WORKING PAPER | VERSION 0.2 | OCTOBER 2026

Activity Levels Across CPAT, SISEPUEDE, CBAM and FTT

One Driver Layer So Four Climate Models Can Be Linked, With the Carbon-Price Response Counted Once

Stephen J. Stretton*

ABSTRACT

Every energy and emissions model starts from activity levels: how much electricity is used, how far people and goods travel, how many vehicles are on the road, how many tonnes of steel and cement are made. Four models used for climate-fiscal work build those levels in incompatible ways. The Climate Policy Assessment Tool grows fuel use with total GDP and prices. SISEPUEDE grows physical activity with income and population but never with prices. The CBAM Policy Tool grows output of steel, aluminum, fertilizer and cement with growth forecasts and a price response. The Future Technology Transformations models take demand as given. This paper maps each model's activity onto ten common activities, sets out one constant-elasticity rule driven by population, income per person and the cost of the energy service, and gives the rules that let the four be linked without counting a carbon price's demand response twice or not at all. Under stylized growth of 1 percent in population and 2 percent in income per person, cars in use grow 64 percent in ten years while travel grows 35 percent, so every link needs a check that car travel stays inside total travel.

Keywords: activity levels, energy demand, model linking, income elasticity, price elasticity, technology diffusion
JEL codes: Q41, Q43, Q47, C53, R41

Working paper, version 0.2, October 2026. Stephen J. Stretton, writing in a personal capacity. The findings, interpretations, and conclusions expressed in this paper are entirely those of the author. They do not necessarily represent the views of the World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent. A short version for policy makers is Climate Fiscal Policy Brief No. 6.4. Companion paper to the public CPAT-AI repository, https://github.com/cpmodelai/cpat_ai.[1]

*Writing in a personal capacity.

[1]This paper condenses the author's design note on activity levels across the four models. Every statement about a model was checked against its public code and documentation on 31 October 2026: CPAT-AI (commit 108f976), SISEPUEDE (`jcsyme/sisepuede`, commit 9bf976c) and the FTT standalone models (`cpmodel/FTT_StandAlone`, commit 896a3ce). SISEPUEDE was developed by James Syme, Edmundo Molina-Pérez, Nidhi Kalra and colleagues at RAND and the Inter-American Development Bank; FTT by Jean-François Mercure and colleagues at Cambridge Econometrics and the Universities of Cambridge and Exeter; CPAT by the World Bank and the International Monetary Fund. The paper uses no country results. Errors are the author's own.

1 Executive summary

Four models that a finance ministry might use to assess a carbon price disagree before any policy is applied, because they start from different activity levels. The Climate Policy Assessment Tool (CPAT), built by the World Bank and the International Monetary Fund (IMF), grows fuel use with total GDP and prices. SISEPUEDE, an open model developed by RAND and the Inter-American Development Bank, grows physical activity with income and population but never with prices. The CBAM Policy Tool (paper 6.2) grows output of the goods covered by the European Union's carbon border adjustment from growth forecasts. The Future Technology Transformations (FTT) models decide which technologies meet demand but take its size as given. Linked as they stand, they give a ministry four baselines and a carbon price whose effect on demand is counted once, twice or not at all.

The fix is a common activity layer: ten activities in common units, one growth rule and one set of drivers. Each activity grows with population, income per person and the cost of the energy service it provides, through constant elasticities, from one base year and one forecast of GDP and population. A sector concordance maps each activity to each model's own variables. Each model keeps its own data and takes the layer's growth.

The carbon price must reach demand exactly once. CPAT and the CBAM Policy Tool already respond to prices; SISEPUEDE and FTT do not. So the layer applies its price term only to the activities it passes to SISEPUEDE and FTT. Applying it in both the layer and a model that has its own response doubles the response in logs. Inside CPAT the same rule means feeding its engineering road model only the baseline from the layer, because that model's own response to the cost of driving would otherwise be added on top.

Under stylized growth, cars outrun travel, and the link needs a check that catches it. With population growing 1 percent a year and income per person 2 percent, cars in use grow 64 percent in ten years and passenger travel 35 percent, so the car share of travel rises by about a quarter. Held at constant elasticities, car travel would exceed all travel within about 56 years from a 30 percent share. A saturating ownership curve such as CPAT's does not prevent this: with CPAT's curve the crossing still comes in year 54. The layer therefore checks every year that cars times kilometers per car times occupancy stay below total passenger travel, and lets kilometers per car absorb the difference.

Where the models' own values fail against the evidence, the layer replaces them. SISEPUEDE's example data grow regional travel with an income elasticity of 2.64 with no ceiling, and move every industrial elasticity to 1 by 2050, against evidence that material demand saturates. The layer uses an elasticity of 1 for total travel, in line with evidence that mobility rises roughly in proportion with income, and lets industrial elasticities fall. CPAT's income elasticity of electricity, 0.8 before its income adjustment, rises to 0.92 at 3,000 dollars of GDP per person and falls to 0.48 above about 22,000 dollars.

The data gaps bind as hard as the method. FTT's standalone transport model has no car fleet data for 21 of its 71 regions, including Nigeria, South Africa, Pakistan and Kenya. SISEPUEDE's public repository holds one example country; country calibrations are not public. CPAT's engineering road model and its co-benefits module use different driving responses unless the user selects the engineering model. The output responses of heavy industrial goods to a carbon price are judgment values.

What to do. Give every model one GDP and population forecast. Grow activities with one rule from one base year, and splice them onto each model's own data. Apply the price response once, in whichever model has it. Check that cars fit inside travel. Use the gap

between CPAT's fuel-based baseline and the layer's activity-based one as a test of both. The policy brief sets out what this means for finance ministries.

2 Introduction

A finance ministry that wants to know what a carbon price does to its budget, its exporters and its power system now has several open or public models to ask, and they rarely agree on the most basic input: how much activity there will be. CPAT projects fuel use for more than 180 countries from total GDP and prices (Black et al. 2023; CPAT Team 2024). SISEPUEDE projects passenger-kilometers, tonne-kilometers, household energy demand and tonnes of industrial output from population and income (Kalra et al. 2023; Syme et al. 2026). The CBAM Policy Tool, described in paper 6.2 (Stretton 2026d), projects output of the goods covered by the European Union's Carbon Border Adjustment Mechanism (European Parliament and Council of the European Union 2023). The FTT models simulate which technologies win market share in power, cars, road freight and heating, but take the size of each market as an input (Mercure 2012; Mercure et al. 2018; Future Technology Transformations developers 2026). Each is good at its job. Linked naively, they give a ministry four baselines built on different growth paths, and a carbon price whose demand response is counted in one model, two, or none.

The problem is practical. FTT's standalone baseline demand is a fixed input, which the coupled E3ME-FTT macroeconomic model "informs" (Future Technology Transformations developers 2026), not a projection from the IMF forecasts that a CPAT or SISEPUEDE run uses. SISEPUEDE has no price response, so a carbon tax changes its emissions only if the analyst writes the change into a scenario. CPAT has a price response for every fuel and sector, but its road transport is described three times, in a fuel equation, an optional engineering fleet model and a co-benefits module, and the co-benefits always follow the fleet model's driving, whichever road model drives fuel use (CPAT-AI Project 2026c, no. I100).

This paper asks how the activity levels in CPAT, SISEPUEDE, the CBAM Policy Tool and FTT can be put on one footing so that the four can be linked. The answer is a common activity layer: one rule that grows each of ten activities with population, income per person and the cost of the energy service it provides, driven by one set of forecasts, set from one base year, and mapped to each model through a sector concordance. The price term is applied only where the receiving model has no price response of its own. Three results organize the paper. The population exponent matters only for activities whose income elasticity is not 1, but there it adds or removes up to 0.4 percentage points of growth a year (Proposition 1). A price response must sit in exactly one place for each activity, and one place in CPAT is easy to miss (Proposition 2). And a car stock that grows faster with income than travel does must eventually carry more passenger-kilometers than there are, even with a saturating ownership curve, so the vehicle chain needs a per-year check (Proposition 3).

The contribution is a concordance and a set of linking rules, not a new model. The idea of giving several models one set of drivers is the approach of the Shared Socioeconomic Pathways, which offer integrated assessment models "a set of common assumptions" for population, economic growth and urbanization (Riahi et al. 2017). The idea of passing one model's outputs to another by rule rather than solving them together is the "softlinking" of energy-economy models (Wene 1996); hard-linked set-ups, which iterate between models until they converge, reach the same equilibrium in tested cases, although different ways of coupling can give different solutions (Helgesen and Tomasgard 2018). This paper applies both ideas to four tools that a ministry may actually run, with the demand evidence each link needs: income elasticities of energy demand (Burke and Csereklyei 2016; Gertler et al. 2016), price elasticities (Labandeira et al. 2017; Goodwin et al. 2004; Jong et al. 2010), vehicle

ownership (Dargay and Gately 1999; Dargay et al. 2007), travel (Schäfer and Victor 2000) and the role of population (Liddle 2014). It complements paper 5.1, which states CPAT's demand system (Stretton 2026e), and paper 5.4, which checks its elasticities against the evidence (Stretton 2026a); paper 6.1 uses the same demand system to measure what carbon prices actually cut (Stretton 2026b).

3 How each model represents activity

3.1 CPAT: fuel use driven by total GDP and prices

CPAT does not model activity directly; it models fuel use, with total GDP standing in for activity. For fuel f in sector s, fuel use F in year t relative to the year before is

$$\frac{F_{sf,t}}{F_{sf,t-1}} = \left(\frac{1}{1+\alpha_{sf}+\alpha^{+}_{s,t}}\right)^{1+\epsilon_U} \Psi_t \left(\frac{Y_t}{Y_{t-1}}\right)^{\epsilon_Y inc_t} \left(\frac{p_{sf,t}}{p_{sf,t-1}}\right)^{\epsilon_U} \left(\frac{p_{sf,t}+sh_{sf,t}}{p_{sf,t-1}+sh_{sf,t-1}}\right)^{\epsilon_F(1+\epsilon_U)}, \quad (1)$$

where α is an autonomous efficiency gain of 0.5 or 1 percent a year, α^{+} any additional efficiency gain assigned to a policy, Ψ_t a demand shock (1 in normal years), Y total real GDP, ϵ_Y the income elasticity, inc_t an income-level factor, p the retail price including taxes, ϵ_U the usage elasticity (how hard existing equipment is used), ϵ_F the efficiency elasticity (what equipment is bought) and sh the shadow price of non-price policies (CPAT Team 2024; CPAT-AI Project 2026b, equation M3.1). Paper 5.1 explains each term (Stretton 2026e). The equations here are those of the rebuilt model, which found 231 bugs in the legacy workbook, and of the corrected edition of the documentation published on 31 October 2026 (CPAT-AI Project 2026c, 2026a).

Income elasticities are set by fuel and sector and then bent by income level. The legacy table holds 32 income elasticities, for eight fuel groups and four sectors, the same for every income group; for electricity they are 0.8 in homes and industry, 1.1 in services and 1.2 in transport (CPAT-AI Project 2026b). The CPAT team chose them so that weighted averages stay close to the sectoral estimates of Burke and Csereklyei (2016), and multiplies them by a factor of income per person that it inferred from Gertler et al. (2016), who show that household purchases of energy-using assets rise nonlinearly with income (CPAT Team 2024). The factor is a quartic in log GDP per person, normalized by the aggregate elasticity of 0.74 and capped at a log of 10, about 22,000 dollars. Recomputed from its constants, it is 1.04 at 1,000 dollars, peaks at 1.17 near 2,250 dollars, is 1.15 at 3,000 and falls to 0.61 at the cap. So CPAT's income elasticity of household electricity is 0.92 at 3,000 dollars per person and 0.48 in rich countries.

CPAT does not separate population from income per person. Because Y is total GDP, a country whose population grows fast while income per person stagnates gets the same fuel growth as one with the reverse (CPAT Team 2024). The CPAT-AI project tracks this as issue I71 (CPAT-AI Project 2026c). In the notation of this paper, CPAT sets the population exponent equal to the income elasticity for every fuel.

Power demand is built bottom-up from the same kind of equation. Electricity demand in each sector grows with GDP, the efficiency trend and last year's blended residential and industrial electricity price, with the combined price elasticity $\epsilon_U + \epsilon_F + \epsilon_U\epsilon_F$ (CPAT-AI Project 2026b, equation M4.1). For homes and industry in low- and middle-income countries the legacy values are -0.3 and -0.2, so the combined elasticity is -0.44. Generation then adds net exports, losses, statistical differences and own use, each held at its base-year share of final demand (CPAT-AI Project 2026b).

Road transport is described three times, and the three do not agree by default. The default road model is equation 1 itself. An optional engineering fleet model, legacy CPAT's "technoeconomic" road model, projects vehicles per 1,000 people along an S-shaped Gompertz curve in GDP per person, lets buyers choose electric vehicles by a logit on the total cost of ownership, and derives fuel use from vehicles, distance and fuel economy; the user switches it on per scenario (CPAT-AI Project 2026b). A third block, the co-benefits module, values congestion, accidents and road damage per vehicle-kilometer (CPAT Team 2024). That module takes its vehicle-kilometers from the fleet model whatever road model the user selects (CPAT-AI Project 2026c, no. I100). So in a default run, fuel use and emissions follow the fuel equation while the co-benefits follow the fleet model's driving, and the change in driving that is valued need not match the modeled change in fuel use. Industry is driven by GDP, not tonnes, and no fuel substitutes for another within an industrial subsector (Stretton 2026e).

3.2 SISEPUEDE: physical activity driven by income, with no prices

SISEPUEDE counts physical activity and multiplies it by energy intensity. SISEPUEDE (SImulating SEctoral Pathways and Uncertainty Exploration for DEcarbonization; the IDB report writes it SiSePuede) is an open Python and Julia framework for emissions accounting by IPCC sector, built for exploring decarbonization strategies across many uncertain futures (Kalra et al. 2023; Syme et al. 2026). Every income-driven demand grows with a factor

$$G_t(\varepsilon, x) = \prod_{s=0}^{t-1} \left(\, 1 + \varepsilon_s \, g_s(x) \, \right), \tag{2}$$

where $g_s(x)$ is the growth rate of driver x in year s and ε_s the elasticity, which may vary by year (the code's "standard" elasticity; a "log" form, $(1+g)^{\varepsilon}$, is also available) (Syme et al. 2026). Passenger-kilometers per person grow with GDP per person and are multiplied by population, in two categories ("private and public" travel and "regional" travel). Megatonne-kilometers of freight grow with total GDP. Household energy demand for heat and for appliances grows per household with GDP per person and is multiplied by the number of households. Production of each industry (metals, cement, chemicals and others) grows with total GDP. Electricity demand is the sum of the electricity used by the end-use sectors, grossed up for transmission losses and exports, and is passed to an embedded least-cost model, NemoMod, developed by the Stockholm Environment Institute (Syme et al. 2026). Mode shares, fuel shares and energy per unit are scenario inputs.

Prices never change demand in SISEPUEDE. No price elasticity appears anywhere in the code. Fuel prices enter the costs of fuel supply and the cost-benefit accounts, not demand. Electrification or mode shift happens because a scenario says so. That is a deliberate design for finding strategies that hold up across many uncertain futures (Kalra et al. 2023), but it means SISEPUEDE cannot say on its own what a carbon tax does to demand.

The public example shows how the elasticities are set. The repository's example country (Costa Rica, from 2015) holds the passenger-travel elasticities at 1.12 for private and public travel and 2.64 for regional travel, and the freight elasticity at 1.0, through 2050. The industrial and household elasticities start from short calibrations and move in a straight line to 1.0 by 2050. In 2025 they are 0.62 for cement, 0.75 for metals, -0.25 for chemicals and 0.08 for household heat (Syme et al. 2026). A negative elasticity for chemical output, and convergence of every material to an elasticity of 1, are the kind of values a common layer should replace.

3.3 The CBAM Policy Tool: output of covered goods

The CBAM Policy Tool grows output of each covered good from a recent base year and cuts it along a constant-elasticity curve when a domestic carbon price raises its cost. For product i,

$$Q_{i,t} = Q^{0}_{i,t}\,\left(\,1+\delta_{i,t}\right)^{\eta_i}, \qquad \delta_{i,t} = \frac{(1-\theta)\,\tau_t\,e_i}{p_i}, \tag{3}$$

where Q^0 grows with IMF growth projections scaled by an output elasticity, δ is the carbon cost share, τ the domestic carbon price, e_i emissions per tonne, θ the share of the carbon cost rebated per unit of output, p_i the product price and η_i a long-run output elasticity; a partial-adjustment block scales the response for a price in force only a few years (Stretton 2026d). In the tool's illustrative runs the author uses judgment values of -0.4 for steel and nitrogen fertilizers, -0.5 for primary aluminum and -0.1 for clinker; the evidence on output elasticities of heavy industrial goods is thin, and a Brazilian estimate of -0.14 for steel is one of few (Schmidt and Lima 2006). The tool covers goods at the product level (steel by process route, gray clinker, ammonia, urea, ammonium nitrate, primary aluminum), which is finer than either CPAT or SISEPUEDE.

3.4 FTT: technology shares of a demand taken as given

FTT decides how demand is met, not how much there is. In each FTT model, heterogeneous agents compare technologies pair by pair on a distribution of perceived costs, and shares move by logistic substitution at a pace set by lifetimes and the speed of capacity expansion; the result is the S-shaped diffusion seen in practice rather than a least-cost mix (Mercure 2012; Mercure et al. 2016; Future Technology Transformations developers 2026). Output of technology i is its share of total demand,

$$G_{i,t} = S_{i,t}\,D_t, \tag{4}$$

and D_t comes from outside. Inside the full E3ME-FTT model, the macroeconometric model E3ME supplies it. In the open standalone Python version, which runs power, cars, road freight and household heat for 71 countries and regions (an industrial heat module ships without baseline data), it comes from input files (Future Technology Transformations developers 2026; Knobloch et al. 2019). Power reads an electricity demand that equals total generation in the base year, so it includes losses and own use and is net of trade (checked for five regions in 2022). Cars read the fleet in use and kilometers per car. Freight reads vans and trucks in use by size, and computes tonne-kilometers from vehicles, kilometers per vehicle and load. Heat reads useful heat demand for homes. FTT never reads passenger-kilometers, and it derives tonne-kilometers rather than reading them.

FTT's data end in different years and have gaps. Power data run to 2022 (generation), transport to 2022, heat to 2020 and freight to 2023 (Future Technology Transformations developers 2026). The car fleet input is zero in 2022 for 21 of the 71 regions, including Nigeria, South Africa, Pakistan and Kenya (checked at commit 896a3ce). The standalone package has no steel model; FTT:Steel runs inside E3ME-FTT (Vercoulen et al. 2023). In policy runs FTT adds to power demand the change in each sector's electricity use against its own baseline run (Future Technology Transformations developers 2026).

4 The concordance

The four models share ten activities once units are aligned, and Table 1 maps each one. The common layer uses the unit the most demanding model needs: vehicles for road transport because FTT needs them, useful heat for homes because SISEPUEDE and FTT both use it, tonnes by product for industry because the CBAM Policy Tool needs them, and terawatt-hours of generation for power because FTT reads generation. Each other model then converts from that unit with its own intensity.

Table 1. The four models share ten activities once units are aligned; only CPAT and the CBAM Policy Tool let prices change demand. "Implicit" means the model carries the activity only inside fuel use. Source: CPAT Team (2024); CPAT-AI Project (2026b); Syme et al. (2026); Future Technology Transformations developers (2026); Stretton (2026d).

Activity (common unit)	CPAT	SISEPUEDE	CBAM Policy Tool	FTT (standalone)
Electricity (TWh generated)	Sector demand (M4.1) plus losses, own use, net exports	End-use electricity, grossed up for losses	Not modeled	Power: demand input, equal to generation
Passenger travel (passenger-km)	Implicit in road and aviation fuel	Passenger-km per person × population, two categories	Not modeled	Not read
Cars in use (vehicles)	Fleet model: Gompertz ownership in income	Not modeled (vehicle-km from passenger-km)	Not modeled	Transport: fleet in use (input)
Km per car (km a year)	Fleet model: income and cost per km	Implicit in vehicle-km	Not modeled	Transport: km per car (input)
Road freight (tonne-km)	Implicit in diesel use	Tonne-km with total GDP	Not modeled	Derived from vehicles, km and load
Vans and trucks (vehicles)	Fleet model: freight vehicles	Not modeled	Not modeled	Freight: vehicles by size (input)
Home heat, hot water and cooking (useful heat)	Residential fuel use, equation 1	Heat demand per household × households	Not modeled	Heat: useful heat for homes (input)
Steel and aluminum (tonnes)	Industrial fuel use with GDP	Metals output with GDP	Steel by route; primary aluminum; equation 3	Steel only inside E3ME-FTT
Cement and clinker (tonnes)	Industrial fuel use with GDP	Cement output with GDP	Gray clinker; equation 3	Not modeled
Fertilizers (tonnes)	Industrial fuel use with GDP	Chemicals output with GDP	Ammonia, urea, ammonium nitrate; equation 3	Not modeled
Does demand respond to price?	Yes, every fuel and sector	No	Yes, output of covered goods	No (prices move shares, not demand)

Three gaps in the concordance matter for policy. First, outside the CBAM Policy Tool and E3ME-FTT nobody models steelmaking routes: CPAT ties industrial fuel to GDP and SISEPUEDE ties industrial energy to tonnes, but neither distinguishes blast furnaces from electric arc furnaces or hydrogen-based steel, which is where the emission cuts in steel come from. Second, none of the four has a cooling demand of its own, although household purchases of energy-using appliances, air conditioners among them, rise steeply as incomes cross middle-income thresholds (Gertler et al. 2016). Third, the two models that do not respond to price, SISEPUEDE and FTT, are the two that hold the most technology detail, so a carbon price reaches their demand only if the layer supplies it.

CPAT's income elasticity of fuel is the sum of an activity elasticity and an intensity elasticity, which is how the concordance maps between them. Write fuel use as activity times energy per unit of activity, $F = A \cdot e$. Taking logs and differentiating with respect to GDP, $\epsilon_Y = \varepsilon_A + \varepsilon_e$: CPAT's single elasticity bundles how fast activity grows with how fast energy per unit changes. When the layer supplies activity to a model that carries its own intensity (SISEPUEDE, FTT), only ε_A should pass. When CPAT is used as a cross-check on the layer, its ϵ_Y should be compared with $\varepsilon_A + \varepsilon_e$, not with ε_A alone. CPAT is a check on the layer, not a target for it. Electricity is the exception: there the activity is the fuel, so the two elasticities coincide.

5 A common activity layer

5.1 The rule

Each activity grows with population and income per person and responds to the cost of the energy service it provides.

$$A_{j,t} = A_{j,0} \left(\frac{P_t}{P_0} \right)^{\beta_j} \left(\frac{y_t}{y_0} \right)^{\varepsilon_j} \left(\frac{c_{j,t}}{c_{j,t}^{\text{base}}} \right)^{\eta_j}, \tag{5}$$

where $A_{j,t}$ is activity j in year t, P_t population, y_t real GDP per person, $c_{j,t}$ the cost of the energy service and $c_{j,t}^{\text{base}}$ its baseline value. The population exponent β_j is 1 for activities that follow people (travel, households, cars) and equal to ε_j for those that follow the economy (freight, industrial output), because total GDP is population times income per person; kilometers per car, a per-vehicle activity, has $\beta_j = 0$. ε_j is the income elasticity and η_j the price elasticity. In the baseline the last term equals 1, so the income and population terms describe the baseline alone.

The cost of the energy service is the price of energy divided by the efficiency of the equipment. For a service delivered by fuel f with price p_f, including taxes and any carbon price, and equipment efficiency e_f (kilometers per unit of energy, or useful heat per unit of fuel), $c_j = \sum_f s_f \, p_f / e_f$, where s_f is the share of the service delivered by fuel f. For electricity it is the electricity price; for covered goods, the product price plus the net carbon cost per tonne. CPAT supplies the prices; FTT's technology mix in the previous year supplies the efficiency, which avoids a loop. equation 5 is the constant-elasticity form; SISEPUEDE's compounding in equation 2 differs from it by second-order terms, $\ln(1 + \varepsilon g) - \varepsilon \ln(1 + g)$, about 0.005 percentage points a year at an elasticity of 1.2 and growth of 2 percent.

5.2 One parameter set

Table 2 sets one value for each parameter, taken from the evidence where it is good and stated as an assumption where it is not.

Table 2. One parameter set for the common layer. Income elasticities are with respect to GDP per person; the electricity values are CPAT's, with respect to total GDP. Price elasticities are long-run, with respect to the cost of the energy service. "Author's assumption" marks a value chosen for the illustration where no public estimate fits; the CBAM Policy Tool's values are the author's judgment. Car numbers: the straight line approximates CPAT's Gompertz curve with a saturation of 800 vehicles per 1,000, whose elasticity peaks at 2.2 near $4,350 per person and 90 vehicles per 1,000; the line is within 0.25 of it between 30 and 750 vehicles per 1,000. Goodwin et al.'s figures are read in Litman (2026)'s summary. Source: as listed; CPAT-AI Project (2026b); Syme et al. (2026); author's calculation (`2.Source/activity_layer.py`).

Activity	Income elasticity ε_j	Source	Price elasticity η_j	Source
Electricity, homes and industry	0.8 base; 0.92 at $3,000 per person; 0.48 above $22,000	CPAT legacy values and income factor	−0.44	CPAT usage and efficiency margins
Passenger travel (passenger-km)	1.0	Schäfer and Victor (2000)	0 (cars carry the response)	Author's assumption
Road freight (tonne-km)	1.0	SISEPUEDE example data	−0.1	Jong et al. (2010)
Home heat, hot water and cooking	0.4	Author's assumption	−0.2	Author's assumption; Sorrell et al. (2009)
Steel and aluminum	1.2, falling	Author's assumption; Ruijven et al. (2016)	−0.4 steel, −0.5 aluminum	CBAM Policy Tool (judgment)
Cement and clinker	0.6	Author's assumption	−0.1	CBAM Policy Tool (judgment)
Fertilizers	0.9	Author's assumption	−0.4	CBAM Policy Tool (judgment)
Cars in use	2.0 up to 300 per 1,000 people, falling to 0 at 800	Dargay and Gately (1999); CPAT fleet model	−0.25	Goodwin et al. (2004)
Vans and trucks	1.0 (as freight)	SISEPUEDE example data	−0.2	Jong et al. (2010)
Km per car	0.1	CPAT fleet model	−0.05	Goodwin et al. (2004)

The travel and car values rest on the best evidence available. Total mobility rises "nearly in proportion with income" because people spend a roughly fixed share of income and a roughly fixed amount of time on travel (Schäfer and Victor 2000), so the layer uses an income elasticity of 1 for passenger-kilometers per person; SISEPUEDE's example values of 1.12 and 2.64, held to 2050 with no ceiling, would let regional travel grow without bound. Vehicle ownership rises about twice as fast as income at low and middle incomes and the elasticity falls toward zero near saturation (Dargay and Gately 1999), at a saturation level that depends on urbanization and density (Dargay et al. 2007); CPAT's fleet model draws the same shape. A sustained 10 percent rise in the real price of fuel cuts traffic by about 3 percent and vehicle

ownership by about 2.5 percent in the long run Litman (2026), so the response of kilometers per vehicle is about -0.05; the evidence is mostly from high-income countries.

The freight values come from one careful synthesis. Jong et al. (2010) reconcile a wide range of European estimates into a consistent set of best-guess long-run elasticities: a fuel-price elasticity of -0.1 for road freight tonne-kilometers and -0.2 for vehicle-kilometers, the difference being better loading and logistics. Because FTT holds kilometers per vehicle and load fixed, the layer passes the vehicle-kilometer response to vans and trucks in use.

The household and industrial values are assumptions, and the layer says so. SISEPUEDE's own values are short country calibrations that converge to 1 by 2050, and the public example shows how noisy they are (Section 3.2). The layer instead uses round values that let material intensity fall as incomes rise, in line with projections in which steel and cement demand rise quickly for decades and then stabilize (Ruijven et al. 2016). The household heat price elasticity of -0.2 sits inside the bound of Sorrell et al. (2009), who conclude that direct rebound for household energy services in OECD countries should generally be below 30 percent, and below CPAT's residential usage elasticities of -0.3 to -0.5; a ministry should test both ends. Heavy-industry output elasticities are judgment values with little empirical support (Schmidt and Lima 2006).

5.3 Three propositions

Proposition 1 (the population exponent). *Replacing $\beta_j = \varepsilon_j$ by $\beta_j = 1$ changes the growth of activity j by $(1 - \varepsilon_j)\, n$ a year in logs, where n is population growth. The choice matters only when $\varepsilon_j \neq 1$ and population grows.*

The models choose $\beta_j = \varepsilon_j$ for every activity driven by total GDP, and CPAT for every fuel. Cross-country studies that separate population from income per person find population elasticities of emissions and energy near one in cross-sections, but widely varying ones over time; the main review cites the argument that, without indirect effects, population acts only as a scale factor with an elasticity of one (Liddle 2014). With population growing 1 percent a year, setting $\beta_j = 1$ adds 0.4 percentage points a year to cement, 0.1 to fertilizers and 0.08 to household electricity, and takes 0.2 a year from steel. Over ten years that is 4 percent more cement (Table 3). The layer keeps the models' choice for industrial output, which follows the economy, and uses $\beta_j = 1$ for electricity used by households.

Proposition 2 (one price response per activity). *If the layer applies η_j to activity supplied to a model that also applies its own elasticity η_j^m to the same activity or to fuel use that moves with it, the total log response to a cost change is $(\eta_j + \eta_j^m)\, \Delta \ln c$. Each activity's response must therefore sit in exactly one model.*

The proof is one line (Appendix A). The rule that follows is mechanical. Apply the price term in the layer only for SISEPUEDE and FTT, which have none. Do not apply it for CPAT or the CBAM Policy Tool, which have their own. The case that is easy to miss is CPAT's fleet model: its usage margin is itself an activity response, because it moves kilometers per vehicle with the cost per kilometer, so feeding it price-adjusted vehicle numbers or distances counts the response twice even when the fuel equation is switched off. When the layer feeds CPAT's road transport, it should pass the baseline only, and the fleet model should also be the selected road model, so that fuel use, emissions and co-benefits rest on one driving response (CPAT-AI Project 2026c, no. I100). The fleet model's own distance response is the road-gasoline usage elasticity of -0.4 (CPAT-AI Project 2026b), about eight times the per-vehicle response of -0.05 implied by the review evidence; paper 5.4 should test it (Stretton 2026a).

Proposition 3 (the vehicle chain needs a check). *Let cars grow with income elasticity ε_K, kilometers per car with ε_k and population exponent 0, and passenger-kilometers with ε_T, cars*

and passenger-kilometers with population exponent 1, and occupancy fixed. If $\varepsilon_K + \varepsilon_k > \varepsilon_T$ and income per person grows without bound, car passenger-kilometers eventually exceed total passenger-kilometers.

The car share of travel grows by the factor $(y_t/y_0)^{\varepsilon_K+\varepsilon_k-\varepsilon_T}$, which is unbounded (Appendix A). With the values in Table 2 held fixed, the exponent is $2.0 + 0.1 - 1.0 = 1.1$, and with income growing 2 percent a year a car share that starts at 30 percent of passenger-kilometers reaches 100 percent in year 56; one that starts at 50 percent, in year 32. A saturating ownership curve does not remove the problem. Recomputed with CPAT's Gompertz curve (saturation 800 vehicles per 1,000, starting at 3,000 dollars per person and 42 vehicles per 1,000), car travel still overtakes total travel in year 54 from a 30 percent share, by which time ownership is 363 per 1,000, and in year 29 from a 50 percent share. Ownership can rise from tens to hundreds of vehicles per 1,000 while total travel grows a few-fold. What keeps car travel inside total travel in practice is that kilometers per car and occupancy fall as ownership rises: in high-income countries ownership responds to income more than traffic does Litman (2026). The layer therefore checks each year that cars times kilometers per car times occupancy (about 1.7 people per car, an assumption) stay below a share of total passenger-kilometers, and where they do not, it lowers kilometers per car.

5.4 Feeding the four models

The layer feeds each model in its own unit, spliced onto its own last data point. For a model input X_m whose last data year is b,

$$X_{m,t} = X_{m,b}\,\frac{A_{j,t}}{A_{j,b}}, \tag{6}$$

so the model keeps its own level and takes the layer's growth. Each FTT input has its own last data year (2020 for heat, 2022 for power and transport, 2023 for freight), and equation 6 handles each separately. Because FTT adds to power demand the change in each sector's electricity use against its own baseline run, the baseline power input and the baseline sector runs must all come from the same layer run. SISEPUEDE takes the layer's base-year levels and elasticities in place of its own, and the price term through its demand scalars, inputs that scale transport demand, industrial energy demand and household heat and appliance energy demand (Syme et al. 2026). The CBAM Policy Tool takes its base-year output and price response from its own data and passes product-level output to the layer for the industrial rows. CPAT takes the layer as a cross-check on its baseline and, if its road transport is unified, as the baseline driver of its fleet model.

Two drivers and one base year make the four baselines comparable. Use the IMF World Economic Outlook for GDP (International Monetary Fund 2026) and the United Nations World Population Prospects for population (United Nations, Department of Economic and Social Affairs, Population Division 2024), for every model. Set base-year levels from one recent year: energy balances for electricity, national data or SISEPUEDE's calibrated inputs for travel and freight, vehicle registrations for cars and trucks, and the CBAM Policy Tool's latest output for covered goods. SISEPUEDE's example inputs start in 2015 and must be rolled forward with observed GDP and population.

6 A worked example

Under stylized growth of 1 percent in population and 2 percent in income per person, activity grows between 2 percent (kilometers per car) and 64 percent (cars in use) over ten years. Table 3 applies equation 5 with the elasticities of Table 2 held fixed. Total GDP grows 35 percent over the decade, which is exactly the growth of passenger travel, road freight and vans and trucks, whose income elasticity is 1.

Table 3. Cars grow much faster than travel, and the population exponent moves industrial and household activity by a few percent. Stylized drivers: population +1 percent and income per person +2 percent a year for ten years; elasticities from Table 2 held fixed. "With $\beta_j = 1$" applies Proposition 1. The last column is equation 5's price term alone, to be applied only where the receiving model has no price response of its own (Proposition 2). Source: author's calculation (`2.Source/activity_layer.py`).

Activity	β_j	ε_j	Change over 10 years	With $\beta_j = 1$	Response to a 20% rise in service cost
Electricity, homes, at $3,000 per person	0.92	0.92	+31%	+33%	−7.7%
Passenger travel	1	1.0	+35%	+35%	0
Road freight	1.0	1.0	+35%	+35%	−1.8%
Home heat, hot water and cooking	1	0.4	+20%	+20%	−3.6%
Steel and aluminum	1.2	1.2	+43%	+40%	−7.0% (steel)
Cement and clinker	0.6	0.6	+20%	+24%	−1.8%
Fertilizers	0.9	0.9	+31%	+32%	−7.0%
Cars in use, below 300 per 1,000 people	1	2.0	+64%	+64%	−4.5%
Vans and trucks	1.0	1.0	+35%	+35%	−3.6%
Km per car	0	0.1	+2%	+2%	−0.9%

Cars carry a rising share of travel, as they do in middle-income countries. Cars in use grow 64 percent and kilometers per car 2 percent while passenger-kilometers grow 35 percent, so the share of travel by car rises by about a quarter in a decade. This is the behavior FTT's transport model needs to see, and it is what a GDP-driven fuel equation cannot show: in CPAT's default road model, each road fuel grows with one elasticity, whatever happens to the mix of cars and buses.

A carbon price that raises the energy cost of driving by 20 percent cuts car use by about 5 percent, almost all through fewer cars. The cut is 4.5 percent in cars and 0.9 percent in kilometers per car, 5.3 percent in all. A 20 percent rise in the energy cost of freight cuts tonne-kilometers by about 2 percent and the number of vans and trucks by about 4 percent. These are the responses the layer supplies to SISEPUEDE and FTT, which have none. In

CPAT, the same carbon price reaches demand through equation 1's usage and efficiency margins, and the layer's price column must not be added on top.

The steel and cement rows show why product detail matters, not which product loses more. At a carbon price of 50 dollars per tonne of CO_2, with stylized intensities and prices (blast-furnace steel at 2 tonnes of CO_2 and 600 dollars per tonne; clinker at 0.8 tonnes and 60 dollars), the cost of steel rises 17 percent and the cost of clinker 67 percent. With the tool's elasticities, output falls 6.0 percent for steel and 5.0 percent for clinker. The two results are similar for opposite reasons: clinker is cheap and carbon-intensive but hard to substitute, and steel is the reverse. Only a model that carries the product price, the emission intensity and the elasticity separately, as equation 3 does, gets either right; a GDP-driven industrial fuel equation gets neither.

7 Linking the four models, and the limits

The layer turns four separate baselines into one, with one carbon-price response per activity. In practice the link runs in one direction each year. CPAT computes prices, including carbon and excise, for every fuel and sector. The layer converts them into costs of energy services, using last year's technology mix from FTT for efficiency. It then grows the ten activities from common drivers and applies the price term for the models that need it. SISEPUEDE and FTT take the activities as demand; the CBAM Policy Tool takes product output; CPAT takes the layer as its cross-check and, for road transport, as one consistent story. Revenue and distribution stay in CPAT, exporter exposure in the CBAM Policy Tool, investment needs in the extension described in paper 6.3 (Stretton 2026c), and technology diffusion in FTT.

The strongest objection is that a soft link is a second-best: why not run E3ME-FTT, or hard-link the models? A hard link solves the models together, so prices, technology and demand are consistent by construction, and E3ME-FTT already does this for demand and technology. The answer is that none of the four models in this paper was built for a hard link, that ministries and the teams who serve them run them separately and on different timetables, and that a soft link with explicit rules is easier to audit than an iteration between models. Hard-linked set-ups that iterate to convergence reach the same equilibrium in tested cases, but the way models are coupled can change the answer (Helgesen and Tomasgard 2018), and in the author's view a hard link of these four tools would need more engineering than a ministry team can maintain. Where macroeconomic feedback matters, E3ME-FTT or a computable general equilibrium model is the right tool; the layer takes growth as given, as all four models do in their standalone forms.

The link inherits the limits of each model it joins. The price elasticities for heavy industrial goods are judgment values with little empirical support. Industrial output in small open economies depends on export demand and on trade policy, such as the border adjustment itself, not only on domestic income. The household and industrial income elasticities are the author's assumptions, and the travel and car evidence comes mostly from high-income countries. None of the four models has a cooling activity, so neither does the layer.

The data gaps are as binding as the method. FTT has no car fleet data for 21 of its 71 regions, so its transport model cannot run for several large developing countries until the number of cars, kilometers per car and the costs of each car type are added. Its standalone industrial heat module ships without baseline data, and its steel model runs only inside E3ME-FTT. SISEPUEDE's public repository holds one example country; country calibrations must be obtained from its developers. CPAT's road transport needs one story, with the fleet model

selected whenever co-benefits are reported, before its co-benefits can be read as more than indicative.

For a ministry, the practical rule is short. Pick one GDP and population forecast and give it to every model. Grow activities with equation 5 from one base year. Splice onto each model's own data with equation 6. Apply the price response once, in whichever model has it. Check every year that cars fit inside travel. Compare CPAT's baseline fuel use with the layer's activity times intensity, and investigate any gap larger than the uncertainty in the elasticities.

8 Conclusion

Four models that a ministry might use to assess a carbon price build activity in four incompatible ways: CPAT from total GDP and prices, SISEPUEDE from income and population without prices, the CBAM Policy Tool from growth forecasts for a handful of goods, and FTT not at all. A single constant-elasticity rule over ten activities, a sector concordance and three checks (on population, on where the price response sits and on the vehicle chain) are enough to put them on one baseline with one price response. The main limitation is the evidence for heavy industry, where output elasticities are judgment values and trade matters more than income, followed by FTT's missing car data for developing regions. The next step is to build the layer as a small open data table beside CPAT-AI, run it for one country through all four models, and use the gap between CPAT's fuel-based baseline and the layer's activity-based one as a test of both.

References

Black, Simon, Ian Parry, Victor Mylonas, Nate Vernon, and Karlygash Zhunussova. 2023. *The IMF-World Bank Climate Policy Assessment Tool (CPAT): A Model to Help Countries Mitigate Climate Change*. IMF Working Paper 23/128. International Monetary Fund. https://doi.org/10.5089/9798400242342.001.

Burke, Paul J., and Zsuzsanna Csereklyei. 2016. "Understanding the Energy-GDP Elasticity: A Sectoral Approach." *Energy Economics* 58: 199–210. https://doi.org/10.1016/j.eneco.2016.07.004.

CPAT Team. 2024. *CPAT Documentation*. World Bank; International Monetary Fund. https://cpmodel.github.io/cpat_public/.

CPAT-AI Project. 2026a. *CPAT Methodology, Corrected Edition: The CPAT Team's 2023 Text with Errors Corrected*. CPAT-AI documentation website, Methodology tab, chapter "Corrections." https://cpmodelai.github.io/cpat_ai/.

CPAT-AI Project. 2026b. *CPAT-AI: Excel-AI and CPAT Python, with Documentation, Bug Register and Issue Register*. GitHub repository, Excel-AI release 1.2.1 (31 October 2026). https://github.com/cpmodelai/cpat_ai.

CPAT-AI Project. 2026c. *Legacy CPAT Bugs and Issues: BUGS.md and ISSUES.md, Register Version 1.1.0*. CPAT-AI repository, excel_ai/core/BUGS.md and excel_ai/core/ISSUES.md. https://github.com/cpmodelai/cpat_ai/tree/main/excel_ai/core.

Dargay, Joyce, and Dermot Gately. 1999. "Income's Effect on Car and Vehicle Ownership, Worldwide: 1960–2015." *Transportation Research Part A: Policy and Practice* 33 (2): 101–38. https://doi.org/10.1016/S0965-8564(98)00026-3.

Dargay, Joyce, Dermot Gately, and Martin Sommer. 2007. "Vehicle Ownership and Income Growth, Worldwide: 1960–2030." *The Energy Journal* 28 (4): 143–70. https://doi.org/10.5547/ISSN0195-6574-EJ-Vol28-No4-7.

European Parliament and Council of the European Union. 2023. *Regulation (EU) 2023/956 of 10 May 2023 Establishing a Carbon Border Adjustment Mechanism*. Official Journal of the European Union, L 130, 16.5.2023, p. 52. http://data.europa.eu/eli/reg/2023/956/oj.

Future Technology Transformations developers. 2026. *FTT_StandAlone: Future Technology Transformation Models in Python*. GitHub repository cpmodel/FTT_StandAlone, commit 896a3ce (28 October 2026). https://github.com/cpmodel/FTT_StandAlone.

Gertler, Paul J., Orie Shelef, Catherine D. Wolfram, and Alan Fuchs. 2016. "The Demand for Energy-Using Assets Among the World's Rising Middle Classes." *American Economic Review* 106 (6): 1366–401. https://doi.org/10.1257/aer.20131455.

Goodwin, Phil, Joyce Dargay, and Mark Hanly. 2004. "Elasticities of Road Traffic and Fuel Consumption with Respect to Price and Income: A Review." *Transport Reviews* 24 (3): 275–92. https://doi.org/10.1080/0144164042000181725.

Helgesen, Per Ivar, and Asgeir Tomasgard. 2018. "From Linking to Integration of Energy System Models and Computational General Equilibrium Models: Effects on Equilibria and Convergence." *Energy* 159: 1218–33. https://doi.org/10.1016/j.energy.2018.06.146.

International Monetary Fund. 2026. *World Economic Outlook Database*. Updated with each World Economic Outlook; latest edition at the time of writing: World Economic Outlook Update, July 2026. https://www.imf.org/en/Publications/WEO.

Jong, Gerard de, Arno Schroten, Huib van Essen, Matthijs Otten, and Pietro Bucci. 2010. *Price Sensitivity of European Road Freight Transport: Towards a Better Understanding of Existing Results*. Report Nos. 9012-1. Significance; CE Delft, for Transport & Environment. https://cedelft.eu/publications/price-sensitivity-of-road-freight-transport-towards-a-better-understanding-of-existing-results/.

Kalra, Nidhi, Edmundo Molina-Pérez, James Syme, et al. 2023. *The Benefits and Costs of Reaching Net Zero Emissions in Latin America and the Caribbean*. Inter-American Development Bank. https://doi.org/10.18235/0005330.

Knobloch, Florian, Hector Pollitt, Unnada Chewpreecha, Vassilis Daioglou, and Jean-François Mercure. 2019. "Simulating the Deep Decarbonisation of Residential Heating for Limiting Global Warming to 1.5 °C." *Energy Efficiency* 12 (2): 521–50. https://doi.org/10.1007/s12053-018-9710-0.

Labandeira, Xavier, José M. Labeaga, and Xiral López-Otero. 2017. "A Meta-Analysis on the Price Elasticity of Energy Demand." *Energy Policy* 102: 549–68. https://doi.org/10.1016/j.enpol.2017.01.002.

Liddle, Brantley. 2014. "Impact of Population, Age Structure, and Urbanization on Carbon Emissions/Energy Consumption: Evidence from Macro-Level, Cross-Country Analyses." *Population and Environment* 35 (3): 286–304. https://doi.org/10.1007/s11111-013-0198-4.

Litman, Todd. 2026. *Understanding Transport Demands and Elasticities: How Prices and Other Factors Affect Travel Behavior*. Victoria Transport Policy Institute. https://www.vtpi.org/elasticities.pdf.

Mercure, Jean-François. 2012. "FTT:Power: A Global Model of the Power Sector with Induced Technological Change and Natural Resource Depletion." *Energy Policy* 48: 799–811. https://doi.org/10.1016/j.enpol.2012.06.025.

Mercure, Jean-François, Aileen Lam, Seb Billington, and Hector Pollitt. 2018. "Integrated Assessment Modelling as a Positive Science: Private Passenger Road Transport Policies to Meet a Climate Target Well Below 2 °C." *Climatic Change* 151 (2): 109–29. https://doi.org/10.1007/s10584-018-2262-7.

Mercure, Jean-François, Hector Pollitt, Andrea M. Bassi, Jorge E. Viñuales, and Neil R. Edwards. 2016. "Modelling Complex Systems of Heterogeneous Agents to Better Design Sustainability Transitions Policy." *Global Environmental Change* 37: 102–15. https://doi.org/10.1016/j.gloenvcha.2016.02.003.

Riahi, Keywan, Detlef P. van Vuuren, Elmar Kriegler, Jae Edmonds, et al. 2017. "The Shared Socioeconomic Pathways and Their Energy, Land Use, and Greenhouse Gas Emissions Implications: An Overview." *Global Environmental Change* 42: 153–68. https://doi.org/10.1016/j.gloenvcha.2016.05.009.

Ruijven, Bas J. van, Detlef P. van Vuuren, Willem Boskaljon, Maarten L. Neelis, Deger Saygin, and Martin K. Patel. 2016. "Long-Term Model-Based Projections of Energy Use and CO2 Emissions from the Global Steel and Cement Industries." *Resources, Conservation and Recycling* 112: 15–36. https://doi.org/10.1016/j.resconrec.2016.04.016.

Schäfer, Andreas, and David G. Victor. 2000. "The Future Mobility of the World Population." *Transportation Research Part A: Policy and Practice* 34 (3): 171–205. https://doi.org/10.1016/S0965-8564(98)00071-8.

Schmidt, Cristiane Alkmin Junqueira, and Marcos André Mendes de Lima. 2006. "A Perda Do Peso Morto e a Elasticidade-Preço Da Demanda Do Setor Siderúrgico No Brasil." *Estudos Econômicos* 36 (1): 127–47. https://doi.org/10.1590/S0101-41612006000100006.

Sorrell, Steve, John Dimitropoulos, and Matt Sommerville. 2009. “Empirical Estimates of the Direct Rebound Effect: A Review.” *Energy Policy* 37 (4): 1356–71. https://doi.org/10.1016/j.enpol.2008.11.026.

Stretton, Stephen J. 2026a. “Are CPAT’s Price Elasticities Right? A Check Against the Evidence.” Unpublished manuscript.

Stretton, Stephen J. 2026b. *Measure Carbon Pricing by the Emissions It Cuts: The Total Effective Carbon Price*. 6.1. Climate Fiscal Policy Briefs.

Stretton, Stephen J. 2026c. “Public Investment in CPAT: Investment Needs and the Efficiency of Public Spending.” Unpublished manuscript.

Stretton, Stephen J. 2026d. “The CBAM Policy Tool.” Unpublished manuscript.

Stretton, Stephen J. 2026e. “The CPAT Mitigation Module: Methods in Brief.” Unpublished manuscript.

Syme, James et al. 2026. *SISEPUEDE: Simulating Sectoral Pathways and Uncertainty Exploration for Decarbonization*. GitHub repository, commit 9bf976c (23 September 2026), with documentation in docs/source. https://github.com/jcsyme/sisepuede.

United Nations, Department of Economic and Social Affairs, Population Division. 2024. *World Population Prospects 2024: Summary of Results*. UN DESA/POP/2024/TR/NO.9. United Nations. https://population.un.org/wpp/.

Vercoulen, Pim, Soocheol Lee, Xu Han, Wendan Zhang, Yongsung Cho, and Jun Pang. 2023. “Carbon-Neutral Steel Production and Its Impact on the Economies of China, Japan, and Korea: A Simulation with E3ME-FTT:Steel.” *Energies* 16 (11): 4498. https://doi.org/10.3390/en16114498.

Wene, Clas-Otto. 1996. “Energy-Economy Analysis: Linking the Macroeconomic and Systems Engineering Approaches.” *Energy* 21 (9): 809–24. https://doi.org/10.1016/0360-5442(96)00017-5.

Appendix A: Proofs

Proposition 1. In logs, equation 5 gives $\Delta \ln A_j = \beta_j \Delta \ln P + \varepsilon_j \Delta \ln y + \eta_j \Delta \ln c$. Changing β_j from ε_j to 1 changes $\Delta \ln A_j$ by $(1 - \varepsilon_j)\Delta \ln P$, which is zero if and only if $\varepsilon_j = 1$ or population is constant.

Proposition 2. If the layer supplies $A_j\,(c/c^{\text{base}})^{\eta_j}$ and the receiving model multiplies its input, or fuel use proportional to it, by $(c/c^{\text{base}})^{\eta_j^m}$, the result is $A_j (c/c^{\text{base}})^{\eta_j + \eta_j^m}$, whose log response to c is $\eta_j + \eta_j^m$.

Proposition 3. Car passenger-kilometers over total passenger-kilometers is $\sigma_t = \sigma_0 (y_t/y_0)^{\varepsilon_K + \varepsilon_k - \varepsilon_T}$, since the population terms cancel when cars and passenger-kilometers both have exponent 1, kilometers per car has exponent 0 and occupancy is fixed. If the exponent is positive and $y_t \to \infty$, σ_t exceeds 1 in finite time, when $\ln(y_t/y_0) > -\ln \sigma_0 / (\varepsilon_K + \varepsilon_k - \varepsilon_T)$. With an exponent of 1.1 and 2 percent income growth, $\sigma_0 = 0.3$ crosses 1 after 55.3 years and $\sigma_0 = 0.5$ after 31.8. Positive growth alone is not enough: if growth rates fall fast enough that y_t converges, σ_t stays bounded. With a Gompertz ownership curve $O(y) = \bar{S} \exp(a e^{by})$, the elasticity $d \ln O / d \ln y = ab\, y\, e^{by}$ peaks at $y = -1/b$ (y in thousands of dollars) and then falls to zero, so once ownership saturates σ_t falls again whenever $\varepsilon_k < \varepsilon_T$. But ownership can multiply many times before it saturates: in the example of Section 5.3 it rises from 42 to 363 vehicles per 1,000 by the year car travel passes total travel, and σ_t peaks at 1.15 from a starting share of 0.3. Code: `2.Source/activity_layer.py`.