrm{s}}} E_{\tau,z,y}\right) s^{\mathrm{sub}}_{z,h} \tag{20}$$

where $\mathcal{T}^{\mathrm{s}}$ is the set of technologies with their own shape and s^{sub} is the shape assigned to the rest of the subsector. Load is divided among distribution feeder types (residential, commercial, industrial) by a feeder allocation by subsector and year. Sector and system load are the sums of subsector loads.

8.3 Reconciliation to system load

The bottom-up electricity profile is reconciled to observed system load in a base year y_b, the later of the mean weather year and the current year. EP forms a top-down profile, base-year electricity demand by zone times the default electricity shape, and takes the ratio of the top-down profile to the bottom-up profile in each zone and hour:

$$\rho_{z,h} = \frac{E^{\mathrm{el}}_{z,y_b}\, s^{\mathrm{sys}}_{z,h}}{\sum_{\mathrm{sub}} \ell^{\mathrm{el}}_{z,h,y_b}} \tag{21}$$

The factor is applied to the share of each subsector's electricity load that existed in the base year, and load growth beyond the base year keeps its bottom-up shape. With $\psi_{z,y} = \min\left(1, E_{z,y_b}/E_{z,y}\right)$ the base-year share of the subsector's demand,

$$\ell'_{z,h,y} = \ell_{z,h,y}\left(\rho_{z,h}\,\psi_{z,y} + 1 - \psi_{z,y}\right) \tag{22}$$

after which the profile is rescaled so that its sum over hours, divided by W, equals the subsector's annual electricity demand. The reconciliation preserves the change in load shape that comes from changes in the technology mix and adjusts the base-year profile to observed load.

9 Outputs and export to RIO

9.1 Outputs

EP writes, by sector, subsector, zone, year, and where applicable technology, vintage, and final energy type: drivers; service demand; stocks; sales; final energy demand; levelized cost and annual spending by cost class; measure savings and costs; hourly profiles by subsector and final energy type for selected years; and the reconciliation factors of Sections 6.3 and 8.3. Each run of EP is a case, defined by the data selected and the measures activated. Results from several cases are concatenated, each row labelled with its case.

9.2 Tables passed to RIO

EP writes a set of tables in the format of the RIO database for the RIO optimization years. Each row carries the name of the EP case that produced it. RIO reads EP final energy through a mapping from each subsector and final energy type to a RIO energy carrier; the mapping can drop a final energy type, or reverse its sign, for a subsector.

Table	Contents	Subsectors
EXO_DEMAND	Annual final energy demand by subsector, zone, and RIO carrier (electricity in GWh)	All
ShapeData, SHAPE_META	Hourly profiles for the weather years, for each RIO year: total electricity load by zone and feeder, profiles of other final energy types named in the run, per-subsector electricity profiles of optimizable and flexible-load subsectors, and technology-specific shapes of optimizable subsectors	As stated
DEMAND_DISPATCH_FEEDER_ALLOCATION	Feeder shares by subsector, zone, and year	All
DEMAND_LEVELIZED_COSTS, DEMAND_SPEND	Levelized demand-side cost and annual spending by subsector, zone, and cost class	All
DEMAND_EE_SAVINGS, DEMAND_FS_SAVINGS, DEMAND_FS_ADDITIONS	Savings and additions of each efficiency and fuel-switching measure by subsector, zone, carrier, and year	Subsectors with these measures
DEMAND_EE_COSTS, DEMAND_FS_COSTS	Levelized cost per unit of energy saved, by measure and vintage	Subsectors with these measures
DEMAND_TECH_MAIN	Technologies, lifetimes, stock units, linked technologies and ratios, and shapes	Optimizable
DEMAND_TECH_SALES_EP, _MIN, _MAX	Sales by technology, vintage, zone, and other index	Optimizable

Table	Contents	Subsectors
DEMAND_TECH_SERVICE_DEMAND	Service demand per unit of total stock, $SD_y / \sum_{\tau,v} N_{\tau,v,y}$, by zone and year	Optimizable
DEMAND_TECH_CAPITAL_COST, DEMAND_TECH_INCENTIVE, DEMAND_TECH_FIXED_OM	Capital plus installation cost of new purchases, incentives, and fixed O&M, by technology and vintage	Optimizable
DEMAND_TECH_EFFICIENCY	Energy intensity by technology, vintage, and carrier, split into main and auxiliary rows by the utility factor	Optimizable
DEMAND_TECH_SERVICE_DEMAND_MODIFIER	Service demand modifier by technology, vintage, and age	Optimizable

Sales and efficiency rows cover vintages from the first RIO year, and efficiency is evaluated in the last RIO year. A measure that spans several end uses with different costs, or, for fuel switching, different energy intensities, is split into one measure per end use, so that RIO can adopt each separately.

The subsectors marked optimizable are those listed as optimizable in the EP run configuration. Only subsectors with a technology stock can be listed. The flexible-load subsectors, whose load RIO can shift in time, are listed in the same way. The list of optimizable subsectors, together with RIO's demand optimization setting, decides which subsectors RIO treats as fixed.

9.3 Mode (a): demand optimization off

When demand-side technology optimization is off in RIO, EP's sales shares determine the technology mix of every subsector. RIO reads EXO_DEMAND and the hourly profiles of one EP case and treats the resulting final energy demand by carrier, zone, year, and hour as fixed. Every subsector is passed as fixed demand. The measures applied are those the EP case activated, and their effect is already in EXO_DEMAND. The demand-side cost tables enter RIO's cost accounting.

9.4 Mode (b): demand optimization on

When demand-side technology optimization is on, RIO chooses sales of the technologies in the optimizable subsectors, within an envelope taken from EP, jointly with supply. EP supplies for these subsectors the technology list, sales, service demand per unit of stock, cost, efficiency, and service demand modifiers of Section 9.2, and RIO conserves each subsector's service demand while changing its technology mix. RIO's formulation is given in the *RIO Model* document. RIO does not advance stocks or project service demand; those come from EP.

The baseline energy of the optimizable subsectors stays in EXO_DEMAND. RIO adds to it the energy implied by the difference between its chosen sales and the EP baseline sales of the selected case, using the exported efficiencies, modifiers, and technology shapes. All subsectors that are not optimizable are passed as fixed demand, exactly as in mode (a). RIO can also treat the adoption of the exported efficiency and fuel-switching measures as a decision, using their savings, additions, and costs; this is a separate RIO setting.

Within one case the three sales tables are identical. The envelope comes from combining cases. RIO reads each EP table from a case named in its own run configuration, and sets the lower bound on sales to the smaller and the upper bound to the larger of the minimum and maximum tables it reads. Two EP cases that differ only in their measures span the range between them: the case with measures supplies the baseline and the maximum, and the case without

them supplies the minimum. Where the measures move sales toward a technology, RIO can reduce its adoption back toward the case without measures, and the technologies it displaces can increase. Where a single case supplies all three tables, sales are fixed at the EP baseline. The other EP tables, such as EXO_DEMAND and the feeder allocation, are read from whichever case the RIO run names.

References

Williams, J. H., Jones, R. A., Haley, B., Kwok, G., Hargreaves, J., Farbes, J., & Torn, M. S. (2021). Carbon-neutral pathways for the United States. *AGU Advances*, 2(1), e2020AV000284. https://doi.org/10.1029/2020AV000284