NBER WORKING PAPER SERIES

MODELING UNCERTAINTY IN CLIMATE POLICY: AN APPLICATION TO THE US IRA

James B. Bushnell
Aaron Smith

Working Paper 32830
http://www.nber.org/papers/w32830

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2024

This research was partially supported supported by a grant from the California Air Resources Board (CARB). Bushnell also received support from NSF Grant 2330450. We thank Wuzheqian Zhao, Tengda Gong, and Jessica Lyu for outstanding research support. The statements and conclusions in this paper are those of the authors and not necessarily those of the any supporting institution. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Modeling Uncertainty in Climate Policy: An Application to the US IRA
James B. Bushnell and Aaron Smith
NBER Working Paper No. 32830
August 2024
JEL No. Q4,Q47,Q58

ABSTRACT

In recent years the analysis of US climate policy on the electricity sector has predominantly deployed electricity planning or capacity expansion models that use deterministic or equilibrium optimization methods. While uncertainty in key input assumptions is considered, it is usually restricted to scenario analysis. In this study we combine time-series econometric forecasting methods with an equilibrium electricity system-expansion model. The goal is to produce statistically rigorous distributions of outcomes, rather than rely upon individually selected scenarios.

We apply these techniques to the case of the US Inflation Reduction Act (IRA) in the context of the western US electricity grid. The most significant power sector financial incentives are tax credits applied to eligible zero-carbon and storage resources. Our results indicate that the impact of the IRA, in terms of additional investment in low-carbon resources, depends heavily on the realization of key exogenous variables. However, the net effect of the IRA is to sharply narrow the range of future carbon emissions, largely by eliminating states of the world where investment in natural gas resources would otherwise be optimal.

James B. Bushnell
Department of Economics
One Shields Ave.
University of California, Davis
Davis, CA 95616
and NBER
jbbushnell@ucdavis.edu

Aaron Smith
Dept.of Agricultural and Resource Economics
One Shields Ave.
University of California, Davis
Davis, CA 95616
adsmith@ucdavis.edu

Modeling Uncertainty in Climate Policy: An Application to the US IRA

James Bushnell and Aaron Smith*

July 30, 2024

1 Introduction

Models of energy systems have been influential inputs to the process of forming climate policy at the state and federal level in the United States and internationally. Typically, modeling of energy systems and the greenhouse gas (GHG) emissions they produce have been dominated by either "bottom-up" models of individual or sector level outcomes using input-output model frameworks or by "top-down" optimization models of energy systems. Examples of the former include the PATHWAYS model used by E3 and of the latter include the NEMS model used by EIA. Other models (e.g. Purdue's GTAP for trade, agriculture and land use) deploy varying levels of sector-level detail and attempt to capture general equilibrium effects between sectors through relationships specified through input-output tables. These models attempt to simulate, often at a sector or local level, relationships between stimuli such as energy prices and policy interventions and outcomes such as CO2 emissions.

Within the electricity sector, several prominent and sophisticated models have been deployed to either scope out potential scenarios for deep decarbonization or to study specific policy proposals. The Energy Information Administration's flagship National Energy Modeling System (NEMS) is deployed for many important analyses, including the EIA's annual energy outlook. The National Renewable Energy Laboratory (NREL) has created an open source model, the Regional Energy Deployment System (ReEDS) model, that adds more electricity system specific elements. Other prominent models have been deployed by research groups at MIT (GenX), Princeton (REPEAT), and the Electric Power Research Institute (REGEN).

*This research was supported by a grant from the California Air Resources Board (CARB). Bushnell also received support from NSF Grant 2330450. We thank Wuzheqian Zhao, Tengda Gong, and Jessica Lyu for outstanding research support. The statements and conclusions in this paper are those of the authors and not necessarily those of the any supporting institution. Email addresses: Bushnell: jbbushnell@ucdavis.edu; Smith: adsmith@ucdavis.edu.

When implemented well, these electricity or energy system models are typically calibrated to a "benchmark" year to provide normalized starting conditions, but they are usually not otherwise calibrated, or otherwise restricted, to fit actual trends in the data over time or between sectors. Because they are based upon discrete optimization frameworks, typical models are also limited to producing "point" estimates of single values for each output, rather than probabilistic ranges for those output values. Uncertainty can be tested via scenario analysis, but frequently the selection of scenarios is not based upon any statistical analysis of the likelihood or relevance of those scenarios.

There are several sources of uncertainty associated with a modeling exercise of this nature. One central aspect is what we will call *model uncertainty*, which we define as uncertainty about how accurately the simulation model represents actual real-world investment and operations decisions, or what econometricians refer to as the data generating process (DGP). One key question in this dimension is the extent to which the decisions of decentralized economic agents are well represented by the solution to a market-wide optimization problem. Another aspect that is extremely difficult to capture quantitatively is *policy uncertainty*, which we define as uncertainty over the future evolution of policies that are influential to the DGP. In the case of environmental policies directed at the electricity sector, such elements might include state-level policies like renewable portfolio standards and carbon pricing, separate but still impactful federal policies such as the moratorium on LNG export expansion, and regulatory policies such as electricity industry restructuring, market design, and policies impacting transmission investment.

A third source of uncertainty, and the main focus of this paper, is what we term *data uncertainty*, which we define as uncertainty about the future values of key inputs into the model that are influential in determining output values. Weather, technology, consumer preferences, and macroeconomic activity are examples of variables over which we have data uncertainty. We incorporate data uncertainty using time series econometric methods to generate forecasts by extrapolating historical trends and correlations into the future. The models are probabilistic, which enables us to focus on predicting the range of possibilities rather than a single number. We use methods from empirical macroeconomics which, since the 1970s and in contrast to energy system modeling, increasingly deploy relatively parsimonious econometric time-series models for policy analysis and forecasting (Stock and Watson, 2001).

We propose and develop time series methods for policy analysis and forecasting of energy systems in multiple regions across multiple sectors. The key insight is that the trajectories of a large number of time series are driven by a small number of common factors (Stock and Watson, 2016). State-level economic activity is driven mostly by factors common to all states. Sector-level economic activity is driven mostly by factors common to all sectors. This approach, known as dynamic factor modeling, promises the twin benefits of better forecasts of any variable of interest and a way to downscale forecasts by region or economic sector.

Dynamic factor models reflect macroeconomic theories of general equilibrium. A central theme of macroeconomics is that a few factors such as technological development, aggregate demand, and the money supply explain how economies evolve. Macroeconomists of different eras and schools of thought may disagree about the importance of particular variables, but their models contain the core feature that the actions of large numbers of economic agents combine to produce economic phenomena that can be explained by a few factors.

From a statistical perspective, dynamic factor models provide a parsimonious way to incorporate large numbers of variables. We want to use a large number of variables both because they may be helpful in forecasting each other (e.g., in identifying the factors) and because we want forecasts of numerous variables across regions and sectors of the economy. Parsimony is particularly important when studying climate change and the associated policies because the relevant historical time series are relatively short. We often have 20-30 years of data on tens or hundreds of variables. Robust statistical analysis is only feasible in such a setting if we can reduce the dimensionality of the problem.

In this paper we demonstrate our approach, which combines a dynamic factors time-series forecast of market inputs with an electricity system optimization model for analysis of climate policies. Our context is the 2022 U.S. Inflation Reduction Act (IRA). The IRA contains dozens of provisions, but our focus here is on the investment and production tax credits provided to clean energy and storage sources. The following section provides more background on the IRA and the modeling efforts that have been used to examine it. Section 3 provides an overview of the electricity market simulation model, which has been adapted from earlier work (e.g. Bushnell et al. (2017)). Section 4 summarizes our data and gives a brief description of our forecasting approach, which is described in more detail in Appendix A. We present our results in Section 5 and then conclude in Section 6.

2 Modeling The Inflation Reduction Act

The Inflation Reduction Act (IRA), passed in 2022, contained dozens of provisions ranging from healthcare subsidies to expanding the budget of the IRS. Much of the policy and media focus was on the elements of the IRA directed at the electricity sector (Krugman, 2022). While these provisions included support for research and development of new technologies and for "onshoring" production of green technologies, the largest expected expenditures included tax credits for the purchase of electric vehicles and other residential electrification and the extension and expansion of production and investment tax credits for new sources of zero carbon electricity production and storage. These provisions were originally estimated to cost around $140 billion over 10 years (CBO, 2022), but estimates of the costs of these credits has steadily grown since then.

While even the climate related provisions of the IRA are extensive, we focus here on the elements most likely to directly influence investment and operations in the wholesale power sector. These elements are the extended and expanded production and investment tax credits for generation from wind, solar, nuclear, and other zero-carbon sources, as well as the expansion of an ITC to grid-scale storage technologies. The production tax credits could range from as low as $5/MWh to as high as $32/MWh, depending upon eligibility for bonuses linked to labor requirements, domestic content, and locations in specific communities. The ITC has a similar broad range based upon similar eligibility criteria. Importantly, the duration of clean electricity production and investment tax credits in the IRA are linked to overall emissions in the U.S. power sector and could be collected "well into the 2030s and potentially longer" (Bistline et al., 2023b). This semi open-ended nature of the tax credits, as well as the fact that total costs will depend upon the quantity of eligible resources claiming the credit, combine to create a large range of uncertainty about the potential cost of these credits.

Many contemporaneous studies modeled the prospective impact of these (and other) IRA provisions on the U.S. wholesale power grid. Most deploy electricity production and investment models similar in character, but in many cases more detailed, to the model we deploy here. These models co-optimize the investment and operation of existing and newly deployed electricity supply technologies subject to the requirement to satisfy electricity demand in a single or sequence of future years. Typically both operational characteristics of specific generation technologies and limits of the transmission network are also captured, with varying degrees of detail and methods.

These models are necessarily very complex, and can be very sensitive to modeler choices of both input assumptions and model specification. Sensitivity to input assumptions is typically tested using specific scenarios. For example, Steinberg et al. (2023) apply their simulation across "seven sets of assumptions with varying projected future electricity market conditions, including technology costs and performance, natural gas prices, and the degree of availability, feasibility, and cost of deployment of renewable resources." Results are usually conveyed in terms of "high, low, and middle" scenarios, where these bounds are often defined by the outcomes being produced by the model. To our knowledge, no other study has attempted to estimate a statistical model of the future distribution of any of the typical inputs used in these types of modeling studies.

Two useful papers, Bistline et al. (2023a) and Bistline et al. (2023c) summarize and compare many of the most prominent studies that examined the potential impacts of the IRA. These papers are illustrative efforts that convey both the common and varying results of both expectations of future electricity market conditions and the impacts of the IRA itself. They help to convey a sense of the range of expectations, but only as far as expressing the range of assumptions and modeling techniques deployed by the various research

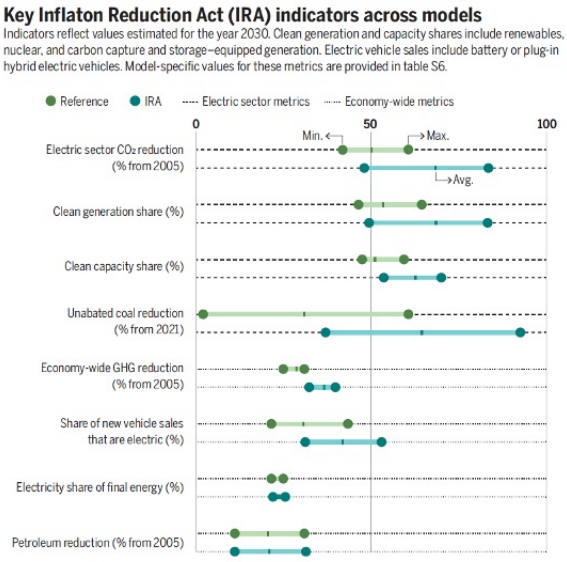

Figure 1: IRA Simulation Results (from Bisteline, et al, 2023)

teams. For example, Figure 1, taken from Bistline et al. (2023a) illustrates the range, and mean, of various outcomes of the IRA across a half-dozen model platforms. They help to convey a degree of uncertainty, but it is a form of uncertainty *across* models rather than of the policy itself *within* a given model.

3 Electricity Market Investment Model

We implement our methodology using a model of the electricity system of the western electricity coordinating council (WECC), encompossing most of the western U.S. The model is an extension of previous versions used to study the clean power plan (CPP) (Bushnell et al., 2017) and leakage from carbon regulation in California (Bushnell and Chen, 2012; Bushnell et al., 2014). We briefly describe key features of the model here.

3.1 Model Specification

The model includes several standard features of electricity system operations and investment. It simultaneously optimizes the dispatch of generation and storage for a sequence of eight 24-hour representative days. Demand is assumed to be linear but with low elasticity and is aggregated to the regional level. The objective function is to maximize social welfare, defined as the weighted sum of hourly consumption less generation costs, including annualized investment costs F_j described below.[1]

[1]Consumer benefit is defined as the area under the demand curve out to the quantity consumed in a given hour and region. In other words it is consumer surplus with the cost of purchasing the electricity added back to reflect the fact that this purchase cost is a transfer to producers and not part of a social welfare calculation.

$$\sum_t \left[\sum_r CB(p_r(t)) - \sum_{i \in r} c_i q_i(t) \right] Tweight(t) - \sum_{jinr} F_j^r * cap_j$$

Where c_i represents the unit level marginal cost of unit i located in region r, $p_r(t)$ is the price in region r, $CB(p)$ is the consumer benefit, $c_i * q_i$ is the variable cost of plant i, and $F_j * cap_j$ is the annualized fixed capital cost of cap_j (new) MW of plant j. The term Tweight is the weight placed upon a given time period based upon the number of actual days in the calibration year it is representing. The production of each generation unit i is constrained to be less than or equal to the capacity of unit i. Within each representative day, the change in hourly output from baseload generation units (coal and combined cycle gas) is further constrained by unit level ramp rates. The model does not, however, impose integer "unit commitment" constraints. For existing generation units, capacity and ramping capabilities are calibrated based upon unit performance in 2019. For new generation, indexed by j, hourly output is also limited by capacity, but capacity for these units is a choice variable.

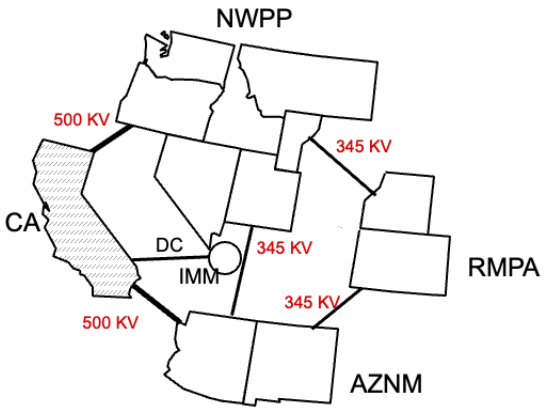

Figure 2: Western Grid Model Regions

The model represents transmission limits between the four sub-regions of the WECC as illustrated in Figure 2. Flows between the four large regions are represented by a four-node direct-current (DC) load model, Limits on transmission between two regions is set based upon 2019 observed seasonal maximums. In addition there are constraints on the total inflows and outflows into and out of a region (sometimes referred to as "nomogram" constraints that capture the maximum simultaneous limit of multiple transmission paths), again based upon the 2019 observed seasonal extreme values.

3.1.1 Model Calibration

Dispatch in the model is calibrated using market and operations data from 2019. Hourly local marginal prices are taken from SNL Global. Hourly demand and generation from hydro, nuclear, and renewables are taken from EIA form 930, which reports hourly loads and supply by technology type. Hourly production from thermal resources is taken from the Continuous Emissions Monitoring (CEMS) data provided by the US EPA. Power flow limits between regions is derived from observed seasonal maximums taken from EIA 930.

Table 1: Summary Statistics of 2019 Western Electricity Market

	AZNM	CA	NWPP	RMPA
Actual Values				
Coal Generation	4248.4	847.0	7283.9	3651.8
Gas Generation	6237.7	7859.7	7825.4	1980.1
Total Generation	18932.5	20584.3	33407.7	7971.5
Wholesale Price	30.5	35.0	34.4	30.7
Net Exports	3575.4	-7045.9	2783.6	838.3
Simulated Values				
Coal Generation	4191.9	885.8	7221.4	3547.3
Gas Generation	6575.0	7679.9	7624.8	1984.9
Total Generation	19213.3	20443.1	33144.7	7871.8
Wholesale Price	31.8	35.3	35.0	30.9
Net Export	3880.2	-7116.0	2495.0	740.0

Table summarizes average hourly generation (MW), prices ($/MWh) and net exports (MW) for each subregion. The top panel summarizes actual 2019 values from data sources. The bottom panel summarizes output from simulation of 2019.

Hourly data from 2019 are collapsed into one peak and one (non-peak) "average" 24 hour day for each season. The intercept of the hourly linear demand in the model is set such that the quantity demanded equals the observed demand at the observed market price. Production from existing nuclear, hydro, renewable, and combined heat and power (CHP) plants is assumed to operate as it did during 2019. The model then dispatches existing thermal generation to meet demand at minimum cost, subject to the constraints described above. Table 1 presents summary statistics for generation, demand, and prices for the four western regions. The upper panel of Table 1 summarizes data for the actual market in 2019 and the lower panel summarizes the output of the calibrated simulation model.

The result is eight sequences of 24 hour representative dispatch 'days', one for each region of the WECC that are connected by the network summarized in Figure 2. Sequential hours are important to capture inter-temporal considerations such as ramping rates and battery storage. The hourly profiles of supply from various resource types is illustrated in Figure 3 using the actual market data (left-hand bars for each hour)

and the calibrated simulation results (right hand bars for each hour).

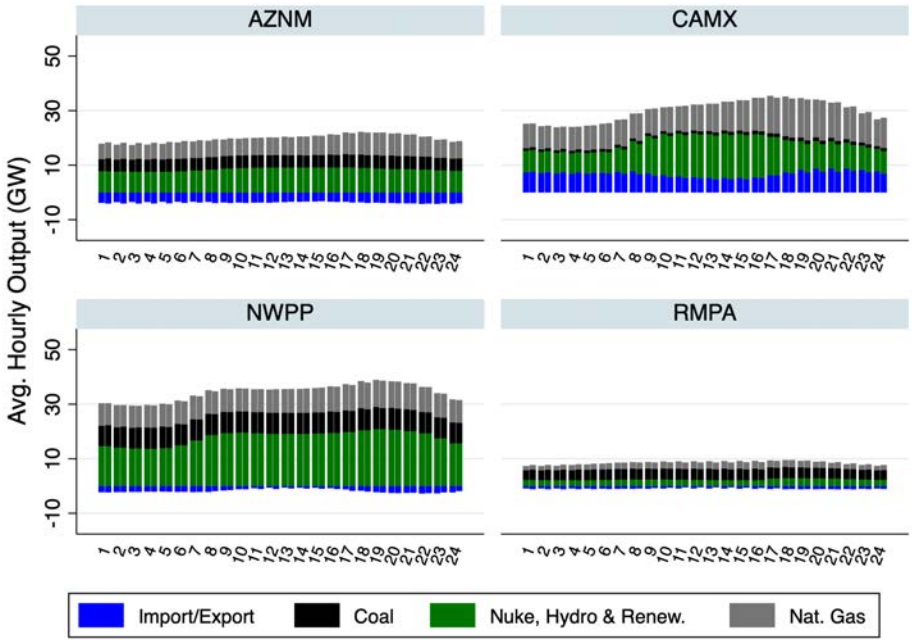

Notes: Figure summarizes actual and simulated average generation by resource type and region, and hour of day for 2019. The left-hand bars for each hour are actual values and the right hand bars are simulation output. The region definitions are Arizona and New Mexico (AZNM), California (CAMX), Pacific northwest (NWPP) and Rocky Mountains (RMPP).

Figure 3: 2019 Hourly Output by Season, Region, and Source: Actual and Simulated Data

3.1.2 Investment in New Generation

The model also calculates the static equilibrium level of investment in new generation technologies. The equilibrium condition for investment in technology j is that operating profits based upon equilibrium prices equals the entry cost of the marginal unit of capacity for that technology. For the renewable technologies, wind and solar, this means that entry increases, and lowers prices, until the time and capacity factor weighted prices in a given region of entry equals the annualized investment cost for that technology and region, or

$$\sum_t p_r(t) * cf_j(t) = F_j$$

where $cf_j(t)$ is the capacity factor of a new unit of capacity of technology j in region r. Hour by season capacity factors are calculated at the state level and are based upon observed 2022 values from EIA 930 data.

Table 2: Capital Cost of New Generation

2022 Values				
Technology	AZNM	CA	NWPP	RMPA
Battery Storage	1259	1325	1300	1310
Nat. Gas CCGT	934	1553	1264	1023
Nat. Gas CT	658	910	816	637
Wind	1411	2447	2057	1411
Solar PV	1300	1440	1332	1315
Assumed 2030 Values (2022 Dollars)				
Technology	AZNM	CA	NWPP	RMPA
Battery Storage	944	994	975	983
Nat. Gas CCGT	934	1553	1264	1023
Nat. Gas CT	658	910	816	637
Wind	1129	1958	1646	1129
Solar PV	1170	1296	1199	1184

Table summarizes values used for overnight capital cost of new generation technologies. The values for 2022 come from the Energy Information Administration's "Cost and Performance Characteristics of New Generating Technologies Values" used in the 2022 Annual Energy Outlook. Units are 2022$/kW of capacity.

In other words, production from new capacity invested in a given state is assumed to mimic the production observed from the same technology in 2022.

There are five technologies for which investment choices can be made. Two are natural gas based, two are renewable, and one is battery storage. We apply the 2022 values for regional overnight capital costs taken from the EIA's Annual Energy Outlook (AEO). These capital costs are summarized by technology and region in Table 2.

One of the key elements of uncertainty about the IRA and any policy targeting the electricity sector is the future trajectory of these relative capital costs. Almost all of the major studies of the IRA assume exogenous technical change that would substantially reduce the capital costs of solar PV and battery storage, and to a lesser extent onshore wind, through 2035. With the benefit of two years of additional data, we can observe that the real costs, at least according to the AEO, of solar and wind have stabilized or even increased slightly since 2020 and only battery storage costs have continued to decline. In the absence of a reliable means of forecasting future technology, we make relatively conservative assumptions about the future costs of all of these technologies. We assume that wind capacity costs will decline by 10%, solar PV by 20%, and battery storage by 25% by 2030. This does differentiate our results from the earlier studies, particularly with regards to the relative market shares of wind versus solar preferred by the models.

Beyond the arc of technology costs, many of the other key market factors that will influence the impact of

the IRA incentives are measurable, plausibly exogenous to the evolution of the wholesale electricity market, and feature long histories of data. It is therefore viable to generate statistically rigorous forecasts of the joint distributions of these market factors. We present a summary of our data and the results of these forecasts in the following section and provide a detailed description of our forecasting methodology in Appendix A. With this methodology we generate a distribution of 1000 forecast 'draws' of market conditions for the year 2030. We then apply each of these forecast draws as inputs into the electricity market investment model to generate a distribution of model outputs.

4 Data and Forecasts

As described in detail in Appendix A, we use a dynamic factor model (DFM). In past work, the authors and other researchers (e.g., Borenstein et al. (2019); Bushnell et al. (2023)) have used cointegrated vector error correction models (VECM). This approach entails choosing a relatively small number of variables to forecast. These variables typically exhibit long-run trends (e.g., GDP increases annually by about 3% on average over time), which this approach models using a random walk (or unit root) process with a drift. This means that each variable tends to increase (or decrease) by a particular amount each period, but it does not tend to follow a fixed trend line. For example, if GDP grows by more than its 3% average this year, the model would not predict smaller growth next year to bring the series back to a trend line. Rather the model would predict continued 3% growth. The model uses cointegration to connect the long-run trends of each variable to each other.

The VECM has several drawbacks. First, it can handle only a small number of variables. This drawback precludes using the model to produce forecasts separately by economic sector or region. It also means that the researcher must choose one variable to capture, for example, the state of the economy, whereas numerous variables exist to measure various aspects of the economy or economic activity in various locations. The dynamic factor model overcomes this challenge by combining a potentially large number of variables into a small number of factors.

The DFM also allows for the model to provide more granularity to the forecast. For example we generate separate forecasts for electricity demand and natural gas prices for different regions of the western grid.

We use six sets of variables covering the Western United States. The data are quarterly and span 1997-2022. All variables enter the model in natural logarithm except temperature variables. The variable groups are:

1. real corporate bond yield
2. real personal income by state

3. real natural gas price by state
4. electricity sales (consumption) by state
5. peak electricity load by region
6. temperature by state

We use only one bond yield variable because interest rates do not vary by state. To capture real yields, we use Moody's Seasoned Aaa Corporate Bond Yield minus 10-year expected inflation. The corporate bond yield series is based on bonds with maturities of at least 20 years. The expected inflation series is produced by the Federal Reserve Bank of Cleveland using a model based on Treasury yields, inflation data, inflation swaps, and survey-based measures of inflation expectations.

For 11 western states, we obtain total personal income from FRED and city gate natural gas prices from EIA. We adjust these variables for inflation using the all items CPI so that the variables are measured in 2022 dollars. We obtain from EIA Form 861M the quantity of electricity sold to customers in each of 11 western states. We also obtain peak load in four distinct regions of the WECC. Finally, we obtain PRISM temperature data from Aaron Smith's Ag Data, including maximum temperature, degree days below 18°C, and degree days above 18°C at the state level.

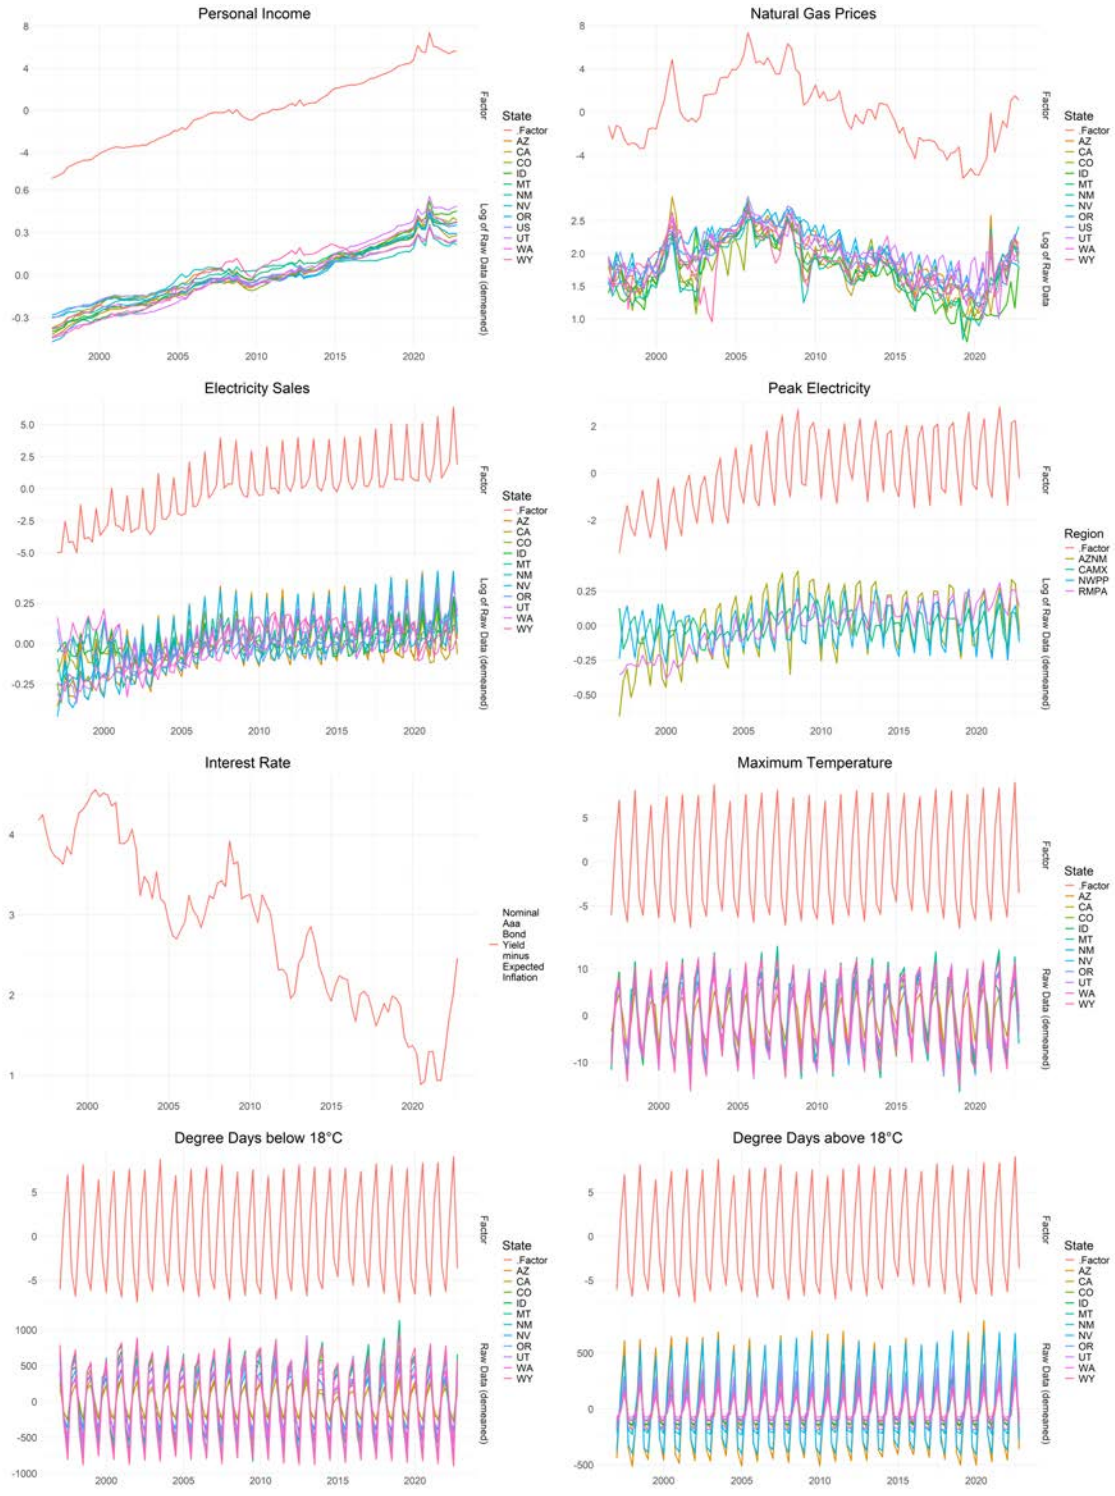

Figure 4: Data and Factors

Notes: For each variable group, bottom panel shows the raw data (demeaned where indicated) for each state or region. The top panel shows the factor for that group. For personal income, natural gas prices, electricity sales, and peak electricity, the factor is the first principal component of the variables in the lower panel. The three temperature variables have a single factor which is the first principal component of all three sets of variables. The real interest rate is its own factor.

Figure 4 shows the data separately by group along with the factor for each group. We take the natural logarithm of each variable except temperature variables. Because personal income, electricity sales, and peak load are substantially different across states, we demean the data to highlight on the same graph how they move together.

Personal income trends upwards in all states at an average of 2.3% per year. The great recession and the COVID-19 pandemic stand out in the personal income factor. The factor dropped significantly at the onset of the great recession and recovered very slowly. In 2020 and 2021, personal income spiked as government transfers buoyed incomes during the pandemic. Natural gas prices decreased from 2005-2020 as the hydraulic fracturing revolution lowered extraction costs, before increasing beginning in 2020. Individual state prices occasionally divert dramatically from the average, reflecting local gluts and shortages. The most notable spike is in Arizona and New Mexico in the first quarter of 2021.

Electricity sales and peak load trended upwards in the sample period, with faster growth before 2010 than after 2010. The early sample trends were much flatter in the northwest than the southwest. Both electricity variables also show prominent seasonality. In most states, electricity consumption is highest in the summer due to the demand for air conditioning. However, states in the northwest tend to have higher consumption in the winter. The peak factors are higher in summer than winter, reflecting the factor that more states peak in summer than winter. The interest rate trended down during our sample period, punctuated by spikes in the 2000 and 2008 recessions and in the 2022 inflation.

Temperature data exhibit clear seasonality. For instance, maximum temperature has spikes and troughs in the summer and winter of each year, respectively. Likewise, degree days above 18°C has a similar pattern whereas degree days below 18°C has the opposite pattern by definition. Their common factor captures this seasonality pattern. There is no clear but a slightly upward trend in the temperature data though.

Figure 5 shows the six factors and our forecasts of them. We generate quarterly forecasts, but plot annualized forecasts for clarity. The random walk component (ϕ_t) in equation (4) causes the forecast ranges to fan out the further into the future we project. The dashed lines indicate the 5^{th} and 95^{th} percentile outcomes. The forecasts are in log units (except for temperature variables), so their magnitudes make sense only relative to the in-sample values.

The model modest increases in electricity sales and peak load. These variables were below trend for much of the last decade, and the model predicts a relatively steady increase in the first few years out of sample. Nevertheless, there is sizable uncertainty around this prediction as shown by the 5^{th} and 95^{th} percentiles.

The model projects interest rates to trend downwards, but by 2034 the median forecast remains above its 2020 low. The 95^{th} percentile interest rate forecast for 2034 is about the level in 2012 and 2014, which is lower than any year prior to 2012 and higher than the years after 2014. The median personal income

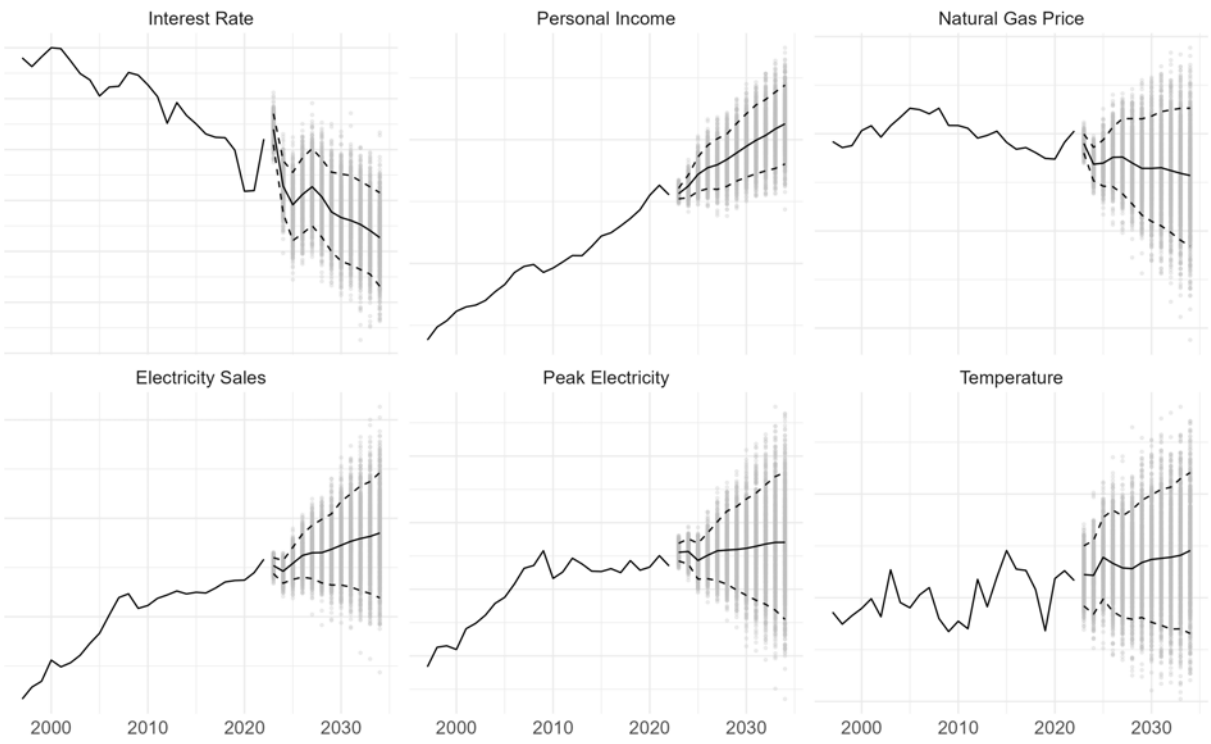

Figure 5: Forecasts of Factors

Notes: Sample period is 1997-2022. Forecast period is 2023-2034. Data observed quarterly. The factors are in log units. Forecast model is a VAR with four lags, and linear trend and seasonal dummy variables. Gray dots show values from 1000 draws of potential future paths, the black line denotes the median draw and the dashed line denotes the 5th and 95th percentiles.

forecast grows somewhat slower than in prior years, which is consistent with interest rates declining at a slower rate than in prior years. Median natural gas prices are forecast to be flat after a quick decline back to 2019 levels; the range of forecasts includes both increasing and decreasing prices.

Figure 6 shows forecasts electricity sales by state. For states such as California and Arizona that experienced a slowdown in the last decade, the model predicts a substantial increase, especially in the first few years of the forecast period. This result demonstrates a reversion to the factor trend. The median forecast for California is an increase in electricity sales from 250m TWh in 2022 to about 320m TWh in 2034, an increase of just over 25%. However, the range from the 5^{th} to the 95^{th} percentile is 275m-355m TWh, which reflects the large uncertainty in the forecast.

Figures 11-14 in the Appendix show forecasts for the other variables. Together, these plots show how the dynamic factor model can produce predictions of an array of variables based on a small number of factors.

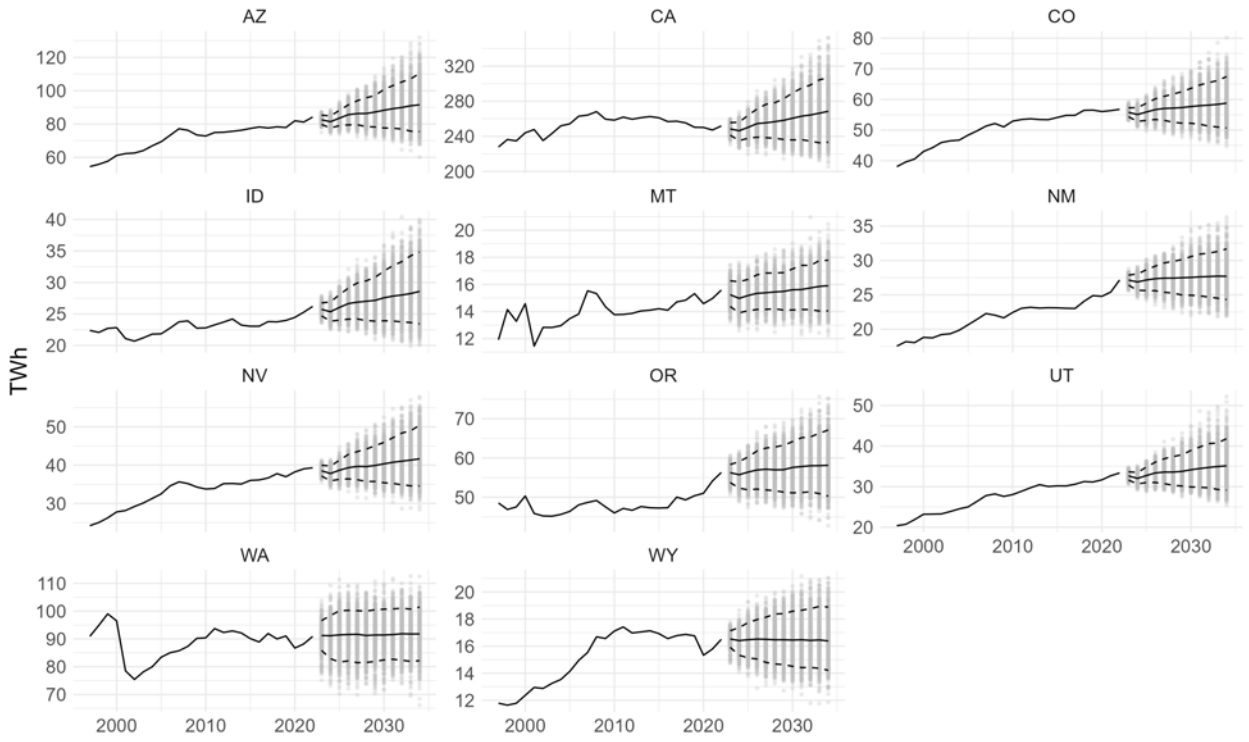

Figure 6: Forecasts of Electricity Sales

Notes: Sample period is 1997-2022. Forecast period is 2023-2034. Data observed and modeled quarterly, but plotted annually. Model estimated using log of data. Forecast model is a regression of the relevant state-level variable on four of its own lags and the electricity sales factor. Gray dots show values from 1000 draws of potential future paths, the black line denotes the median draw and the dashed line denotes the 5th and 95th percentiles.

5 Results

We combine the forecasts generated by the methods described in Section 4 with the electricity market equilibrium model described in Section 3 to simulate a hypothetical electricity market equilibrium in the year 2030. Starting with the model that has been calibrated to 2019 market conditions, we inflate (or deflate) natural gas prices, electricity demand in both the average and peak days, and the real interest rate according to our forecast prediction of these values for 2030. Each forecast draw produces a joint prediction of these values, and we run the simulation for each of 250 joint realizations.

We simulate this hypothetical 2030 market assuming no IRA incentives are in place, and then with the assumption that all new wind and solar investments qualify for $26/MWh PTC, and that all battery storage investments qualify for a 30% investment tax credit. The resulting resource disposition over the hours of our representative days are illustrated in Figure 7. This figure summarizes simulated 2030 output by region and hour of day across several important resource types. For each hour of the day, the left-hand bars in Figure 7 present output under a simulation of 2030 without PTC or ITC incentives. The right-hand bars summarize

average hourly output from a 2030 simulation in which PTC or ITC incentives are applied to all new wind, solar, and battery investments.

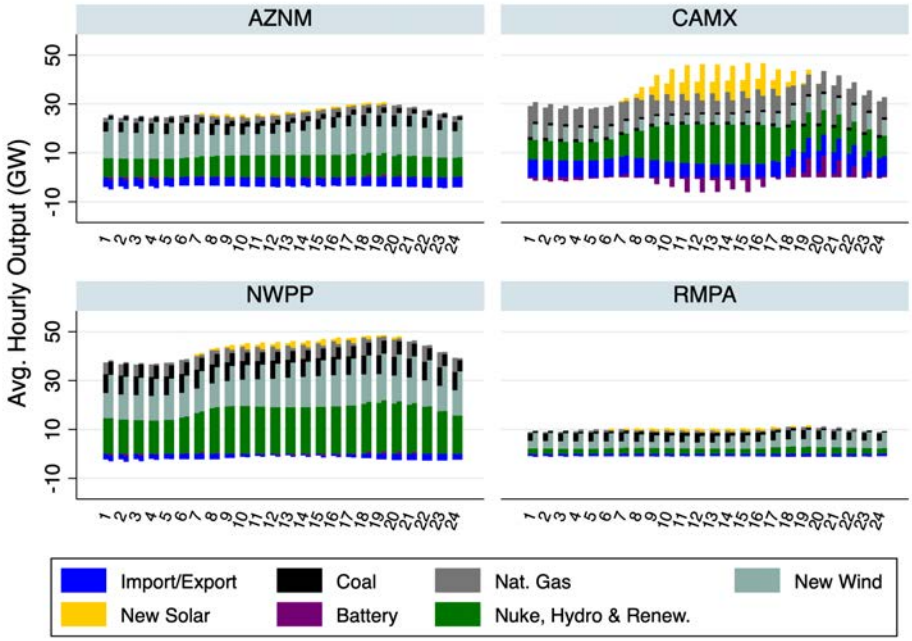

Notes: Figure summarizes simulated average generation by resource type and region, and hour of day for 2030. The left-hand bars for each hour assume no IRA incentives and the right hand bars assume IRA tax incentives are applied. The region definitions are Arizona and New Mexico (AZNM), California (CAMX), Pacific northwest (NWPP) and Rocky Mountains (RMPP).

Figure 7: Impact of PTC and ITC Incentives on Simulated 2030 Equilibrium

The results illustrated in Figure 7 utilize the mean forecast values for the relevant market parameters (peak demand, natural gas price, etc.). However the focus of this exercise is to capture the range of possible outcomes bounded by the range of the joint distribution of market variables in our forecast model. We first examine the sensitivity of the IRA impacts to variation in key inputs. Figure 8 illustrates the impact of the PTC on investment in new capacity of the various resources. The new wind generation capacity induced by the PTC (e.g. built in model runs where the PTC is applied but not built in the absence of the PTC) is far larger at realizations of low natural gas prices. The intuition is straightforward: at very low gas prices, absent the PTC subsidy, new CCGT capacity is preferred based upon pure economic considerations while at high natural gas prices, wind investment crowds out new and even existing gas generation regardless of the PTC incentive. In other words, most of the new wind capacity earning the PTC in a hypothetical future with high natural gas prices is not additional with respect to the PTC. It would have been built anyway.

While wind capacity dominates outcomes in the western grid, a similar qualitative point emerges for solar and batteries as well.

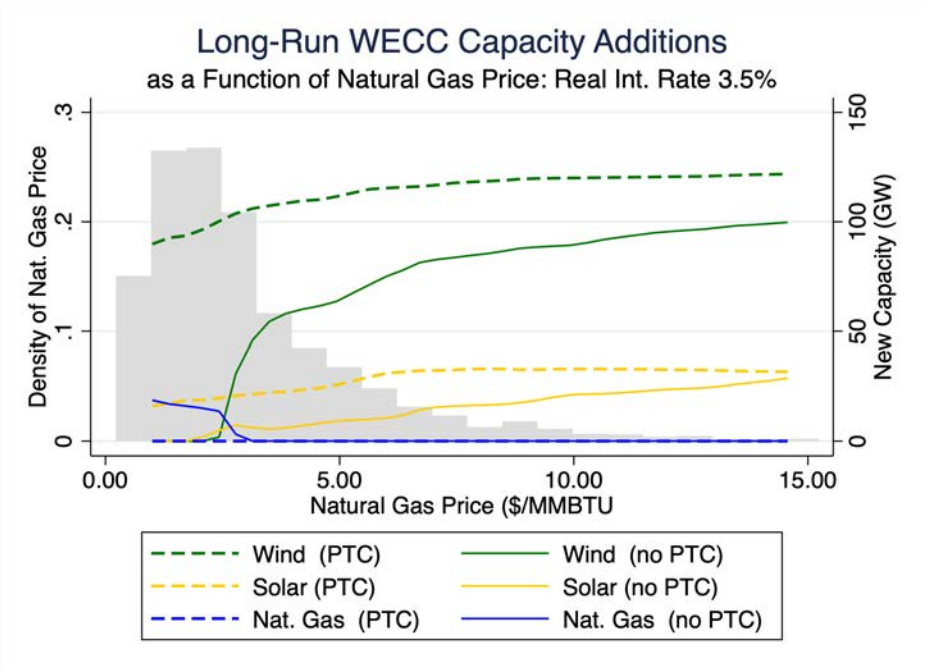

Notes: The histogram plots the forecast distribution of the future natural gas price. Lines illustrates investment in new generation capacity for 2030 by technology for a range of realized natural gas price. Dashed lines illustrate investment assuming IRA tax incentives are in place. Solid lines illustrate investment if no IRA tax incentives are in place.

Figure 8: 2030 New Generation Capacity by Fuel Type and Future Natural Gas Price

In order to assess the range of potential impacts of the IRA, we evaluate the distribution of several outcome variables from the market simulation. These outcomes include annual energy (TWh) and new capacity (GW) by fuel type, and annual carbon emissions for the WECC. Figure 9 summarizes the range of new investment and energy output from gas, solar, wind, and battery sources emerging from the simulation of 2030. The energy values for wind and solar are only for newly built resources, whereas energy output from natural gas plants includes existing resources. Intuitively, results that include PTC incentives yield substantially more wind capacity and energy, and a concurrent reduction in natural gas output.

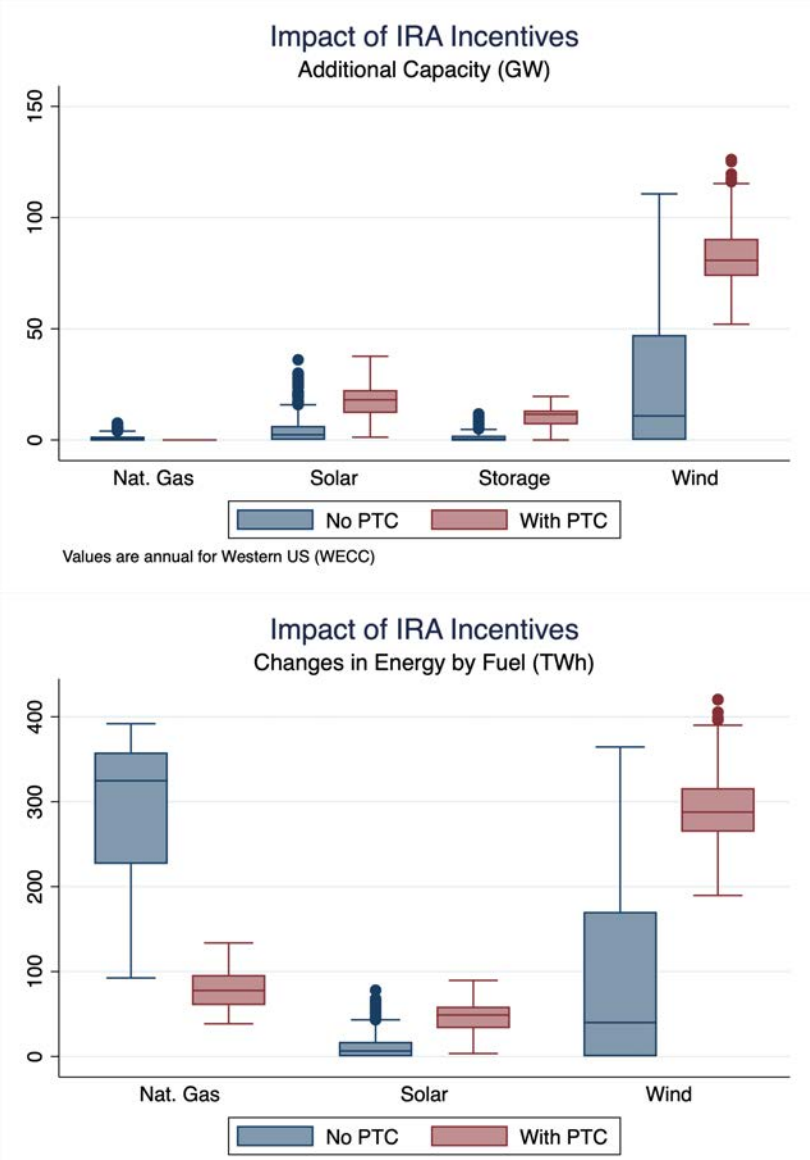

Notes: Top panel reports simulated capacity additions in the WECC region for 2030. The lower panel reports energy output (in TWh) from existing and new natural gas plants, new solar, and new wind plants. The solid bar ranges from 25th to 75th percentile of simulation results. Upper and lower whiskers represent the 75th (or 25th) percentile value plus (or minus) 1.5 times the interquartile range. Dots are outlier values.

Figure 9: Impact of PTC and ITC Incentives on Simulated 2030 Equilibria Outcomes

The most consistent result of this exercise is the fact that, in addition to shifting the mean outcome, IRA incentives narrow the range of many of these outcomes. This is best illustrated in Figure **??**, which summarizes the range of 2030 western CO2 emissions from electricity generation with and without the IRA incentives. Whereas grid emissions range from around 120 to over 275 MMT in the absence of a PTC, this range narrows substantially to between roughly 40 to 75 MMT when PTC incentives are included in the simulation. This effect is largely due the fact that the PTC has its greatest impact when market factors, such as interest rates and natural gas prices, are least favorable to renewable energy investments. By contrast, the PTC has a more modest impact on CO2 reductions under market conditions that would likely induce substantial renewable energy investment regardless of the IRA incentives.

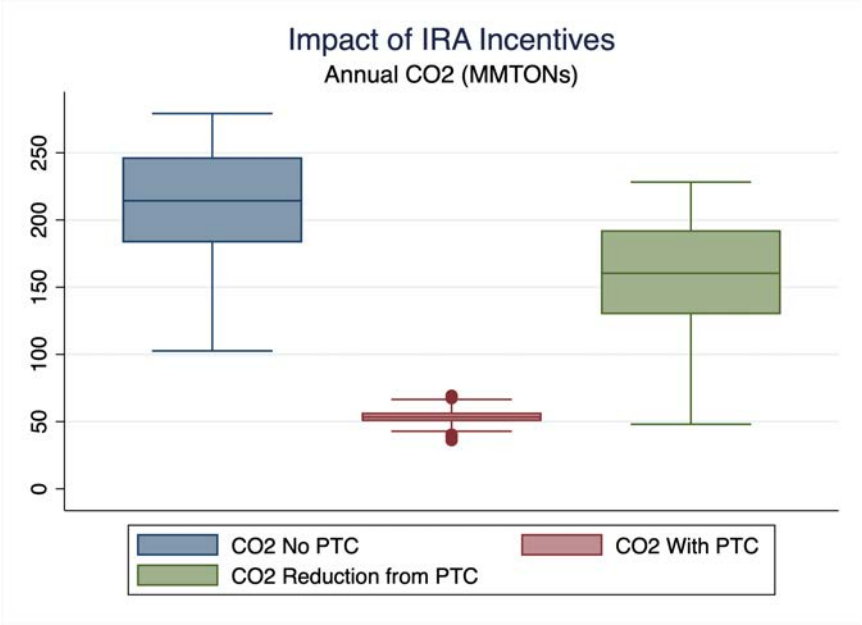

Notes: Figure reports simulated annual emissions in the WECC region for 2030. The solid bar ranges from 25th to 75th percentile of simulation results. Upper and lower whiskers represent the 75th (or 25th) percentile value plus (or minus) 1.5 times the interquartile range. Dots are outlier values.

Figure 10: Impact of PTC and ITC Incentives on 2030 CO2 Emissions

6 Conclusion

The need for analysis of the prospective mid to long-term impact of climate policies on the energy sector raises substantial challenges for empirical economists. The typical methods that the field has continually improved and refined over decades are largely designed for backward looking analysis, such as measuring the impacts of weather or energy prices on various outcomes such as mortality, economic activity, or sector-wide emissions. Many excellent papers have deployed these methods that provide insight into the "world as it is." However, analysis of proposed climate policy often requires an assessment of a future that is both "out of sample" and subject to non-marginal changes. These challenges can restrict the value of typical micro-econometric methods.

A growing number of increasingly complex simulation models has come to fill the void left by the weaknesses of traditional microeconomic models. These models offer the advantage of being able to calculate solutions for a wide range of hypothetical futures, and to be able to provide an almost absurd amount of detailed output data. These advantages come at the cost of transparency, replicability, and the strong influence of often subtle and opaque assumptions.

The detailed nature of these models can also convey a false sense of precision. Optimization models provide a precise answer to one specific problem, with a specific set of input assumptions. Unlike statistical models, they are not designed to measure uncertainty or convey distributional information such as confidence intervals.

In this paper we take a first step toward trying to address this latter weakness by combining statistical forecasts of many key inputs with an optimization model that simulates operations and investment in the western U.S. electricity system. The context for our analysis is the impact of selected elements of the 2022 Inflation Reduction Act. We simulate annual market outcomes of the western electric grid in 2030, comparing results that include the production and investment tax credits provided by the IRA with results that assume those incentives are not available. The analysis is conducted for a wide range of prospective realizations of key input variables including real natural gas prices, real interest rates, and electricity demand. These realizations of 2030 input variables are generated by a dynamic factor model that incorporates state level data on many of these inputs and provides geographic variation in input forecasts.

Our results indicate that, in addition to substantially changing the mean value of key outcomes such as renewable energy production and overall carbon emissions, the IRA incentives also substantially compress the range of these outcomes. The intuition behind this latter result is that the IRA is most impactful under future conditions least favorable to renewable energy investment. While the IRA incentives yield only modest carbon reductions under some potential future conditions, they substantially curtail emissions under the 'tail'

of outcomes in which natural gas generation sources would otherwise enjoy substantial market advantages.

This exercise is meant to be as much methodological as empirical. Our analysis is far from a comprehensive study of the full package of provisions in the IRA. Many other studies have examined the IRA with more scope in both geographic and economic reach. We do believe that this work illustrates the value of incorporating more rigorous consideration of uncertainty, at least the modeling elements for which reasonable statistical forecast models can be developed.

References

Bistline, J., Blanford, G., Brown, M., Burtraw, D., Domeshek, M., Farbes, J., Fawcett, A., Hamilton, A., Jenkins, J., Jones, R., et al. (2023a). Emissions and energy impacts of the inflation reduction act. *Science*, 380(6652):1324–1327.

Bistline, J., Mehrotra, N., and Wolfram, C. (2023b). Economic implications of the climate provisions of the inflation reduction act. *Brookings Papers on Economic Activity*, 1(5):77–182.

Bistline, J. E., Brown, M., Domeshek, M., Marcy, C., Roy, N., Blanford, G., Burtraw, D., Farbes, J., Fawcett, A., Hamilton, A., et al. (2023c). Power sector impacts of the inflation reduction act of 2022. *Environmental Research Letters*, 19(1):014013.

Borenstein, S., Bushnell, J., Wolak, F. A., and Zaragoza-Watkins, M. (2019). Expecting the unexpected: Emissions uncertainty and environmental market design. *American Economic Review*, 109(11):3953–3977.

Bushnell, J. and Chen, Y. (2012). Allocation and leakage in regional cap-and-trade markets for co2. *Resource and Energy Economics*, 34(4):647–668.

Bushnell, J., Chen, Y., and Zaragoza-Watkins, M. (2014). Downstream regulation of co2 emissions in california's electricity sector. *Energy Policy*, 64:313–323.

Bushnell, J., Lade, G., Smith, A., Witcover, J., and Xiao, W. (2023). Forecasting credit supply demand balance for the low-carbon fuel standard program. Technical Report WP-340, University of California Energy Institute.

Bushnell, J. B., Holland, S. P., Hughes, J. E., and Knittel, C. R. (2017). Strategic policy choice in state-level regulation: the epa's clean power plan. *American Economic Journal: Economic Policy*, 9(2):57–90.

Bushnell, J. B., Mansur, E. T., and Saravia, C. (2008). Vertical arrangements, market structure, and competition: An analysis of restructured us electricity markets. *American Economic Review*, 98(1):237–66.

CBO (2022). Estimated budgetary effects of h.r. 5376, the inflation reduction act of 2022. Technical report, Congressional Budget Office.

Krugman, P. (2022). Did the democrats just save civilization? *The New York Times*.

Müller, U. K. and Watson, M. W. (2018). Long-run covariability. *Econometrica*, 86(3):775–804.

Steinberg, D. C., Brown, M., Wiser, R., Donohoo-Vallett, P., Gagnon, P., Hamilton, A., Mowers, M., Murphy, C., and Prasana, A. (2023). Evaluating impacts of the inflation reduction act and bipartisan infrastructure law on the u.s. power system. Technical Report NREL/TP-6A20-85242, National Renewable Energy Laboratory.

Stock, J. H. and Watson, M. W. (2001). Vector autoregressions. *Journal of Economic Perspectives*, 15(4):101–115.

Stock, J. H. and Watson, M. W. (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In *Handbook of Macroeconomics*, volume 2, pages 415–525. Elsevier.

Wolfram, C. D. (1999). Measuring duopoly power in the british electricity spot market. *American Economic Review*, 89(4):805–826.

Appendix

A.1 Dynamic Factor Model

We observe a potentially large number (N) of variables X_{jt}. The dominant trend and correlations among these variables can be captured by a small number (r) of latent factors, which themselves evolve according to a vector autoregression model (VAR). Collecting the X_{jt} variables into a vector X_t, we can write the DFM as

$$X_t = \mu + \Lambda F_t + u_t \tag{1}$$

$$F_t = \sum_{\ell=1}^{L} A_\ell F_{t-\ell} + B_0 + B_1 t + C s_t + \varepsilon_t \tag{2}$$

where F_t denotes the vector of factors t is a time trend, and s_t are seasonal dummy variables. The "idiosyncratic component" u_t represents the deviations of the observed variables from the factors. This term may be be serially correlated. We assume that the VAR includes sufficient lags (L) such that the shocks (ε_t) are serially uncorrelated.

To fit this model and generate forecasts from it, we take four steps:

1. obtain the factors from X_t
2. fit a VAR for the factors as in (2)
3. forecast the factors using the VAR
4. generate forecasts of the variables X_{jt} using the forecasts of the factors

Next, we describe how we implement these steps.

A.1.1 Obtaining the factors

We obtain the factors using principal components analysis. There are multiple ways to normalize the factors depending on the purpose of the analysis. One way, known as principal components normalization, is to estimate the factors as $\hat{F}_t^{PC} = N^{-1}\hat{\Lambda}^{PC\prime}X_t$, where $\hat{\Lambda}^{PC}$ contains the eigenvectors associated with the largest r eigenvalues of the covariance matrix of X_t. The resulting factors are orthogonal to each other and capture as much as possible of the variance of X_t.

This approach is not invariant to the scale of the variables in X_t. To make the factors invariant to scale, we could use eigenvectors of the correlation matrix instead of the covariance matrix of X_t. Principal components normalization, whether computed from the covariance or correlation matrix, produces factors that are not easily interpretable and therefore cannot easily be used for structural analysis.

An alternative to aid interpretation is "named factor normalization," which imposes the structure

$$\hat{\Lambda}^{NF} = \begin{bmatrix} I_r \\ \hat{\Lambda}^{NF}_{r+1:N} \end{bmatrix} \tag{3}$$

Each of the first r variables in X_t is indexed to a factor, which enables the factors to be named by the variable they are indexed to. The named factors can be computed as $\hat{F}^{NF}_t = \hat{\Lambda}^{PC}_{1:r} \hat{F}^{PC}_t$ and the corresponding loadings (coefficients) as $\hat{\Lambda}^{NF} = \hat{\Lambda}^{PC} \left(\hat{\Lambda}^{PC}_{1:r}\right)^{-1}$, where $\hat{\Lambda}^{PC}_{1:r}$ denotes the first r rows of $\hat{\Lambda}^{PC}$. The resulting factors are correlated with each other, unlike in principal components normalization. Importantly, this method only affects the normalization of the factors; they span the same space however they are normalized.

We use a hybrid of the two approaches. We organize the variables X_{jt} into groups based on their role in the model. For example, in a multi-region model, a group may be economic activity and the X_{jt} variables in that group would be total income in each region; another group could be gas prices and the X_{jt} variables in that group would be the gas price in each region. The factor for group g is the first principal component of the variables in that group. In words, it is the linear combination of the observed variables that has the highest correlation with the set of variables. This approach is more restrictive than "named factor normalization" because it allows only variables in the group to determine the factor. A benefit of this restriction is that it more strongly aligns each factor with a defined set of variables.

To obtain the first principal component, we standardize each variable in group g by subtracting its mean and dividing by its standard deviation. Then the factor for that group is a weighted average of the standardized variables, where the weights are the first eigenvalue of the correlation matrix of the variables.

A.1.2 VAR on the factors

We model the VAR as a trend stationary process using a linear trend function. An alternative approach would be a cointegration model, which we discuss further in Section A.2.1. The VAR coefficients (A_ℓ) capture how long the variables take to revert back to trend after deviating from it and how the deviations from trend correlate across the variables. The model allows the factors to differ by quarter through the seasonal dummy variables. We estimate the model parameters by ordinary least squares.

A.1.3 Forecasting the factors

We use the coefficient estimates from the VAR to predict the distribution for each factor through the period under study, 2023–2034. Rather than generating a single forecast of each factor, we generate 1000 different potential trajectories based on the models. This approach allows us to characterize the range of potential

outcomes and their likelihood. To this end, we assume that the potential shocks (ε_t) that may occur in the forecast period have the same distribution as the shocks during our estimation sample period, 1997–2022. Using this assumption, we simulate potential future shocks by sampling randomly with replacement from the 1997–2022 shocks. For each random draw, we use the VAR model to generate a hypothetical path for the five factors. We repeat this exercise 1000 times to give a distribution of potential paths.

The VAR for the factors in (2) is trend stationary, which means that long-run projections from it will be driven by the linear trend. The long-run forecast of the factors is $F_t = \left(I - \sum_{\ell=1}^{L} A_\ell\right)^{-1} (B_0 + B_1 t + C s_t)$. However, the trends may not remain constant over time. To allow for variation in the forecast paths due to changing trends, we add a random walk component to (2) when generating future trajectories. Specifically, we generate potential future values of the factors using the model

$$F_t^* = \sum_{\ell=1}^{L} \hat{A}_\ell F_{t-\ell}^* + \hat{B}_0 + \hat{B}_1 t + \hat{C} s_t + \phi_t + \varepsilon_t^* \tag{4}$$

$$\phi_t = \phi_{t-1} + \eta_t \tag{5}$$

where ε_t^* denotes a random draw of the residual vector and $\eta_t \sim N(0, \theta B)$ and B is a diagonal matrix with diagonal elements equal to $\hat{B}_1^2$. Thus, the standard deviation of the random walk component (ϕ_t) is proportional to the trend coefficient ($\hat{B}_1$). We select all ε_t^* terms together to preserve contemporaneous correlations across factors.

A.2 Forecasting the individual variables

To generate forecasts of the individual variables, we use an autoregressive model

$$X_{jt} = \sum_{\ell=1}^{L} \alpha_{j\ell} X_{j,t-\ell} + \beta_j F_{gt} + \gamma_j s_t + v_{jt} \tag{6}$$

This approach allows an individual variable to deviate from the path of the factor for multiple periods, i.e., it accounts for serial correlation in the idiosyncratic component (u_t). The speed at which it returns to the path of the factor depends on the magnitude of the coefficients $\alpha_{j\ell}$ in (6); the closer $\sum_{\ell=1}^{L} \alpha_{j\ell}$ is to one, the longer X_{jt} tends to take to revert back to the path of the factor. The fact that there are no trends in (6) implies that all trends in the data are driven by the factors. Individual series cannot trend away from their factor. This means, for example in a multi-region model, that one region cannot have long-run economic growth different from another region. The inclusion of seasonal dummy variables in (6) implies that an individual series may have different seasonality from the factor.

A benefit of our factor normalization approach is that only one factor enters the forecasting equation (6). The factors each help to forecast each other in the VAR, but then only the relevant factors enter into the forecasting equation. Given the relatively short time series in climate studies, parsimony is important.

To simulate future trajectories of the individual variables, we follow the same approach as above for forecasting the factors. When taking draws, we select all v_t terms together to preserve contemporaneous correlations across variables. We do not select ε_t and v_t together because the factor structure makes these terms uncorrelated in the model.

A.2.1 Comparison to Traditional Time Series Method: Vector Error Correction

In past work, the authors and other researchers (e.g.,Borenstein et al. (2019); Bushnell et al. (2023)) have used cointegrated vector error correction models (VECM). This approach entails choosing a relatively small number of variables to forecast. These variables typically exhibit long-run trends (e.g., GDP increases annually by about 3% on average over time), which this approach models using a random walk (or unit root) process with a drift. This means that the each variable tends to increase (or decrease) by a particular amount each period, but it does not tend to follow a fixed trend line. For example, if GDP grows by more than its 3% average this year, the model would not predict smaller growth next year to bring the series back to a trend line. Rather the model would predict continued 3% growth. The model uses cointegration to connect the long-run trends of each variable to each other.

The VECM has several drawbacks. First, it can handle only a small number of variables. This drawback precludes using the model to produce forecasts separately by economic sector or region. It also means that the researcher must choose one variable to capture, for example, the state of the economy, whereas numerous variables exist to measure various aspects of the economy or economic activity in various locations. The dynamic factor model overcomes this challenge by combining a potentially large number of variables into a small number of factors.

Another drawback of the VECM is that the long-run, or cointegrating, relationships cannot generally be interpreted as structural economic equations. For example, a transportation model may correlate gasoline prices and consumption, which is why it forecasts well, but is not specific about whether the relationship quantifies supply or demand. This shortcoming makes counterfactual policy analysis difficult because we cannot model, for example, the response of consumers to an increase in fuel prices because no equation represents consumer demand. The DFM presented in this section is not set up explicitly to do counterfactual policy analysis, but we plan to add this feature following the approach of Stock and Watson (2016).

In addition, cointegration implies that variables are tied together in the long-run. For example, fuel prices may be tied to vehicle miles traveled because drivers may reduce their miles if fuel prices are higher.

However, these two variables may still have divergent long run trends because of other factors such as growing incomes, improving fuel efficiency, and changing propensity to work from home. So, we specify the trends exogenously and use the VAR to explain variation around the trends. This will work when we come to analyze policy shocks, such as a fuel tax, because the model will capture the response of VMT to fuel prices without trend confounding the response. Finally, climate applications are characterized by relatively short time series. In our application, we use 26 years of data.

For these reasons, we model the long-run trends as deterministic functions (Müller and Watson (2018)), but we allow the trends to evolve out of sample as detailed in (4).

Figures 11-15 show forecasts for the variables other than electricity sales.

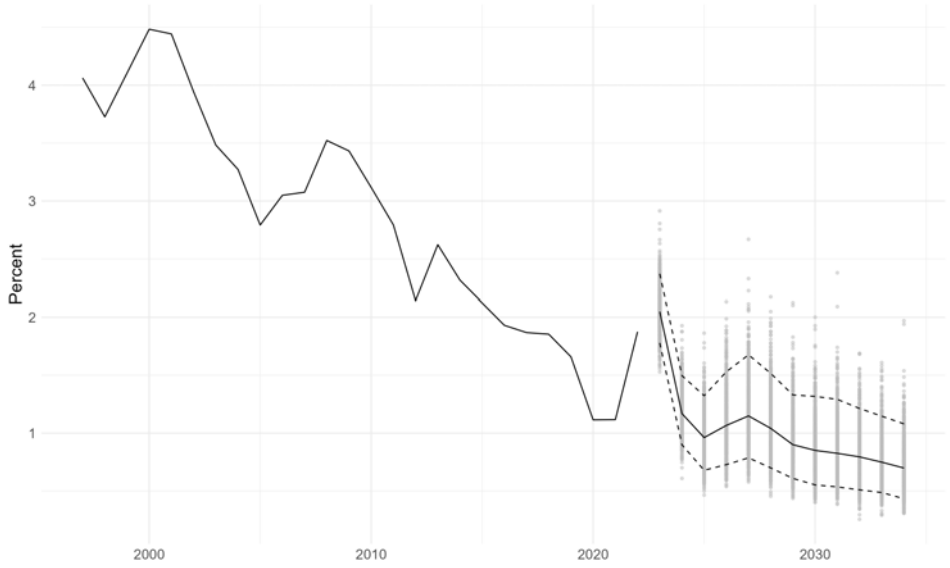

Figure 11: Forecasts of Corporate Bond Yield

Notes: Sample period is 1997-2022. Forecast period is 2023-2034. Data observed and modeled quarterly, but plotted annually. Model estimated using log of data. Forecast is directly from factor VAR because the interest rate is its own factor. Gray dots show values from 1000 draws of potential future paths, the black line denotes the median draw and the dashed line denotes the 5th and 95th percentiles.

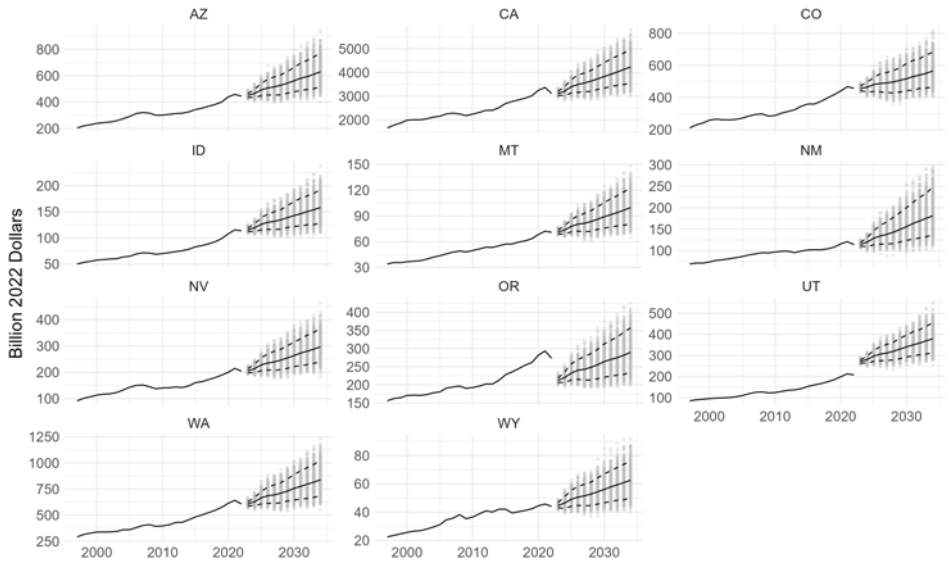

Figure 12: Forecasts of Personal Income

Notes: Sample period is 1997-2022. Forecast period is 2023-2034. Data observed and modeled quarterly, but plotted annually. Model estimated using log of data. Forecast model is a regression of the relevant state-level variable on four of its own lags and the electricity sales factor. Gray dots show values from 1000 draws of potential future paths, the black line denotes the median draw and the dashed line denotes the 5th and 95th percentiles.

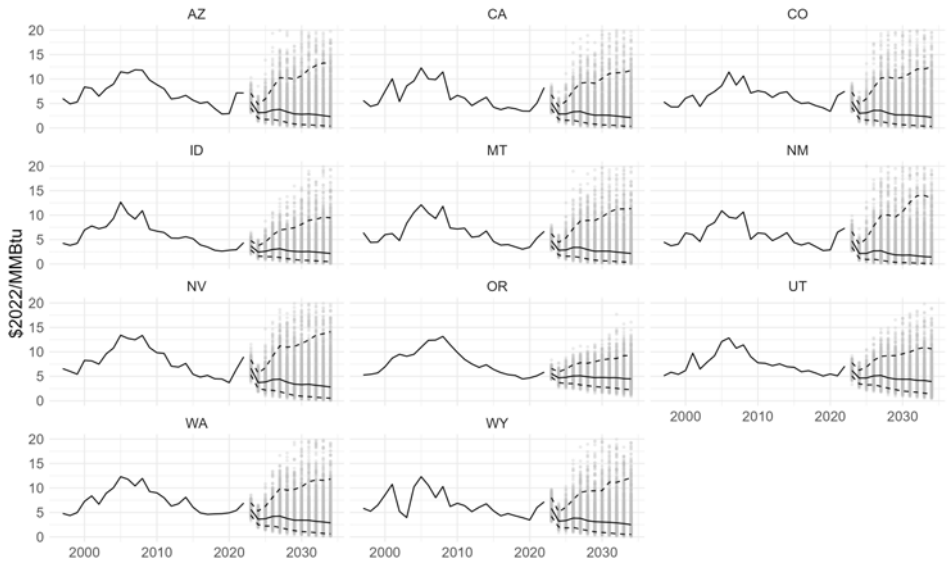

Figure 13: Forecasts of Natural Gas Prices

Notes: Sample period is 1997-2022. Forecast period is 2023-2034. Data observed and modeled quarterly, but plotted annually. Model estimated using log of data. Forecast model is a regression of the relevant state-level variable on four of its own lags and the electricity sales factor. Gray dots show values from 1000 draws of potential future paths, the black line denotes the median draw and the dashed line denotes the 5th and 95th percentiles.

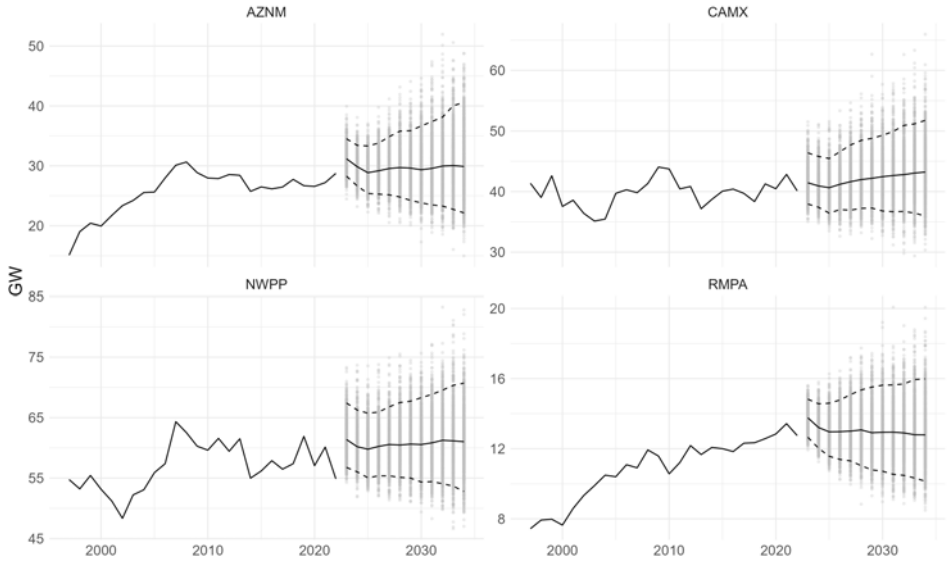

Figure 14: Forecasts of Peak Electricity Load

Notes: Sample period is 1997-2022. Forecast period is 2023-2034. Data observed and modeled quarterly, but plotted annually. Model estimated using log of data. Forecast model is a regression of the relevant region-level variable on four of its own lags and the electricity sales factor. Gray dots show values from 1000 draws of potential future paths, the black line denotes the median draw and the dashed line denotes the 5th and 95th percentiles.

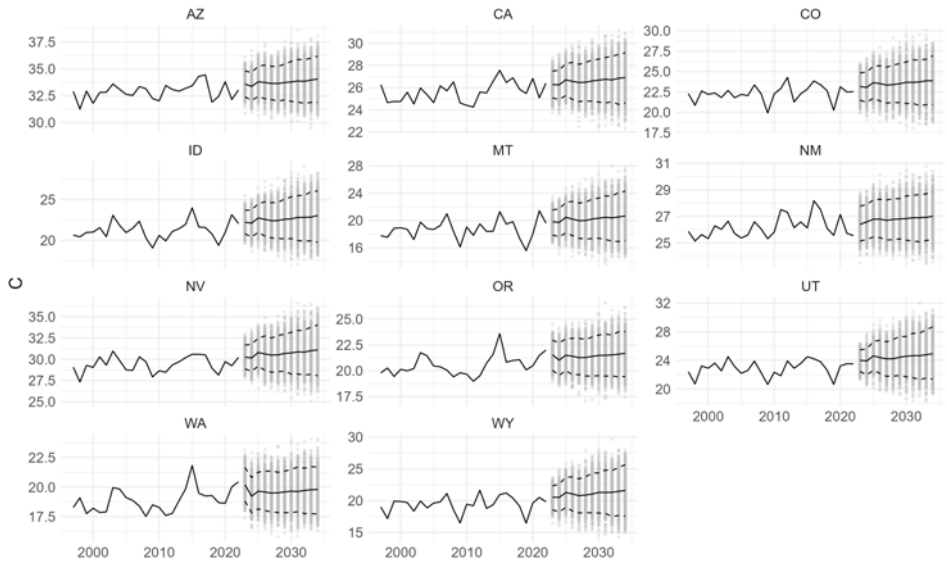

Figure 15: Forecasts of Max Temperature

Notes: Sample period is 1997-2022. Forecast period is 2023-2034. Data observed and modeled quarterly, but plotted annually. Model estimated using log of data. Forecast model is a regression of the relevant state-level variable on four of its own lags and the electricity sales factor. Gray dots show values from 1000 draws of potential future paths, the black line denotes the median draw and the dashed line denotes the 5th and 95th percentiles.

B.1 Details of electricity market investment model

B.1.1 Model Specification

The model includes several standard features of electricity system operations and investment. It simultaneously optimizes the dispatch of generation and storage for a sequence of eight 24-hour representative days. Demand is assumed to be linear but with low elasticity and is aggregated to the regional level. The objective function is to maximize social welfare, defined as the weighted sum of hourly consumption less generation costs, including annualized investment costs F_j described below.[2]

$$\sum_t \left[\sum_r CB(p_r(t)) - \sum_{i \in r} c_i q_i(t) - \sum_{jinr} F_j^r * cap_j \right] Tweight_t$$

Where c_i represents the unit level marginal cost of unit i located in region r, $p_r(t)$ is the price in region r, and Tweight is the weight placed upon a given time period based upon the number of actual days in the calibration year it is representing. The production of each generation unit i is constrained to be less than or equal to the capacity of unit i. Within each representative day, the change in hourly output from baseload generation units (coal and combined cycle gas) is further constrained by unit level ramp rates. The model

[2]Consumer benefit is defined as the area under the demand curve out to the quantity consumed in a give hour and region. In other words it is consumer surplus with the cost of purchasing the electricity added back to reflect the fact that this purchase cost is a transfer to producers and not part of a social welfare calculation.

does not, however, impose integer "unit commitment" constraints. For existing generation units, capacity and ramping capabilities are calibrated based upon unit performance in 2019. For new generation, indexed by j, hourly output is also limited by capacity, but capacity for these units is a choice variable.

B.2 Market demand

To construct our demand functions, we assume linear demand that passes through the mean price and quantity for each representative time period and region. End-use consumption, as defined above, in each region is represented by the demand function $Q_{r,t} = \alpha_{r,t} - \beta_r p_{r,t}$, yielding an inverse demand curve defined as

$$p_{rt} = \frac{\alpha_{r,t} - \sum_i q_{rit} - y_{i,t}}{\beta_r}$$

where $y_{r,t}$ is the aggregate net imports into region r.

The parameter α_{rt} is calibrated so that, for a given β_r, $Q_{r,t}^{actual} = \alpha_{r,t} - \beta_r p_{r,t}^{actual}$. In other words, the demand curve is shifted so that it passes through the average of the observed price quantity pairs for that collection of hours. To derive actual demand, EIA 930 provides hourly total end-use consumption by control-area which we aggregate to the North American Electric Reliability Commission (NERC) sub-region level. For electricity prices, we use hourly market prices calculated as part of the western Energy Imbalance Market (EIM) and taken from SNL Global.

B.3 Fossil-fired generation costs and emissions

We explicitly model the major fossil-fired thermal units in each electric system. Because of the legacy of cost-of-service regulation, relatively reliable data on the production costs of thermal generation units are available. The cost of fuel comprises the major component of the marginal cost of thermal generation. The marginal cost of a modeled generation unit is estimated to be the sum of its direct fuel, CO_2, and variable operation and maintenance (VO&M) costs. Fuel costs can be calculated by multiplying the price of fuel, which varies by region, by a unit's 'heat rate,' a measure of its fuel-efficiency.

The capacity of a generating unit is reduced to reflect the probability of a forced outage of each unit. The available capacity of generation unit i, is taken to be $(1 - fof_i) * cap_i$, where cap_i is the summer-rated capacity of the unit and fof_i is the forced outage factor reflecting the probability of the unit being completely down at any given time.[3] Unit forced outage factors are taken from the generator availability data system (GADS) data that are collected by the North American Reliability Councils.

[3]This approach to modeling unit availability is similar to Wolfram (1999) and Bushnell, Mansur and Saravia (2008) .

B.4 Transmission network

Our regional markets are highly aggregated geographically. The region we model is the electricity market contained within the U.S. portion of the Western Electricity Coordinating Council (WECC). The WECC is the organization responsible for coordinating the planning investment, and general operating procedures of electricity networks in most states west of the Mississippi. The multiple sub-networks, or control areas, contained within this region are aggregated into four "sub-regions." Between (and within) these regions are over 50 major transmission interfaces, or paths. Due to both computational and data considerations, we have aggregated this network into a simplified 5 region network consisting primarily of the 4 major subregions.[4]

Mathematically, we adopt an approach utilized by Metzler, et al. (2003), to represent the transmission arbitrage conditions as another set of constraints. Under the assumptions of a direct-current (DC) load-flow model, the transmission 'flow' induced by a marginal injection of power at location l can be represented by a power transfer distribution factor, $PTDF_{lk}$, which maps injections at locations, l, to flows over individual transmission paths k. Within this framework, the arbitrage condition will implicitly inject and consume power, $y_{l,t}$, to maximize available and feasible arbitrage profits as defined by the difference in prices across locations.

Transmission models such as these utilize a "swing hub" from which other marginal changes in the network are measured relative to. We use the California region as this hub. In other words, an injection of power, $y_{r,t} \geq 0$, at location r is assumed to be withdrawn in California. The welfare maximization objective function is therefore subject to the flow limits on the transmission network, particularly the line capacities, T_k:

$$-\overline{T}_k \leq \sum_r PTDF_{r,k} \cdot y_{r,t} \leq \overline{T}_k.$$

Given the aggregated level of the network, we model the relative impedance of each set of major pathways as roughly inverse to their voltage levels. The network connecting AZNM and the NWPP to CA is higher voltage (500 KV) than the predominantly 345 KV network connecting the other regions.

For our purposes, we assume that these lower voltage paths yield 5/3 the impedance of the direct paths to CA. Flow capacities over these interfaces are based upon inter regional transfers reported in EIA 930, calculated as the sum of exports from control areas located within one of our 4 NERC subregions going to control areas in a neighboring subregion. Transfers between control areas that share a NERC subregion are

[4]The final "node" in the network consists of the Intermountain power plant in Utah. This plant is connected to southern California by a high-capacity DC line, and is often considered to be electrically part of California. However under some regulatory scenarios, it would not in fact be part of California for GHG purposes, it is represented as a separate location that connects directly to California.

Table 3: Regional Transfer Capacity

Source	Destination				Simulaneous
	AZNM	CA	NWPP	RMPA	Max/Min
AZNM	NA	5581	1131	940	4753
CAMX	5581	NA	1300	1310	-10248
NWPP	1216	6061	NA	803	5495
RMPA	921	NA	766	NA	1255

Table 4: Capacity and 2019 Mean Production of Incumbent Generation

	AZNM	CA	NWPP	RMPA
Capacity (MW)				
Coal	6271	1729	9680	5004
Natural Gas	11474	24212	12247	5067
Renewable, Hydro & Nuclear	11107	23109	25152	4131
Mean Houlry Production (MW)				
Coal	4630	1098	7792	3915
Natural Gas	6629	10509	8806	2229
Renewable, Hydro & Nuclear	8471	11445	17884	2370

considered internal to that subregion.

We summarize the inter-regional and simultaneous maximum flows and net injections in Table 3. These values reflect the average annual maximum observation between 2018 and 2022 taken from EIA 930. In addition to the limits on individual flows we also impose a seasonal "nomogram" style constraint on the total netflow into or out of a region. In other words, for each hour t in season s we constrain net injections from region r such that

$$ymin_{r,s} \leq y_{i,t} \leq ymax_{r,s}.$$

B.5 Hydro, renewable and other generation

Incumbent generation capacity and annual energy production for each of our regions is reported by technology type in Table 4. Fossil-fired generation (coal and natural gas) is taken from CEMS. Non-CEMS generation is taken from EIA 930, which reports hourly production by resource category (e.g. wind, solar, nuclear, etc.) for each control area. Regional demand is derived by aggregating production from CEMS and non-CEMS generation by NERC sub-region and subtracting hourly net-exports.

Investment in new capacity

Table 5: Average Capacity Factors of New Renewable Sources by State

	Wind	Solar
AZ	0.383	0.254
CA	0.377	0.311
CO	0.433	0.246
ID	0.312	0.268
MT	0.408	0.213
NM	0.389	0.278
NV	0.185	0.318
OR	0.255	0.271
UT	0.446	0.294
WA	0.277	0.262
WY	0.429	0.292

Equilibrium investment is market driven, and requires sufficient market revenues to cover the (annualized) capital costs of new generation. Formally, hourly production from generation plant i in state s is constrained to not exceed the product of the installed capacity of that plant and the plants hourly capacity factor (CF).

$$q_{sit} \leq CAP_{si} * CF_{sit} \forall i, t.$$

For natural gas technologies, CF reflects an outage rate and is identical for all hours. For solar and wind investment, CF is taken from the observed hourly production portfolios of those resources in each State during the calendar year of 2022. These values are summarized in Table 5

Overnight capital costs of new generation are given in Table 2. We calculate an annualized capital cost by applying the ammortization formula assuming a payback period of 20 years.

$$Annual\_Cost_{si} = OCC_{si} \frac{r(1+r)^n}{(1+r)^n - 1}$$

where OCC_{si} is the overnight capital cost of technology i in state s, r is the interest rate and n is the payback period. We assume that the interest rate is the forecasted real AAA bond rate plus 200 basis points to account for the additional risks and costs of project financing.

The social welfare maximization objective function in the model combines the annual capital costs with the annual production cost from all generation sources. Increasing the quantity of new renewable generation increases capital costs and decreases production costs from the increasingly displaced existing generation sources. At the optimal solution, the revenues earned by the marginal additional MW of renewable capacity from resource i in State s is equilibrated with the marginal annual capacity cost of that renewable technology.

$$\sum_t P_{st} * CFsit * Tweight_t = Annual\_Cost_{si}$$

.