Comparing the Energy Policy Simulator to Other Energy and Climate Models

INTRODUCTION

The Energy Policy Simulator (EPS) is a free and open-source energy and climate policy model developed by Energy Innovation. It is available online at https://energypolicy.solutions. The EPS has been developed in consultation with universities, national laboratories, think tanks, and advocacy groups and deployed in many regions, including many of the world's largest emitters: China, India, the United States, Indonesia, Mexico, Canada, and Saudi Arabia. Versions have also been created for sub-national areas, such as U.S. states and Chinese provinces.

The EPS is developed with a system dynamics framework using a program called Vensim. A free version of Vensim, called Vensim Model Reader, can be used to run the EPS. The EPS is built to be highly modular, allowing users to update and change any input data without the need for commercial software or programming expertise.

The EPS is also deployed in a powerful web application interface, allowing users to interact with the model via a web browser. The web application interface is designed to be user friendly and extremely fast, with analysis capabilities beyond what is possible when running the EPS locally in Vensim.

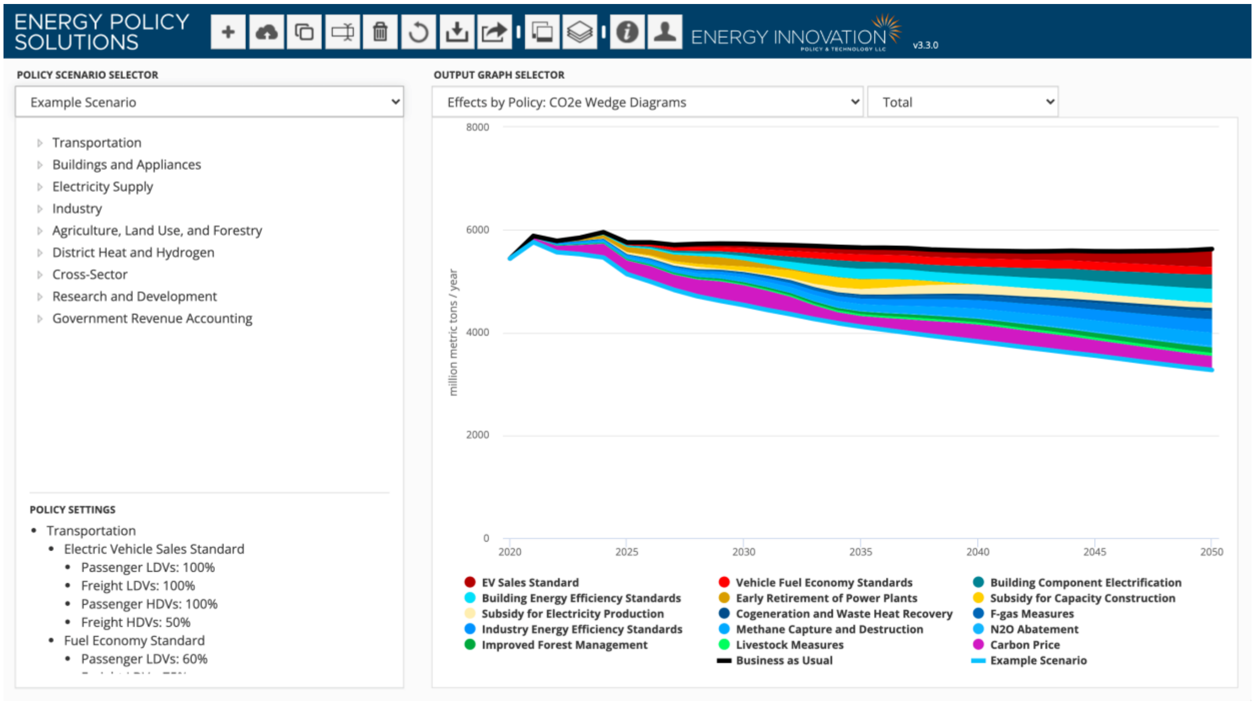

POLICY SETTINGS

- Transportation
  - Electric Vehicle Sales Standard
    - Passenger LDVs: 100%
    - Freight LDVs: 100%
    - Passenger HDVs: 100%
    - Freight HDVs: 50%
  - Fuel Economy Standard
    - Passenger LDVs: 60%

Example of the EPS web interface

There are many energy and climate models, each of which operates differently, with certain strengths and weaknesses. This document outlines the general structure and theory of the EPS and compares it to some of the most widely used and referenced energy models in the U.S., including the Global Change Assessment

Model (GCAM), the National Energy Modeling Systems (NEMS), and stock turnover models like the Low-Emissions Analysis Platform (LEAP) and its derivatives.

THE STRUCTURE AND THEORY BEHIND THE EPS

Models are designed with different use cases and objectives in mind. For example, GCAM is one of several Integrated Assessment Models (IAMs), whose goal is to show technology pathways and potential economic and energy impacts under carbon and climate constraints. NEMS is designed to represent the U.S. energy system in detail and specific U.S. policies (such as corporate average fuel economy standards or tax credits). LEAP and its derivatives are designed to evaluate detailed technology stock turnover and potential impacts of technological solutions such as different deployment rates or efficiency levels. (Note: these are all simplifications of what these models can do but are useful as a point of comparison). While each model can be used for other purposes, these are the core capabilities and design criteria.

The EPS is designed to estimate combined impacts of policies on technology deployment, energy, emissions, public health, and the economy for any region in which it might be useful. Because it is used internationally, the model is designed to be highly adaptable and open source. We structured the model to be as broadly applicable as possible.

To achieve these goals, the EPS is designed as an annual, single-region, forward-simulating system dynamics model (with an embedded macroeconomic input-output model). The model simulates a business-as-usual case (BAU) and a policy case, with policies set based on user input, and computes differences between the scenarios to estimate policy impacts on emissions of 12 pollutants, cash flows, jobs, gross domestic product, and public health outcomes, among other outputs. The model runs in annual timesteps to 2050 by default, but it can be configured to run through 2100, provided the necessary input data is available. The theory and structure of the EPS can generally be broken down as follows: 1) separation of model structure and input data; 2) energy and service demand projections; 3) optimized technology deployment; 4) policy implementation; and 5) estimating policy impacts.

Separating model structure from data

All the input data used by the EPS is contained in comma separated value (CSV) files read by the model at run-time. Thus, to change any input data, a model user need only update the relevant CSV file(s) and run the model (and a single model run takes a fraction of a second). CSV files are accompanied by Excel spreadsheet files that provide full bibliographic information for data sources and handle data formatting tasks. This means the EPS can be easily updated to use the most recent available data, such as cost projections for wind, solar, or batteries, or to conduct sensitivity analysis (even without using built-in features in Vensim, though the EPS running in Vensim DSS also supports Monte Carlo analysis).

This contrasts with other models, which often include data embedded in the model structure or have significant barriers to routine data updates. For example, GCAM contains input data in highly complex XML files, requiring the use of Rstudio and an accompanying dataset to produce input data files. The input data is not transparent. There are also inputs embedded into the code itself, which make it difficult to update data as desired. Ease of data updates is very important for energy models because technology costs change rapidly. Using the most up-to-date cost data is critical to accurately assessing emissions trajectories and policy impacts. By separating the data from the structure, users can also quickly test multiple sensitivities and variations on input data and explore how results change.

Another benefit of separating model structure from data is the ease with which the EPS can be adapted to different geographies. Customized versions for different U.S. states, regions, or countries can be developed and deployed simply by updating the input data and control settings in CSV files. This makes it possible to deploy new features and fix bugs in the core model structure without having to manually port these changes to numerous model regions. This avoids a shortcoming of models such as the UK DECC-authored Global Calculator, which is implemented in Excel and has no separation between model structure and data.

What Are the Best Policies to Solve Climate Change?

Designing energy policy correctly is tough work. Well-designed energy policies reduce pollution, cut consumer costs, and minimize dependence on foreign energy supplies. Done wrong, they can do the reverse, and increase pollution, lock in dirty technologies, or waste money.

Which energy policies can lower greenhouse gas emissions in Zhejiang, and what will it cost? Discover the effects of various energy policies, build your favorite policy package, and share it with friends using the Energy Policy Simulator. Try it using the button below:

© 2022 Energy Innovation. All rights reserved.

Landing page for the EPS customized for Zhejiang Province, China

Energy and service demand projections

Unlike models such as GCAM, NEMS, and LEAP, the EPS does not endogenously calculate energy and service demand, but rather relies on outputs of other models to serve as this input. Specifically, demand for energy in buildings (by building type, equipment type, and energy source) is read into the model as input data. For example, the U.S. model relies on data from the Energy Information Administration’s Annual Energy Outlook (AEO) as input data. Building energy consumption therefore matches data from the AEO. We use the same approach in industry, where industrial energy demand is pulled from external data, such as the AEO.

In the transportation sector, we use service demand (passenger-miles and freight-miles), not energy demand, as the input data. The model then endogenously selects vehicle technologies and deploys new vehicles (this is discussed in more detail in the next section), starting with the existing stock. Similarly, in the electricity sector, the EPS determines electricity demand (i.e., as the sum of demand from the other sectors), then decides what to build using an approach that accounts for costs, grid flexibility constraints, etc.

The approach used for energy and service demand in the EPS is distinctly different from the approaches used by GCAM, NEMS, and LEAP because it does not try to calculate demand endogenously. Rather, it relies on sector-specific models with much greater detail and uses the produced information directly. This avoids duplicating successful structures from other models, makes it easier for policymakers to use their own national or state data sources, and facilitates extremely fast run times in the EPS (less than a second per run).

Optimized technology deployment

For certain sectors, the EPS endogenously selects new technologies to meet energy or service demand. In transportation, the EPS optimizes the deployment of new vehicles to meet growing demand plus retirements using total cost of ownership and includes non-monetary barriers (e.g., range anxiety). This approach is similar to the operation of other models, such as GCAM, NEMS, and LEAP. Demand for energy is computed based on the evolution of the vehicle fleet.

Similarly, the EPS endogenously estimates new power plant additions based on a modified levelized cost of electricity, accounting for seasonal peak demand requirements and grid flexibility constraints (which can cause curtailment of variable solar and wind, if there is insufficient flexibility). This approach is similar to approaches used in other models, such as GCAM and NEMS. The EPS includes several methods for deploying capacity, including least-cost, policy-driven (e.g., through a renewable portfolio standard or clean energy standard), and mandated capacity additions.

The EPS also includes multiple methods for dispatching electricity. Users can utilize the least-cost dispatch mechanism, which mimics a competitive electricity market; the guaranteed dispatch mechanism, which sets minimum or exact levels of dispatch for different power plant types, or a hybrid of both. This flexibility is valuable because while regions may in theory operate an electricity market, plants are often run uneconomically, and a pure least-cost dispatch model will be unable to capture this effect. Note these approaches are simplified and not a replacement for detailed power system dispatch or capacity expansion models.

The dynamic nature of the power sector in the EPS is distinct from some other types of models. For example, in previous iterations, GCAM dispatched power plants at fixed capacity factors and was unable to simulate anything approaching least-cost dispatch. (GCAM-USA has more detailed dispatch now for U.S. states, but for other countries the old approach is still used).

Energy Innovation is working to update the EPS’s building and industry sectors to add optimized technology deployment in the future.

Policy implementation

Unlike many other models, the EPS was designed from the outset with the idea that users would interact with the model through implementing an array of policy options, and not through strict technology adoption or a carbon cap or climate target. This differs from the other models discussed here. For example, GCAM is an IAM focused on pathways to meet carbon constraints, principally through a carbon price. While it can be used to evaluate other policies, it was not originally designed for this purpose, and thus policy modeling can be very challenging, with users needing to translate policies into a format (often markets or hard-coded deployment constraints) that conforms with GCAM’s specific structure.

By contrast, the EPS offers hundreds of policy options, with the entire design centered on application of policies. It therefore offers a much easier user experience and far longer list of policy options. This is perhaps the most unique feature of the EPS: it provides users with numerous policy options that can influence technology adoption, but technology adoption is not the model's sole driver of change. This means the EPS can handle many more policies than other models, can estimate their impacts much more quickly, and does not have to be heavily customized to evaluate common policies.

For example, reflecting coal plant retirements as a policy in GCAM can be very cumbersome given the limited way in which GCAM allows for retirement (only through economic retirement). However, the EPS has a policy lever for coal plant retirements, which can reflect, for example, state Integrated Resource Plans that mandate retirement of specific units. Specifying exact retirement data in the EPS is therefore very quick and easy compared to a model like GCAM.

Another example of differences in ease of policy modeling is the zero-emissions vehicle (ZEV) standard. In the EPS, users simply set the ZEV standard percentage and implementation timeline for different vehicle types, and the model instantaneously calculates deployment. In GCAM, users must configure several input XML files, setting up a new market and assigning credit values to certain vehicle types in non-intuitive, undocumented units, all of which takes significant time (in fact, Energy Innovation hired a consultant to build out software to better automate this process, including for several other policies, given how cumbersome it can be to model these in GCAM). Further, in our experience running GCAM with this policy, it failed to solve, and we had to iteratively hard code deployment constraints to get the correct ZEV standard impact.

Estimating policy impacts

The EPS tracks two scenarios in every run: a BAU scenario that is unaffected by policy choices, and a policy scenario that reflects the policy choices made by the user. The EPS therefore provides real-time comparisons between the policy scenario and the BAU scenario, and the scenarios can be run iteratively thousands of times within a few minutes.

The EPS is designed to account for the interactive nature of policies, including sequencing when relevant, and to correctly capture the dynamic effects of policies across technologies and sectors (e.g., increased electrification of end-use technologies will increase electricity demand and require additional power plants, and increased costs to build out the infrastructure are reflected in electricity prices and mitigate some demand growth through price elasticities). These capabilities allow the EPS to produce powerful, unique types of graphs, such as "wedge diagrams" (which break a policy package's effects out by component policies) and "cost curves" (which illustrate the cost-effectiveness of each policy within a policy package). They are generated dynamically based on the user's policy settings, unlike published curves, which are calculated only once (often based on technical potential) and then remain static.

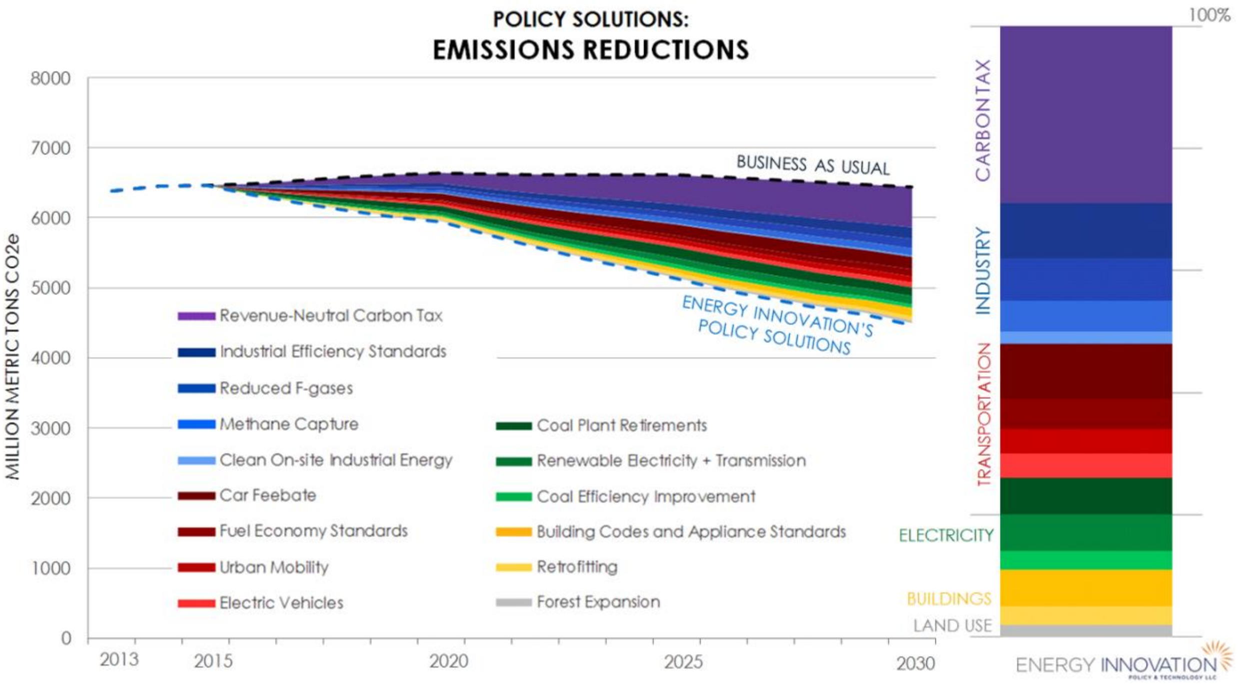

Example wedge diagram showing policy contributions to GHG abatement

The EPS also includes an embedded macroeconomic input-output model that allows for estimation of jobs and GDP impacts as a result of policy decisions. The model includes macroeconomic feedbacks with a one-year time delay to account for interactive effects. In addition, the EPS calculates avoided premature mortality and avoided morbidity outcomes, based on risk-per-ton datasets.

Conclusion

The EPS is designed to integrate the outputs of sector-specific models and to provide insight into how policy measures lead to changes in emissions, cash flows, public health impacts, jobs, and economic output. This design allows for incredibly fast computation. The model was also designed at the outset with *policy* as the primary interaction point, and not technology adoption or carbon price/carbon constraints. As such, it provides much more flexibility than other models in modeling policy options, and can do so much faster, and while accounting correctly for policy interactions.

To achieve these results, the EPS has reduced complexity in certain areas relative to other models. For example, it is only applied to a single region at a time and may not have the same level of technological detail as other models.

Yet the EPS adds capabilities that other models typically lack, such as an embedded input-output model, estimation of health impacts and monetized avoided health and climate damages, and much more policy detail. It can also be fully run in an interactive web browser, which can calculate complex outputs such as policy wedge diagrams (see above) and marginal abatement cost curves.