On the Blue States' Race to the Top in Energy Policy: A Structural Model of Renewable Portfolio Standards

Bentley Coffey[1]
Visiting Assistant Professor
Economics Department
Moore School of Business
University of South Carolina

9.29.14

Abstract:

In the past dozen years, most states have adopted some form of a Renewable Portfolio Standard (RPS) mandating that utilities generate a larger fraction of energy from renewable sources. Although an RPS is quite effective at replacing relatively dirty generation from old coal-consuming plants with cleaner energy, it can also be quite costly. The immediate partial equilibrium effect of utilizing cleaner energy is an increase in energy prices that reduces a state's output (especially from energy-intensive industries). The larger general equilibrium effects would be broader price changes and more substitution but might also include the loss of business to states without an RPS, which might lead to an interesting and increasingly observed phenomenon: coalitions of states filing suit against the Environmental Protection Agency (EPA) for not regulating the emissions of fossil-fired power plants more aggressively. To explain the RPS adoption decision that could lead to this phenomenon, we develop a theoretical model that captures the partial equilibrium effects of an RPS in terms of lost income from the state's industries and attributes the benefits of an RPS to environmentally conscientious voters. We find evidence to support this model in an analysis of data on state voting patterns, state energy regulations, a battery of electric power sector details, and the effects of energy prices on state GDP (by industry) from 1997-2010.

[1] Correspondence contact: 3231 Wilmot Ave., Columbia, SC 29205. bentleygcoffey@gmail.com.

I. Introduction

Federal policymakers in the United States have had climate change on their radar for decades but have continually delayed taking sweeping action because the politics have been poisoned, the uncertainties undermining, the damages distant, and the costs colossal. Revolutionary reductions in greenhouse gas emissions would require a radical revision of environmental /energy policy from the energy efficiency of heating/cooling buildings to lighting and from the fossil fuel consumed in transportation to electric power generation. Dissatisfied with federal inaction, States dominated by greener Democrats have taken matters into their own hands. To reduce their own carbon footprint, many Blue States (as well as a handful of Red States) have pursued progressive policies such as setting an RPS (Renewable Portfolio Standard) that require each electric power utility to generate a minimum fraction of their energy from renewable sources. Observers generally consider an RPS to be both quite effective at elevating cleaner energy but also quite costly because the generation capital for cleaner energy is much more expensive.

In a world absent of interstate externalities, federalists would be rather satisfied with the outcome of greener states securing greater environmental quality via higher environmental standards. While these greener states are free to adopt more stringent environmental policies than those promulgated by the Environmental Protection Agency (EPA), we increasingly observe coalitions of states (usually Blue States) filing suit against the EPA to compel the federal government to further regulate the emissions of electric power utilities, which tends to decrease the cost-effectiveness of coal-fired generation in favor of natural gas and renewables. Idealists have ample reasons for supporting these lawsuits, which attempt to supplant a second-best solution (setting energy policy state-by-state) with a first-best solution (centralized policy that internalizes externalities with limited leakages). Although any such lawsuit against the EPA necessarily invokes direct harm from an interstate environmental externality, a cynical economist could suggest an interesting alternative explanation.

In an economy where states compete for economic development opportunities by landing firms that are deciding where to site their production, the energy price differential induced by an RPS could be pivotal. If economic development in energy-intensive industries declines in Blue States relative to Red States due to their electricity rising under an RPS, then Blue States may well sue EPA to compel Red States to further regulate their energy utilities in order to level the playing field. A general equilibrium model, sufficiently rich in detail to capture this business-stealing effect, might produce some interesting implications: the apparent polarization across states could become starker over time as energy-intensive industries take an increasingly dominant role in the economy of a state without an RPS, increased regulatory activity in Blue States would inadvertently decrease regulatory activity in Red States, and federal inaction could even be prolonged despite the present precipice upon which national energy policy teeters.

Developing a general equilibrium model that is capable of accommodating such a grand scale of RPS effects is just beyond the scope of this paper. Instead, we focus on the more immediate characterization of the partial equilibrium costs of an RPS and explaining the decision to pass an RPS. By applying a revealed preference argument to a state legislature's median voter, we can infer benefits of an RPS to environmentally conscientious voters. Our theoretical model captures the partial equilibrium effects of an RPS in terms of lost income from the state's industries and attributes the RPS's benefits to political observables. We find evidence to support this model in an analysis of data on state voting patterns, state energy regulations, a battery of electric power sector details, and the effects of energy prices on state GDP (by industry) from 1997-2010.

From this research, we learn about the determinants of an RPS mandate – both its costs and its benefits. In particular, we find that centrists opposed mandating renewables in 1997 (i.e. $\beta_1 < 0$), that leftists in 1997 highly valued an RPS in opposition to right-wingers (i.e. $\beta_3 > 0$), that people from all political stripes seem to increasingly value renewable mandates over time (i.e. $\beta_2 > 0$), and that conservatives have warmed up to renewable mandates over time while liberals have cooled off a bit. Most importantly, these trends are convergent so that we may actually be approaching a consensus. Further inspection suggests that red states are increasingly implementing RPS mandates, despite considerable polarization and ongoing opposition within these red states.

There is already a growing literature on RPS mandates. Chen, Wiser, and Bolinger (2007) attempt a broad characterization of costs and benefits of RPS mandates; our study engages in a somewhat similar endeavor but we are more focused on producing benefits estimates grounded in rigorous inference from data sets that can support identification with enough innovative groundwork. Prasad and Munch (2012) estimate the effect of an RPS mandate on carbon emissions but not the actual decision to pass a mandate, which is assumed to be exogenous. Lyon and Yin (2009) is closest to this study but it is a truly reduced-form analysis of whether an RPS is adopted instead of a structural modeling of the magnitude of the RPS mandate; moreover Lyon and Yin (2009) is an older study and hence could not take advantage of newly available data from this rapidly evolving area.

In addition to answering research questions about RPS mandates, this study also makes a methodological contribution. The estimation of a system of nonlinear equations derived from structural model by Bayesian MCMC that is robust to strong sample size limitations is not an innovation unique to this study (e.g. see examples from Gelman et al. 2004); however, the particularities of this setup are a novel contribution. More generally, this paper continues the effort of work like Delarue et al. (2011) and Sue Wing (2006) to insert more realistic technical detail in economic models (both empirical and theoretical) of energy markets.

The layout for the remainder of the paper is as follows. Section II develops the theoretical model motivating our empirical design, Section III describes the data that we have gathered to identify key elements of the theoretical model, and Section IV details the econometric methodology for identifying parameters of the structural model from the data. Section V presents results and Section VI concludes.

II. Theoretical Model

a. State Governments

The political preferences of voters are defined along a spectrum ($S \in [0,1]$) that describes how they trade-off their valuation of the non-market benefits of more generation from renewables against a hypothetical RPS' costs (in terms of decreased per-capita income). Legislators, at both the state and national levels, are assumed to have preferences that reasonably represent their constituents. Given that state legislatures ultimately make their decisions by majority rule, the median voter theorem applies: the legislature will try to pass the RPS most favored by the median voter.[2] However, because the actual legislative process can get exceedingly messy (e.g. agenda-setting and parliamentary rules such as the Senate's filibuster), the bargaining costs of legislative deal-making may be too great to pass an RPS even when it is perceived by the median voter as net beneficial. Hence, the state legislature will pass an RPS that mandates a growth rate (ρ) in the share of generation from renewable sources, if the benefits exceed the sum of the variable costs to the economy (increasing in ρ) and the fixed bargaining costs:

$$\left(\beta_1 + \left[t \otimes \left(\frac{1}{2} - med(S_{it})\right)\right]\begin{bmatrix}\beta_2\\ \beta_3\\ \beta_4\end{bmatrix}\right)\rho \geq \sum_{j=1}^{J} \Delta\left[\frac{\widehat{GDP}_{ijt}\left(P_{Eit}\left((1+\rho)\theta_{it}\right)\right)}{POP_{it}}\right] + \kappa\, pol(S_{it}) + \varepsilon_{3it}$$

Where $\beta_1\rho$ is the valuation of a true centrist, β_2 describes how the centrist's valuation increases over time (e.g. because people are becoming increasingly alarmed about climate change), β_3 describes how that valuation changes along the political spectrum (as measured by the score S, which is centered on 0.5), β_4 describes how the partisan valuation is changing over time, ε_{3it} is a random utility term that captures costs (e.g. a period of acrimony that raises bargaining costs) less idiosyncratic differences in benefits valuation (e.g. due to deviations in the representativeness of the legislators), GDP_{ijt} is the revenue earned by industry j in state i at time t (so that the sum across all industries is the state's income), POP_{it} is the population of state i, κ captures the fixed cost (i.e. not varying with ρ) associated with legislative deal-making (which is increasing in the polarization of political preferences), the $\otimes$ operator is a Cartesian product (not a Kronecker product), and Δ is an operator representing the difference in GDP resulting from increases in energy prices resulting from an RPS mandate to increase θ by a factor of $(1+\rho)$.

The first-order condition implicitly describes the value of ρ that equates marginal benefit to cost:

[2] Naturally, if the median voter perceives the costs as dominating the benefits even at the lowest levels of ρ, then the legislature would even take up an RPS in that session.

$$\beta_1 + \left[t \otimes \left(\frac{1}{2} - med(S_{it})\right)\right]\begin{bmatrix}\beta_2\\ \beta_3\\ \beta_4\end{bmatrix} = \frac{1}{POP_{it}}\left[\sum_{j=1}^{J}\frac{\partial GDP_{ijt}}{\partial P_E}\right]\frac{\partial P_E}{\partial \rho}$$

Where P_E is the price of electricity, that channels all of the economic effects of an RPS mandating ρ. In a general equilibrium framework, changes in energy prices would cascade to other prices throughout the economy. Hence, this model is a form of partial equilibrium where only adjusts to the shock of the legislature passing an RPS that mandates that renewables' share of generation grows at ρ.

b. Electric Power Industry

Each state is served by a single electric power utility that produces electricity from renewable sources (R_{it}) and electricity from non-renewables (N_{it}). Each type of generation (i.e. R_{it} and N_{it}) cannot exceed its respective capacity constraint: the stock of capital that generates from renewable sources (K_{Rit}) and the stock of capital that generates from non-renewable sources (K_{Nit}). Because this electric utility is a monopolist, it is subject to a variety of Public Utilities Commission (PUC) regulations: careful setting of rates (i.e. the price of electricity, P_{Eit}) according to a standard formula, an obligation to serve the amount of electric power demanded (E_{it}) given that regulated price, and direct oversight of capital investment. PUC oversight of capital investment requires that the utility meets the RPS requirements (internalizing some externalities), that the utility treats E_{it+1} as exogenous to planning decisions (to keep the monopoly from exercising its market power); and that the utility invests in technologies with the lowest marginal cost (so that it doesn't game the PUC's rate making formula).[3]

In the short-run (i.e. 1 year), the utility dispatches K_{Rit} and K_{Nit} to generate amounts of R_{it} and N_{it} that minimize operating costs (f_t for fueled operating costs and o_t for other) subject to the technology and obligation to serve constraints, which nicely conforms to a simple linear programming problem:

$$\min_{R_{it},N_{it}} o_t R_{it} + f_t N_{it} \quad s.t. \quad \begin{aligned} R_{it} + N_{it} &\geq E_{it} \\ K_{Rit} \geq R_{it} &\geq 0 \\ K_{Nit} \geq N_{it} &\geq 0 \end{aligned}$$

Given $0 \leq o_t < f_t$, the solution to this operations problem is to first utilize all of the renewable capacity because it can generate electricity at a lower operating cost than production of electric power from non-renewable sources that must purchase a fuel (coal, gas, oil, or uranium) to be consumed; the remaining balance of E_{it} is then generated using the non-renewable production technology.[4]

[3] These simplifying assumptions are not far from reality. PUCs require every utility to regularly submit an Integrated Resource Plan (IRP) to ensure that RPS requirements are met, that the utility is not padding its assets with investment in gold-plated generation assets, and that demand is treated as an exogenous load forecast.

[4] In theory, we could include availability shocks to the renewable and non-renewable capacity; both can become unavailable when machinery breaks down, non-renewables can become unavailable if there is a fuel disruption, and renewables are notorious for becoming unavailable (e.g. the wind is not blowing enough). However, although this inclusion might provide more realistic detail about the relative merits of the renewable versus non-renewable technology, it does not affect our results here.

Under the influence of PUC regulation, the planning decision will result in adding enough capacity to ensure that the utility meets its obligation to serve E_{it+1}. In the real world, utilities that are short of their obligation to serve simply purchase from the grid from utilities that are long – this happens all of the time as operations decisions are made every few minutes; however, in this model there is no uncertainty so there is no risk of a utility being short (or long). Hence, this model needs no grid (demand is entirely supplied by the state's monopolist, i.e. $R_{it} + N_{it} = E_{it}$) and utilities will not have any excess capacity (i.e. $N_{it} = K_{Nit}$). The only remaining role for planning is to decide how much of that additional capacity (needed to meet future demand) should be invested in K_{Rit} versus K_{Nit}. Under the influence of PUC regulation, the utility will decide on its investment (I_{Nit} and I_{Rit}) so as to minimize the marginal cost of investment (unless it needs to meet an RPS requirement). The marginal cost of investment is implicitly contained within the following laws of motion on the capital stocks:

$$K_{Nit+1} = (1-\delta)K_{Nit} + g_N(I_{Nit}; \psi_N)$$
$$K_{Rit+1} = (1-\delta)K_{Rit} + g_R(I_{Rit}; \psi_{R0}, \psi_{R1}, K_{RMaxi}, C_t)$$

Where δ is depreciation, K_{RMaxi} is the maximum amount of renewable capacity that it technically feasible for state i (i.e. it's the renewable potential for state i), C_t is a measure that captures the technologic al improvements that have brought down the cost of renewable generation capital over the past decade, the g()'s are capital adjustment cost functions (akin to those that are fairly common in macroeconomics for slowing down capital accumulation), and the investment cost parameters are the ψs.

Rather than directly specify these capital adjustment costs, it is more convenient to specify the amount of investment required to grow the capital stock at time t+1 to a particular level from a given level at time t:

$$I_N(K_{Nit+1}, (1-\delta)K_{Nit}) = \psi_N[K_{Nit+1} - (1-\delta)K_{Nit}]$$
$$I_R(K_{Rit+1}, (1-\delta)K_{Rit}) = \psi_{R0}C_t[K_{Rit+1} - (1-\delta)K_{Rit}] + \psi_{R0}\ln\left(\frac{K_{RMaxi} - (1-\delta)K_{Rit}}{K_{RMaxi} - K_{Rit+1}}\right)$$

The ψ_N and ψ_{R0} describe the constant component of the marginal cost of increasing capacity. The remainder of the renewable investment cost function is nonlinear to capture the fact that the marginal cost of investment should increase as one approaches a state's maximum potential. Hence, ψ_{R0} captures the generic cost of the equipment invariant of site (e.g. the solar panels) and ψ_{R1} captures how that cost increases as utilities resort to less efficient sites (e.g. more cloud cover means less solar generation per panel). Note that, even without an RPS, utilities would have some stock of renewable generation that grows over time (e.g. conventional hydropower plus some recent modest increases in select low-cost opportunities to deploy renewables); this stylized fact is captured by some units of renewable capacity being cheaper than non-renewable units (until increasing marginal costs of renewables overtakes the constant marginal cost of non-renewables) and that more units will become cheaper as the marginal cost of renewable (equipment) falls over time.

The rates, P_{Eit}, for electric power utilities are set according to a standard formula; the rates are set so that the utility earns enough revenue to cover its operating costs (O_{it}) and get a fair rate of return on its capital investments:

$$P_{Eit}E_{it} = O_{it} + r \times V_{it}$$

Where the left-hand side is the revenue requirement, V is the value of the utility's capital stocks, and r is a "fair" rate of return – the utility earns a revenue of r×V (after covering operating costs) and then could conceivably sell the generation capital assets for V to net a return of 1+r on its investment of V.[5] The evaluation of the investment cost functions at the current stock with no initial capital yields the value of the utility's assets as measured by their replacement costs:

$$V_{it} = I_R(K_{Rit}, 0) + I_N(K_{Nit}, 0)$$

Simple algebra (including dividing the rate setting equation by E_{it} and substituting in θ_{it} for the share of electric power generated from renewables in state i during year t) yields:

$$P_{Eit} = o_t\theta_{it} + f_t(1-\theta_{it}) + r \times \left[\psi_N(1-\theta_{it}) + \psi_{R0}C_t\theta_{it} + E_{it}\psi_{R1}\left[-\ln\left(1-\frac{\theta_{it}E_{it}}{K_{RMaxi}}\right)\right]\left[\frac{1}{E_{it}}\right]\right]$$

The change in the price from the state legislature imposing an RPS at time t to grow θ_{it} by a rate of ρ can be found by substituting in $(1+\rho)\theta_{it}$ into the previous equation, and differentiating with respect to ρ:

$$\frac{\partial P_{Eit}}{\partial \rho} = [o_t - f_t]\theta_{it} + r \times \left[-\psi_N + \psi_{R0}C_t\theta_{it} + \ \psi_{R1}C_t\left(\frac{\theta_{it}}{K_{RMaxi} - \theta_{it}(1+\rho)E_{it}}\right)\right]$$

c. Other Industries Comprising the State's Economy

In order to keep the model tractable, each of the J industries are assumed to produce their output from energy and a composite of all other goods (Z) inputting into a Cobb-Douglas technology with input-output elasticities of γ and ξ (respectively). From the well-known solution to this profit maximization problem, we can substitute the factor demands into the production function and subsequently into revenue to describe revenue as a function of the price of electricity, other prices (P_j and P_Z, which are held fixed in this partial equilibrium analysis) and technology parameters (including A_{jt}):

$$GDP_{jit} = A_{jt}^{\frac{1}{1-\gamma_j-\xi_j}}\gamma_j^{\frac{\gamma_j}{1-\gamma_j-\xi_j}}\xi_j^{\frac{\xi_j}{1-\gamma_j-\xi_j}}P_{Eit}^{\frac{-\gamma_j}{1-\gamma_j-\xi_j}}P_Z^{\frac{-\xi_j}{1-\gamma_j-\xi_j}}P_j^{\frac{1}{1-\gamma_j-\xi_j}}$$

[5] PUCs set electricity rates so that the "fair" rate of return at 8% – observed values exhibit some variation around that 8% due to the realization of (unmitigated) shocks born by the utility.

III. Data

Table 1 presents a tabular overview of the data that we have assembled for this project. In the remainder of this section, we step through that table to describe each major source.

a. RPS Data from DSIRE

An RPS results from an act of legislation. Due to the very nature of the legislative process, the details of each RPS can be quite different from those in other states. For instance, some states may not count energy generated from hydrogen (an energy carrier) unless that hydrogen was itself produced from renewable energy (instead of natural gas and steam, currently the cheapest way to produce hydrogen) while other states have specific carve-outs for a minimum fraction that must come from hydrogen; some states do not even count the existing stock of conventional hydropower but do count new (small-scale) hydropower. Fortunately, quantitative analysis has been facilitated by an effort to create a single database, DSIRE (Database of State Incentives for Renewable Energy), which is funded by the Department of Energy (DoE) and maintained by a partnership of the North Carolina Solar Center and the Interstate Renewable Energy Council (IREC).

For each mandatory RPS, DSIRE contains: the starting date, the share of renewables (that count towards the RPS) at the starting date, the targeted share of renewable, and the date by which that target must be achieved. In order to make these directly comparable with a single parameter, we compute the necessary average annual growth rate of the share of renewables (ρ and θ, respectively, from the model), graphed as the line connecting the DSIRE data points in Figure 1:

$$\theta_{it+\Delta} = \theta_{it}(1+\rho_i)^{\Delta}$$

Figure 1 summarizes this data for states with a mandatory RPS. For each voluntary RPS, DSIRE contains some of this same data but, crucially, does not include the share of renewables (that count towards the RPS) at the starting date; hence, we were unable to perform the same transformation for each voluntary RPS. While that prevents us from making the same use of the voluntary RPS, this does keep our research design cleaner – it is not clear how much a voluntary RPS is just cheap talk versus something that might actually trigger a mandatory RPS if voluntary targets are not met. In addition to the complication of whether the RPS is voluntary, 3 mandatory states (Iowa, Texas, and Michigan) have an RPS that specifies an absolute amount of generation capacity instead of a share of generation. For Iowa, we decided to discard our observations on its RPS because it is quite old (i.e. from 1991) and specifies a negligible quantity of renewable generation, perhaps suggesting that it is a law without teeth that no one has bothered to sunset; for Texas and Michigan, we converted the absolute capacity requirements into comparable terms of a share of generation using the generation data that we gathered from EIA.

b. Annual Generation Data from EIA

DoE's Energy Information Agency (EIA) provides a wealth of data on many aspects of energy production and consumption. Among the data collected by EIA is the annual generation (in MegaWattHours) from 1990-2010 for the total electric power industry by source (coal, petroleum, natural gas, other [fossil] gases, nuclear, hydroelectric conventional, wind, solar thermal and photovoltaic, wood and wood derived fuels, other biomass, and other). We converted these levels in to TeraWattHours, for convenience, and then also aggregated coal, petroleum, natural gas, other [fossil] gases, and nuclear into non-renewable energy and computed the share of each state's generation in each year from all renewable sources. Figure 2 presents a time series graph for most states, split in an ad hoc fashion into natural regions, with dotting to indicate the lack of an RPS (and dashing to indicate if an RPS is voluntary). By inspection, alone, it is not immediately clear whether an RPS actually increases the deployment of renewable from the previous year (even for the mandatory states). Indeed, Figure 3 suggests that the DSIRE and EIA data, although strongly correlated, have a noisy connection.

c. Political Data from DW-Nominate

The gold standard in measuring political preferences is DW-Nominate, initially created by Poole and Rosenthal (1985) and steadily updated through Carrol et al. (2009). Essentially, DW-Nominate scores every legislator in the US House of Representatives since 1789 according to their voting record using a Random Utility Model (RUM); in theory, legislators who have voted the same way have revealed that their stable preferences are similar. [6] To convert DW-Nominate scores to the unit interval (i.e. the 0 to 1 scale assumed in the theoretical model), we perform the inverse-logit transform mapping 0 to a new center of 0.5 and preserves the variance:

$$S_{itk} = \frac{1}{1 + e^{-s_{itk}/0.28}}$$

This function maps an extreme liberal to a value near 0 (e.g. Dennis Kucinich to 0.05) and an extreme conservative to a value near 1 (e.g. Ron Paul to 0.99).[7] Although this transformation on DW-Nominate provides a fine measure of the preferences of members of Congress from a state, our theoretical model is focused on the state legislature. To make the connection between these two bodies of representatives, we assume that the distribution of preferences elected to the national delegation is representative of the state legislature. Hence, the distribution of scores across Ki members of Congress is representative of the distribution of scores within state i's own legislature. Under the median-voter

[6] DW-Nominate actually contains 2 dimensions of scores. The first dimension, which we use here, nicely lines up with liberal versus conservative in the sense of preference over government interference in private markets. The second dimension, which we do not use here, appears to mostly line up with attitudes on slavery and civil rights but loses its significance altogether by 1980.

[7] Although Ron Paul maps to 0.99 and Dennis Kucinich only maps to 0.05, this should not be interpreted as a weakness in the data. Because every Congress includes some incumbents from previous Congresses, the DW-Nominate scores are comparable over time. Hence, the most liberal member of the Congress in 2009 may be to the right of the most liberal member of Congress ever. Hence, Ron Paul mapping to 0.99 is a reflection of how much America's conservatives have tacked to the right (relative to the rest of U.S. History).

theorem, motivated by majority rule, the decisive preferences within the state legislature belong to the median representative. Hence, the theoretical model's reliance on med(S_{it}):

$$med(S_{it}) = \underset{k}{\text{median}}(S_{itk})$$

Although such majority-ruled passes as a first-order approximation of state legislatures, it is a stylized fact of recent history that little legislation is passed when the legislature is excessively polarized. We construct a measure of polarization from the variance of scores. Our transformation of DW-Nominate provides an absolute maximum to the variance of scores, which occurs when half of the delegation has a score of 0 and the other half has 1 (yielding a variance of ¼, the maximum variance of a Bernoulli random variable), allowing us to use a measure of polarization (between 0 and 1) equal to the variance of S_{it} divided by ¼ :

$$pol(S_{it}) = \frac{\underset{k}{\text{var}}(S_{itk})}{1/4}$$

Note that, for those few states that are so small that they only send a single representative to Congress (Alaska, Delaware, Montana, North Dakota, South Dakota, Vermont, and Wyoming), pol(S_{it}) will be undefined. For all other states, Figure 3 provides visualization of the role of the median and polarization measures on the distribution of political preferences within a state in a given year.

d. State Industry Data from BEA

The Bureau of Economic Analysis (BEA) collects data on the GDP of every industry, at the level of 62 different 3-digit NAICS codes, for every year from 1997-2010.[8] Figure 6 depicts this annual GDP data for a subset of states without an RPS and a subset of industries. Figure 7 compares the non-RPS state average for each industry, i.e. mean within industry and year conditional upon being a non-RPS state, to corresponding time and industry for select RPS states. As can be gleaned from these figures, each state-industry appears to have a different intercept even though slope appears to be common across states to a particular industry.

e. Electricity Price Data from EIA

EIA collects average retail prices for each state from 1997-2010, measured in cents per hour. We convert this measure into millions of dollars per TeraWattHour.[9] In theory, an RPS should increase this price and dampen the GDP of industries most sensitive to energy prices. Figure 8 displays the comovement of prices for non-RPS states and select RPS states, as well as the corresponding figure for GDP and a

[8] A handful of values are already in the data for 2011 but most are N/A (Not Available).
[9] As with all other nominal variables in our data set, we are unconcerned about converting into real dollars of a particular year because our specifications effectively absorb these transformations.

vertical line indicating the date that the RPS began. However, because of classic endogeneity problems, we need to instrument for price in such analyses.

f. Renewables Cost Data from Lawrence Berkeley National Labs

The classic instruments for price are exogenous cost shifters. In this context, the cost of the hardware serves as a reasonable instrument. In theory, this cost is set by the state of the industry and the usually dwarfed size of the domestic market relative to the world market. Lawrence Berkeley National Labs (LBNL), funded mostly by DoE, has maintained a data set of the average cost of a watt of solar photovoltaic capacity from 1998 to 2010. Clearly decreasing but somewhat noisy due to the nature of LBNL's sampling, we smooth out this series by replacing non-monotonic points (e.g. 2000) and missing points (1997) with the linear trend as depicted in Figure 9.

g. Renewable Potential Capacity from National Renewable Energy Labs (NREL)

In addition to LBNL, DoE supports dozens of other national labs; perhaps the most important national lab to renewable energy is NREL. NREL has produced an estimate of the total potential, in technical terms (i.e. ignoring economic feasibility), for each major renewable resource in each state. We have simply summed these across the set of renewable resources considered by NREL. Figure 10 presents the distribution of potential (in log scale) across states. Just as we can use LBNL's measure of the cost of solar panel hardware as an instrument for price, we can also use this NREL potential measure as an exogenous cost shifter. Because renewable energy generation requires considerable amount of land, it is congestible and moreover land is heterogeneous; hence, we would expect that the sites first selected for renewable would be

h. Factor Cost Shares from BEA

In addition to GDP by state and industry, as specified by 62 3-digit NAICS codes, the BEA collects data on the nation-wide flow of resources from 1998-2010 across industries in the form of input-output (I/O) tables. We utilize the BEA's I/O data in the form of their "Use Table", which describes how much industry j used of the output of industry –j (dollar denominated). We use this data to ground-truth our second stage estimates of the elasticity of industry output to the price of electric power. In theory, if production does follow a Cobb-Douglas structure, then the utility's fraction of the overall expenses of industry j should be equal to the absolute value of that elasticity; therefore, we would at least expect a monotonic relationship between these factor shares and the sensitivity of GDP to P_E. Hence, we present this I/O factor-share data in the form of tornado diagrams in Figures 11 and 12.

i. State Population from US Census Bureau

The US Census Bureau's population data is widely revered as a gold standard in data quality. Even though the decennial Census is only collected twice (1999 and 2009) during our period of study, the Census provides some highly accurate estimates of population for each state from 1990 to 2010. We have taken those estimates without alteration – they will ultimately divide state GDP across inhabitants to reflect a decision-making process by policymakers where income per capita is central.

j. Operating Costs from EIA

Each regulated utility must submit detailed costs to the Federal Energy Regulatory Commission (FERC), which only makes them available through a database that is so unwieldy that it has given rise to a cottage industry of 3rd parties offering proprietary interfaces. Fortunately, EIA subscribes to one of these services and has queried FERC data to produce national averages for detailed costs (Operations, Maintenance, Fuel, and Total) for each year from 1999-2010 and the major generating technologies (hydropower; gas, solar, and wind; nuclear; fossil steam). To make operational use of this data, we assume that these inputs are all traded on perfectly competitive markets that equalize their factor costs across states (but no over time).

Given these bold assumptions, we separate out solar and wind by assuming that they have the same O&M costs as gas but the same fuel cost as hydropower (i.e. no fuel cost). We then construct our measure for renewable and non-renewable operating cost as the weighted average of these technologies with the weights set to the share of generation from the technology (according to EIA's generation data), all multiplied by an overhead rate to cover indirect costs. We estimate the overhead rate from the ratio of total O&M + fuel costs weighted by generation technology to the total costs taken from another EIA table from the same source (8.1 versus 8.2, both in EIA's Annual Energy Outlook and both produced from the same FERC data), which lacks the breakdown by technology but is more extensive and (presumably) inclusive in its breakdown of costs: Electric Utility Operating Expenses, Operation, Production, Cost of Fuel, Purchased Power, Other, Transmission, Distribution, Customer Accounts, Customer Service, Sales, General Administrative, Maintenance, Depreciation, Taxes, and Other Utility. Considering the common data source and variable labels, we have assumed that the latter table covers the indirect/overhead costs omitted from the other table. Hence, we can estimate the overhead rate by taking a ratio of these costs. The resulting overhead rate averages around 75% (i.e. the technology-specific operating costs should be scaled up by a factor of 1.75). Multiplying that overhead by the weighted average of operating costs yields an operating cost of around 1.2 ¢/kWh for renewables and around 5.7 ¢/kWh for non-renewables. As with the initial missing year for the LBNL cost data, the missing operating costs for 1997 and 1998 are filled with imputed values fitting the linear trend (decreasing for non-renewables and flat for renewables).

IV. Econometric Methodology

The estimation proceeds in 3 stages: a non-parametric prediction of energy prices from costs, a reduced-form panel regression of state-industry GDP on these predicted energy prices, and a Bayesian estimation of a structural model. In this section, each of these stages is discussed in turn.

a. Non-parametric Prediction of Energy Prices from Costs

In order to capture the loss in per capita income due to an RPS, we need to characterize how the price of electric power changes with an RPS and how GDP changes with those prices. Although regressing GDP on price is the most direct approach, it is plagued by the fact that energy prices may have some covariance with the disturbance term (which includes prices of other goods that are omitted from the specification, according to our very own theoretical model). Hence, we seek to instrument for the price of electric power using a couple of exogenous cost shifters: the time-invariant technical potential for renewables in each state (K_{RMaxi}) and the state-invariant average cost (C_t) of adding additional renewable capacity (in the form of solar PV panels) in year t. Hence, we specify that the within state variation in log price, controlling for time fixed effects, should be some monotonic function of these exogenous cost shifters:

$$\ln P_{Eit} = \alpha_{0i} + \alpha_{1t} + g(K_{RMaxi}, C_t) \times 1\{RPS_{it}\} + \varepsilon_{1ijt}$$

Because we do not know the actual functional form of g() and because our objective from this step is to obtain a good fit so that our instruments are not weak, we take a non-parametric approach to fitting that functional form. As is well-known, a sufficiently high order polynomial can approximate any continuous functional form arbitrarily well. As is slightly less well-known, high order polynomials suffer from near multicollinearity (even for relatively low level polynomials) when one uses the monomials as a basis. Hence, we employ Chebyshev polynomials due to their orthogonality properties:

$$\ln P_{Eit} = \alpha_{0i} + \alpha_{1t} + \sum_{m=1}^{M}\sum_{k=1}^{M} \alpha_{2mk} \times Cheby(C_t; m) \times Cheby(K_{RMaxi}; m) \times 1\{RPS_{it}\} + \varepsilon_{1ijt}$$

Where the α's are utilitarian parameters, M is the order of the Chebyshev polynomial in each direction, and Cheby() is a term in an order M Chebyshev polynomial with the following functional form:

$$Cheby(x; n) = \cos\left(n \times \cos^{-1}\left(2\left[\frac{x - \min x}{\max x - \min x}\right] - 1\right)\right)$$

We still must choose M, the order of the Chebyshev polynomials. What makes this approach non-parametric is that n would increase with successive increases in sample size and the fit would get arbitrarily close to the true underlying relationship. Although there are more objective means of selecting M, such as cross-validation or plug-in estimators, we select it based on a subjective criterion. We select the highest M so long as the predicted relationship remains monotonic because any departure from monotonicity is a clear indication of over-fitting the data (given that theory dictates a monotonic relationship).

b. Reduced-form Panel Regression of State-Industry GDP on Instrumented Energy Prices

Given our estimates from the first step, it is a relatively straight-forward exercise to recover the γ's using panel methods. We can use a basic model that accommodates our findings from inspecting the raw data (e.g. industry-specific time trends and large differences in intercepts across states and industries):

$$\ln GDP_{ijt} = \alpha_{3ij} + \alpha_{4j}t + \gamma_j \ln \hat{P}_{Eit} + \varepsilon_{2ijt}$$

Given our structure of the right-hand side, which is based on our inspection of the data, this amounts to a SUR of J=62 regressions on the data demeaned by state-industry. Because we are performing these regressions on predicted log price, each of these J regressions is a two-stage least squares (2SLS) estimation (which is equivalent to instrumenting log electric power prices with its exogenous regressors). Because the SUR then performs an addition feasible generalized least squares (FGLS) to take advantage of cross-equation correlations in the disturbance terms, so we will have effectively performed 3 stages in these two steps.[10]

c. Bayesian Estimation of Structural Model

The structural model consists of 2 estimation equations. The first equation comes from utility revenues (given by the product of electricity prices times the quantity generated); it identifies the investment cost parameters from the 700 observations on revenue (net of operating cost), quantities, and the exogenous cost shifters that previously served as instruments in the reduced-form panel estimation:[11]

$$P_{Eit}E_{it} - O_{it} = \psi_N[r(1-\theta_{it})E_{it}] + \psi_{R0}[rC_t\theta_{it}E_{it}] + \psi_{R1}\left[-\ln\left(1-\frac{\theta_{it}E_{it}}{K_{RMaxi}}\right)\left(\frac{1}{E_{it}}\right)\right] + \varepsilon_{5it}$$

The other estimation equation comes from the state's FOC, which contains the previously derived expression for the partial derivative of energy price with respect to ρ, equating marginal benefit to marginal cost. The equating of marginal benefit and cost yields an implicit equation for the optimal ρ (censored unless an RPS is passed in year t) that identifies the benefit parameters (β_1, β_2, β_3, β_4) from the 23 observations on ρ given estimates of the cost parameters:[12]

[10] For this reason, many analysts call this procedure 3SLS (following Zeller).
[11] Because r cannot be separately identified from the ψ parameters, we set it to 0.08, which is an innocuous assumption that does not affect our results but does aid in interpretation.
[12] The 23 observations come from observing 23 states adopting an RPS between 1997 and 2010.

$$-\left[\sum_{j=1}^{J}\gamma_j \frac{GDP_{ijt}}{Pop_{it}} \times \frac{1}{P_{Eit}}\right] \times \frac{\partial P_{Eit}}{\partial \rho} = \beta_1 + \beta_2 t + \beta_3\left[\frac{1}{2} - med(S_{it})\right] + \beta_4 t\left[\frac{1}{2} - med(S_{it})\right] + \varepsilon_{4it}$$

Note that the left-hand side variable (i.e. marginal cost) is constructed from the parameters estimated in the revenue equation but that these benefit parameters do not enter the revenue equation. Although one could exploit this diagonal structure to perform a 2-step regression, we do not actually estimate it that way for two reasons. First, it can be difficult to produce valid estimates of the standard errors of the parameters from the second step. Second, more importantly, joint estimation means that each parameter is partially identified by the respective data for every equation in which it is included. Hence, the 700 observations provide additional information for identifying the benefits parameters, which would otherwise suffer from a severely small sample size.

In addition to the parameters from the revenue equation, the left-hand side variable is also constructed from the elasticity parameters identified in the previous stage of analysis. Hence, our estimates should also account for the fact that these elasticity parameters estimates have some uncertainty. Moreover, these parameters non-linearly interact on the left-hand side. To accommodate this non-linearity and also jointly estimate this system of equations, we require a simulation-based estimator. While other simulation estimators might suffice, we believe that Bayesian MCMC is particularly well-suited to this problem because it is robust to small sample sizes and we have observed only a small sample of mandatory RPS policies (which we describe by the growth rate ρ). In order to minimize the influence of the priors on our parameter estimates, we specify non-informative (i.e. [virtually] flat) priors.

Because we have already performed a reduced-form estimation of γ_j, we appeal to Empirical Bayes methods in providing a joint normal prior on the 62 γ_j that is centered on the point estimates from the reduced-form estimation with variation around that point estimate described by the variance-covariance matrix also taken directly from their reduced-form estimation. The estimates will, then, receive some Bayesian updating from the information contained in the data about the estimation equations; however, most of the information will go to identifying the parameters with flat priors.[13]

To complete specification of the model's likelihood, we must make a distributional assumption for each of the ε's:

$$\varepsilon_{nit} \sim Logistic(\mu_n, \sigma_n^2)\ \forall n \in \{4,5\}$$

Where we make the fairly standard identifying assumption that $\mu_4 = 0$.[14]

[13] The β and ψ parameters are normals with miniscule "precisions" (i.e. the reciprocal of variance) but centered on reduced form estimates in order to speed convergence. The "precisions" for each equation are specified as a gamma(0.01,0.01).
[14] Because the net revenue regression has no intercept, μ_5 is identified.

V. Results

a. Non-parametric Prediction of Energy Prices from Costs

Goodness of fit is very important in this first stage because a weak fit means that these instruments are weak and will provide a weak signal on the quantities of interest relative to the noise, resulting in large standard errors.[15] As can be seen in Table 2, our first stage regression explains most of the variation in log energy prices as measured by an R^2 of 0.95 and the plot of observed versus predicted log energy price in Figure 15. In addition to the goodness of fit, two other features stand out in that figure. First, the black points (prices in states with a mandatory RPS after the RPS has started) tend to be higher than the blue points (prices in states without a mandatory RPS [yet]). Second, a handful of points seem to be outliers that are relatively distant from the core cluster; these points all belong to Hawaii which is absolutely distant from the mainland and hence it is both unconnected to the grid and it is much more costly to ship non-renewable fuel to Hawaii (which is endowed with virtually no fossil fuel deposits).

Although the particular estimates in Table 2 have little interpretable meaning and their individual significance is irrelevant, their joint ability to predict log energy prices is important. Because the R^2 of 0.95 is already achieved with a first-order Chebyshev polynomial, there is little variation left to be explained by adding the $[(2+1)^2 – 1] – [(1+1)^2 – 1]$ additional regressors of the second-order Chebyshev polynomial; indeed, the increase in R^2 from the additional Chebyshev terms is negligible.[16] Moreover, the second-order Chebyshev polynomial's over-fitting of the data produces a highly non-monotonic saddle point, as can be seen in Figure 14 and contrasted with Figure 13. For these reasons, we selected a first-order Chebyshev polynomial (with interaction term) for our first stage. If more data became available, we would likely select a higher order polynomial with the order increasing in the sample size to produce an arbitrarily good fit to the underlying surface from an asymptotic quantity of data – hence this remains a non-parametric procedure even though the procedure has selected a low order model that coincides with an off-the-shelf reduced-form specification.

b. Reduced-form Panel Regression of State-Industry GDP on Instrumented Energy Prices

Our second stage regression of log GDP on the log energy price predicted by the instruments, as well as state-industry specific intercepts and industry-wide time trends, produces 2 full tables of results: Table 2 and Table 3. Of all of the industries, only 4 goods industries appear to be on a significant decline: textiles, printing, leather (& apparel), and motor vehicles (& parts). Most of these are light manufacturing that has been moving overseas for decades. The decline in wood products is not

[15] A strong correlation between the instruments and the "endogenous" regressor is one of two assumptions required for IV regression. Unlike this one, the second assumption – that instruments are orthogonal to the unobserved variation in the disturbance term –is not testable. Instead, it is justified with a theoretical argument. Exogenous cost shifters, as used here, are a commonly employed and usually accepted argument.
[16] We investigated this for Chebyshev polynomials up to order 4, at which point the $[(4+1)^2 – 1]$ regressors had thinned out the data so much that identification became an issue.

statistically significant, neither is the increase in primary metals. All other goods and service industries enjoy an increasing trend nation-wide.

The point estimate for the elasticity of output to energy prices is negative for the majority of industries. Most of those negative elasticities obtain some level of significance. Yet, there are some notable exceptions that are significantly positive. Water transportation, warehousing (& storage), air transportation, electrical equipment/appliances, computers/electronics, and utilities are all significantly positive. Except for utilities, each of these industries can be viewed as a substitute for energy (or an energy-intensive industry); for instance, higher electricity prices might spur consumers to purchase more energy efficient appliances and retire vintage units. Microeconomic theory provides a straightforward reason why utility profits appear to be increasing in energy prices: current market conditions tend to have consumers on the inelastic portion of their demand curve. Because electric power utilities are natural monopolies in the locality of their service area, they would like to charge higher prices to bring in more revenue than what's lost to a lower quantity demanded at that price. Despite PUC regulation targeting an economically efficient outcome from a perfect competition counterfactual, demand for electric power remains price elastic.[17]

Figure 16 plots the elasticity of output to energy prices for the 62 industries against the utility's share of costs born by industries. If the technologies employed by various industries conform to the Cobb-Douglas production function, then the factor shares would equal the elasticity (in magnitude, the sign of factor shares cannot be negative but the elasticity should be negative under the standard partial equilibrium model). In this case, the points would fall along a line that leaves the origin with a slope of negative 1. Contrary to evidence presented by van der Werf (2008), our results provide some support for Cobb-Douglas technologies: the best linear fit to these 62 points gets extremely close to the origin and has a slope that is statistically indistinguishable from -1. Except for right near the origin, where it is hardest for the confidence region to avoid the positive orthant, the confidence region remains negative. Yet, there are half a dozen individual industries with confidence bands that do not drop below 0 (water transport is the point highest up, the fourth highest with narrowest confidence bands is the utility industry). For these industries with most (or virtually all) of their uncertainty distribution's probability mass above 0, these estimates can be taken as evidence of the importance of general equilibrium effects. Indeed, each elasticity measures the partial change in GDP in its industry resulting from the change in an RPS state's price from the exogenous shifters (the cost of renewable capital and renewable potential constraints) while holding all other prices constant, even though those other prices are not actually held constant in the actual economy (the fundamental data generating process).

c. Bayesian Estimation of Structural Model[18]

[17] Because our location on the demand curve is related to the cost of supplying electric power, it is reasonable that relatively higher cost of production make it prohibitive for PUCs to target rates where consumers would be price-elastic.

[18] Our results were generated using R's interface to JAGS (Just Another Gibbs Sampler) over a parallel environment with 1 cluster node per MCMC chain.

As with any simulation based estimator that relies on iterations, it is important that the method has reached convergence (i.e. inside some tolerance) over the course of those iterations. For Bayesian MCMC estimation, the standard practice is to allow the model to execute for a large number of iterations of burn-in (e.g. 10k, as used here) so that the Markov Chain sequence has converged from the starting values to the ergodic distribution.[19] To gauge convergence, the standard practice is to launch 3 (or more) independent chains from difference starting values and plot the timepath for each chain. If the chains have converged then the timepaths will be wandering around the same stationary mean, which is sometimes referred to as a good mixing of the chains in the parameter space. Examining the timepath for our 3 key parameters (ψ_N, ψ_{R0}, ψ_{R1}), displayed in the top row of Figure 17, convergence of our 3 chains (colored as red, blue, and green) is obvious with such high-frequency bouncing around that mean that the timepaths resemble a solid horizontal line (with some fuzzy whiskers) as if made by a magic marker along a ruler.

When chains do not rapidly bounce around the ergodic mean, there is a fair amount of persistence in the Markov chain. This autocorrelation affects how long the MCMC sampling needs to run – we seek [virtually] independent draws from the ergodic posterior distribution of uncertainty and we obtain those via sequential sampling (i.e. sampling every n^{th} draw from the Markov Chain). Hence, we need to make sure that n is big enough that we have virtually independent draws (and that the MCMC sampling has ran long enough for us to obtain enough draws when we're thinning out all but every n^{th} draw) As can be discerned from the autocorrelograms in the second row of Figure 17, our autocorrelation has died out by the 10^{th} lag. To be safe, we thin out the MCMC samples to every 30^{th} draw.

The last major diagnostic examined here is the distribution generated by each chain. The third row of Figure 17 presents kernel densities for each of the 3 key parameters; these densities are virtually indistinguishable. The [joint] distribution of MCMC draws from the uncertainty distribution is the basic result of Bayesian estimation. Again we represent the distributions for these key parameters, in histogram form with thinning to every 30^{th} draw (giving us 1002 draws from the 3 independent chains), in Figure 18 along with the benefits parameters and a couple elasticity parameters that are representative of the entire set. Along with these histograms, we have also plotted the density obtained from our preliminary reduced-form estimates using classical techniques.[20] We should note that the classical density for the uncertainty distributions on the β parameters is a t distribution with 19 degrees of freedom, which is probably a misspecification because the derivation of the t-distribution requires that all of the regressors and regressands be normal (because the disturbance is linear in these variables, this implies that the disturbance will be normal) and we have made no effort to transform our regressors and regressand so that they (arbitrarily, in terms of economic theory) conform to normality.

[19] Adaptive sampling schemes are also calibrated during this burn-in period.

[20] Because the classical uncertainty distribution on ψ_{R1} was so broad that it dwarfed the histogram, we suppressed it in this plot.

Running the two equations together has the effect of shifting the ψ parameters slightly downward, perhaps with some negligible reduction in uncertainty. We might expect little change due to the fact that these parameters are identified by such a large sample size. The mean of the posterior distribution for each ψ is worth consideration. A ψ_N value of around 325 tells us that an additional TWh of non-renewable generation requires an investment of around $325M, which amounts to about 25 GW of additional capacity that is 50% utilized or about $13M/GW. Clearly, this is a gross underestimate relative to the standard figures from EIA, which suggests that identification of this investment function could be greatly aided by harnessing more data. Likewise, a ψ_{R0} value just beneath 1 suggests that the LBNL estimates of renewable hardware are fairly accurate, although perhaps a bit of an overestimate since LBNL focuses on photovoltaic panels and those tend to be one of the more expensive forms of renewable energy Finally, a ψ_{R1} value of around 722k suggests that renewable congestion costs can become quite steep; once a state is at 10% of its potential, this non-linear portion of investment cost reaches as high as $76B. Although this figure sounds quite large, there is little data directly identifying it because most states are so far down on their potential that they are in the virtually flat portion of the hyperbola. Again, more data and more flexible functional forms might improve this estimate.

Likewise, the γ posteriors are slightly jostled from priors (obtained from our classical 3SLS).[21] However, change in the posterior distributions of the benefits parameters is dramatic. The new means appear to be about a standard error (as measured by the preliminary classical reduced-form methods) away from the old mean. The movement on β_1 is so dramatic that its sign has flipped. Most importantly, the credible intervals (the Bayesian analogue of confidence intervals is just the interval spanned by symmetric quantiles of the posterior) are narrow enough that results are now statistically significant. Utilizing Bayesian MCMC, which can still be effective in small samples and seamlessly handle non-linear simulation-based estimation, has therefore enabled us to reach 4 significant findings that we had previously lacked when we simply relied on classical reduced-form methods.

Our findings on benefits are that centrists opposed mandating renewables in 1997 (i.e. $\beta_1 < 0$), that leftists in 1997 highly valued an RPS in opposition to right-wingers (i.e. $\beta_3 > 0$), that people from all political stripes seem to increasingly value renewable mandates over time (i.e. $\beta_2 > 0$), and that conservatives have warmed up to renewable mandates over time while liberals have cooled off a bit. Most importantly, these trends are convergent so that we may actually be approaching a consensus. Further evidence can be gleaned from Figure 21, where we see the state of the political variables in 21 of the 23 states that have adopted a mandatory RPS.[22] It appears that red states are increasingly implementing an RPS mandate despite considerable polarization of political opinion within these red states.

[21] Credible intervals for each γ is presented in the

[22] Two states (Delaware and Montana) are excluded because they have only 1 representative in Congress and hence we cannot measure their polarization; their median score is around 0.85 (just to the left of Arizona).

VI. Conclusion

In this paper, we assembled a wealth of data on industry, income, various aspects of the electric power industry, renewable energy, politics, and RPS mandates for the 50 states from 1997-2010. We used this data to fit a structural model of legislature decisions on whether to implement an RPS, including a submodel of investment and operations for the electric power industry, and a simple characterization of other industries as Cobb-Douglas technology operating in markets fixed in partial equilibrium. We find evidence to support this model in successive steps of analysis: a nonparametric regression of log energy prices on instruments (exogenous cost shifters), a reduced-form panel data regression of log GDP on those predicted log prices with a SUR FGLS weighting, and a Bayesian MCMC estimation of a system of two equations: utilities' revenue requirement as a function of their operating costs plus a fair rate of return on the value of their renewable and non-renewable generation assets and the first-order condition of the median voter in the state who equates the marginal benefits of an RPS (varying with time and their place on the political spectrum) to its marginal costs (a reduction in GDP per capita due to the increases in energy prices resulting from an RPS).

Future research should hone the investment cost functions with more flexible functional forms and, where available, direct data on the cost of investment in generation capacity. Reexamining the treatment of voluntary RPS legislation might provide additional points for identifying benefits parameters. One could even try to add the censoring equation to the structural model to identify the parameter on the bargaining costs resulting from political polarization. Of course, our treatment of political polarization is rather monolithic – variations in particular rules, coalition forming, logrolling, and the like may well result in the passage of an RPS even when the electorate is polarized. Nonetheless, our best understanding of the relationships that we have identified in the data, given the caveats, is rather encouraging.

We find that centrists opposed mandating renewables in 1997 (i.e. $\beta_1 < 0$), that leftists in 1997 highly valued an RPS in opposition to right-wingers (i.e. $\beta_3 > 0$), that people from all political stripes seem to increasingly value renewable mandates over time (i.e. $\beta_2 > 0$), and that conservatives have warmed up to renewable mandates over time while liberals have cooled off a bit. Most importantly, these trends are convergent so that we may actually be approaching a consensus, which may well lead a long overdue overhaul to national energy policy.

References

Carroll, Royce; Jeffrey B. Lewis; James Lo; Keith T. Poole; and Howard Rosenthal. 2009. "Measuring Bias and Uncertainty in DW-NOMINATE Ideal Point Estimates via the Parametric Bootstrap." *Political Analysis* 17:261-27.

Chen, Cliff; Ryan Wiser and Mark Bolinger. 2007. “Weighing the Costs and Benefits of State Renewables Portfolio Standards: A Comparative Analysis of State-Level Policy Impact Projections”. LBNL Working Paper #61580.

Delarue, Erik; Cedric Jonghe, Ronnie Belmans, William D’haesleer. 2011. “Applying portfolio theory to the electricity sector: Energy versus Power”. *Energy Economics*, 33: 12-23.

Gelman, Andrew; John B. Carlin; Hal S. Stern; and Donald B. Rubin. 2004. *Bayesian Data Analysis, Second Edition*. Chapman Hall / CRC.

Lyon, Thomas P. and Haitao Yin. 2009. “Why Do States Adopt Renewable Portfolio Standards?: An Empirical Investigation”. Working Paper.

Poole, Keith T. and Howard Rosenthal. 1995. "A Spatial Model For Legislative Roll Call Analysis". *American Journal of Political Science*, pp. 357-384.

Prasad, Monica and Steven Munch. 2012. “State-level renewable electricity policies and reductions in carbon emissions”. Energy Policy, 45: 237-242.

Sue Wing, Ian. 2006. “The synthesis of bottom-up and top-down approaches to climate policy modeling: Electric power technologies and the cost of limiting US CO2 emissions”. Energy Policy, 34: 3847–3869.

Van der Werf, Edwin. 2008. “Production functions for climate policy modeling: An empirical analysis”. *Energy Economics*, 30: 2964-2979.

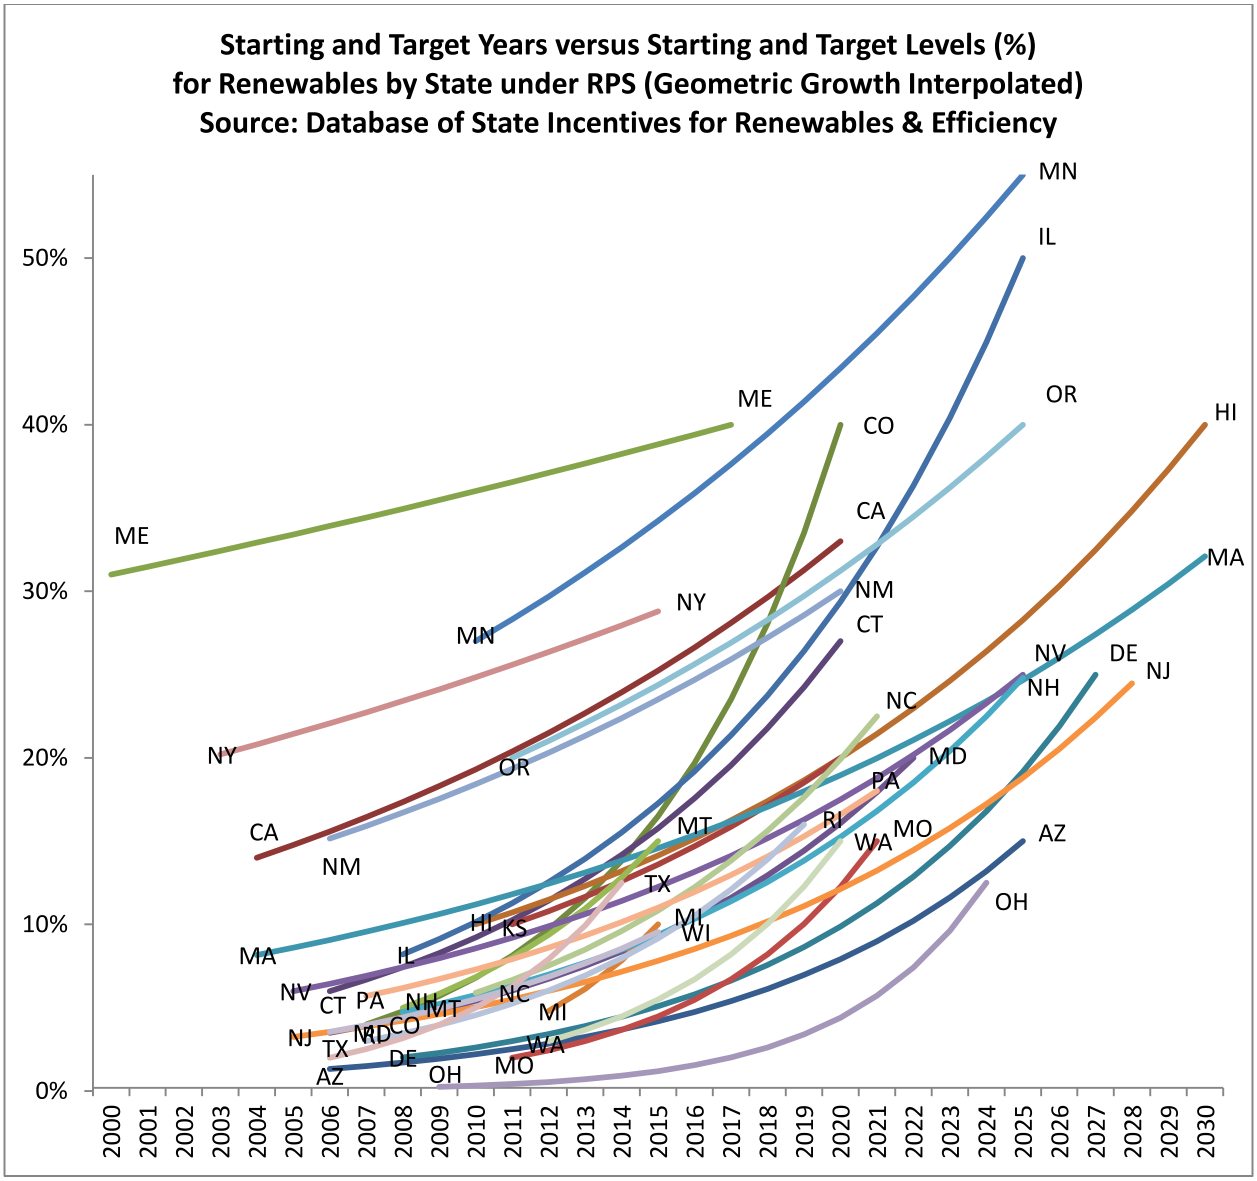

Figure 1. Starting date, starting level, target level, target date, and average growth rate necessary to comply with state's mandatory RPS. It appears that the following states have some of the most ambitious growth rates to reach their RPS targets: Minnesota, Illinois, Maine, Colorado, and Hawaii.

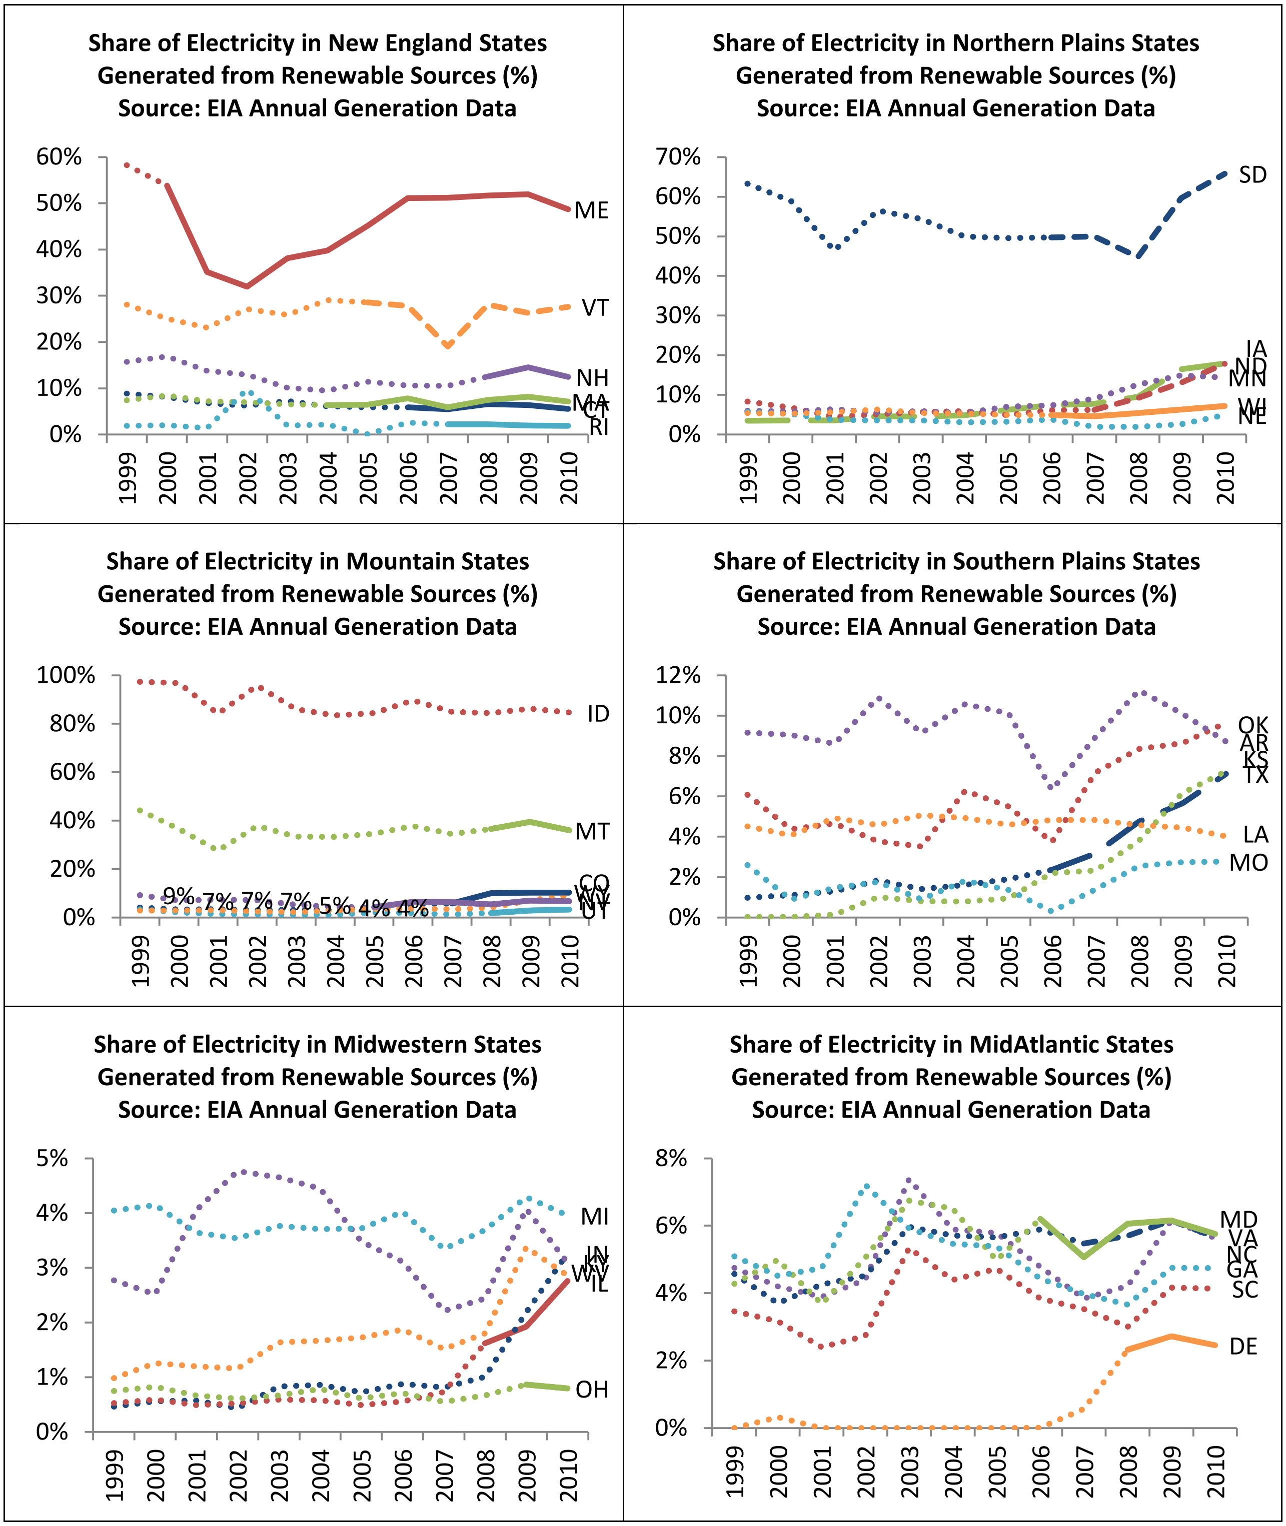

Figure 2. Share of renewables for select states by region. Dotted indicates no RPS, dashed indicates a voluntary RPS, and solidu indicates a mandatory RPS. The breadth of heterogeneity is notable.

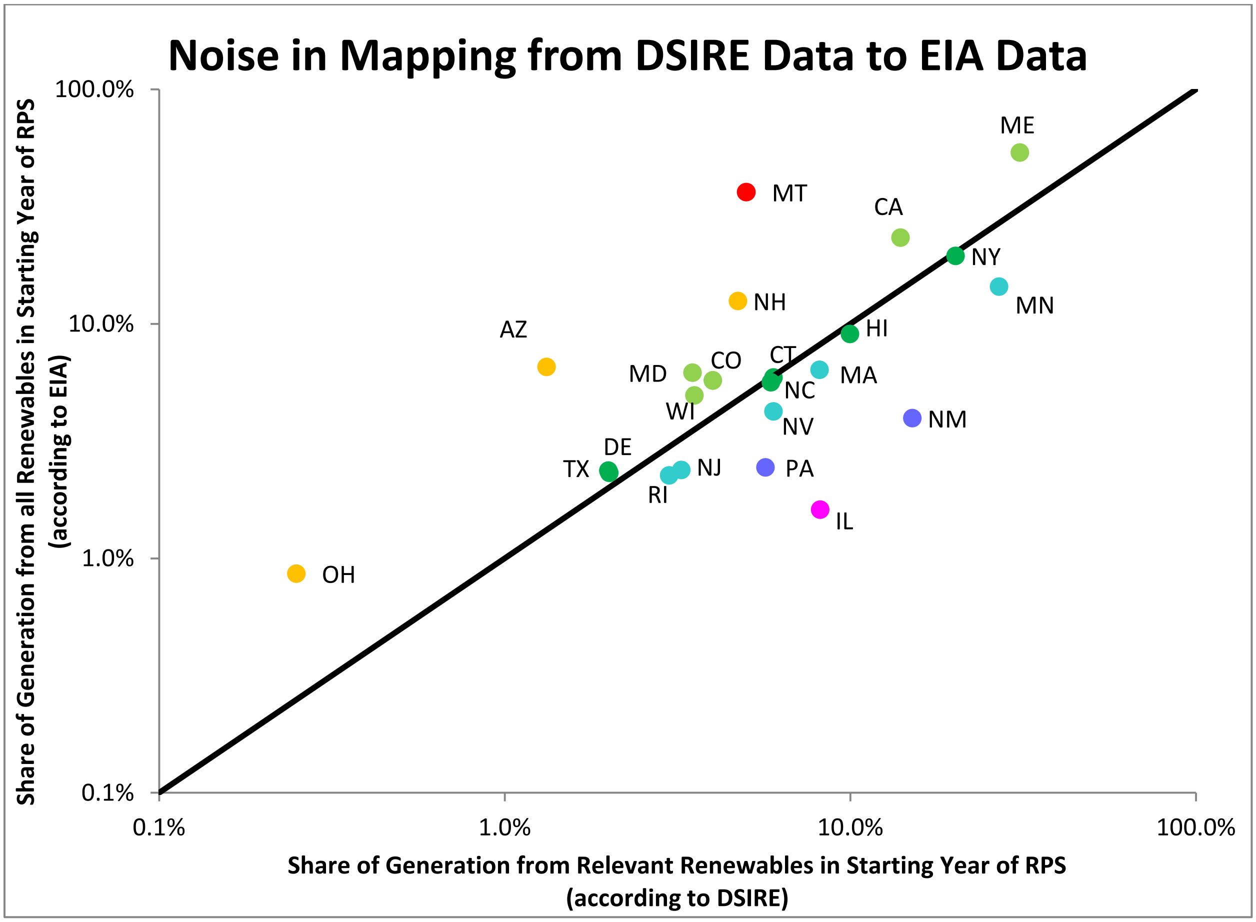

Figure 3. The share of renewables (counted by the state's RPS) in the starting year of an RPS is positively correlated to (but not perfectly so) the more inclusive EIA measure of renewables share.

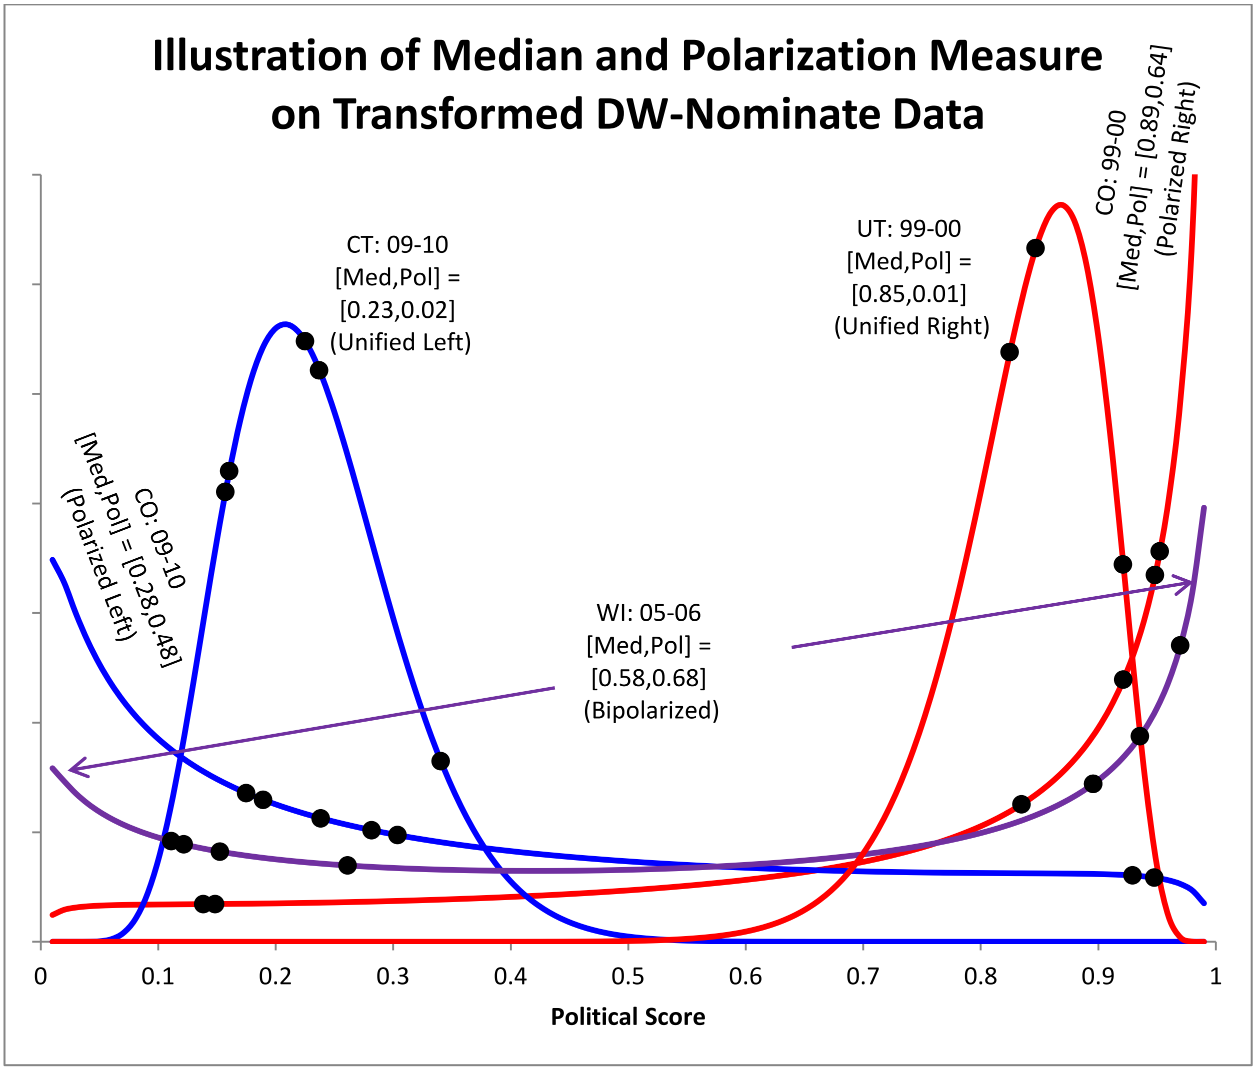

Figure 4. Distribution of political scores across a state in a year for 4 particular states/years with black dots to indicate individual members of Congress, selected so as to illustrate the appropriate of our transformed measure of political score. When the mode(s) of the distribution occur(s) at the extreme(s), then the distribution is polarized. When both extremes are modes, then we dub it "bipolarized". Wisconsin is bipolarized because 3 members are clearly leftists, with scores beneath 0.3, while 3 members are clearly right-wingers, with scores above 0.8 (including Paul Ryan, who is just below the most conservative member of the Wisconsin delegation, Jim Sensenbrenner. In addition to having swung from red in 1999-2000 to blue in 2009-2010, Colorado is less polarized than Wisconsin but more polarized than Connecticut and Utah (both of whom exhibit a mode interior to the extremes).

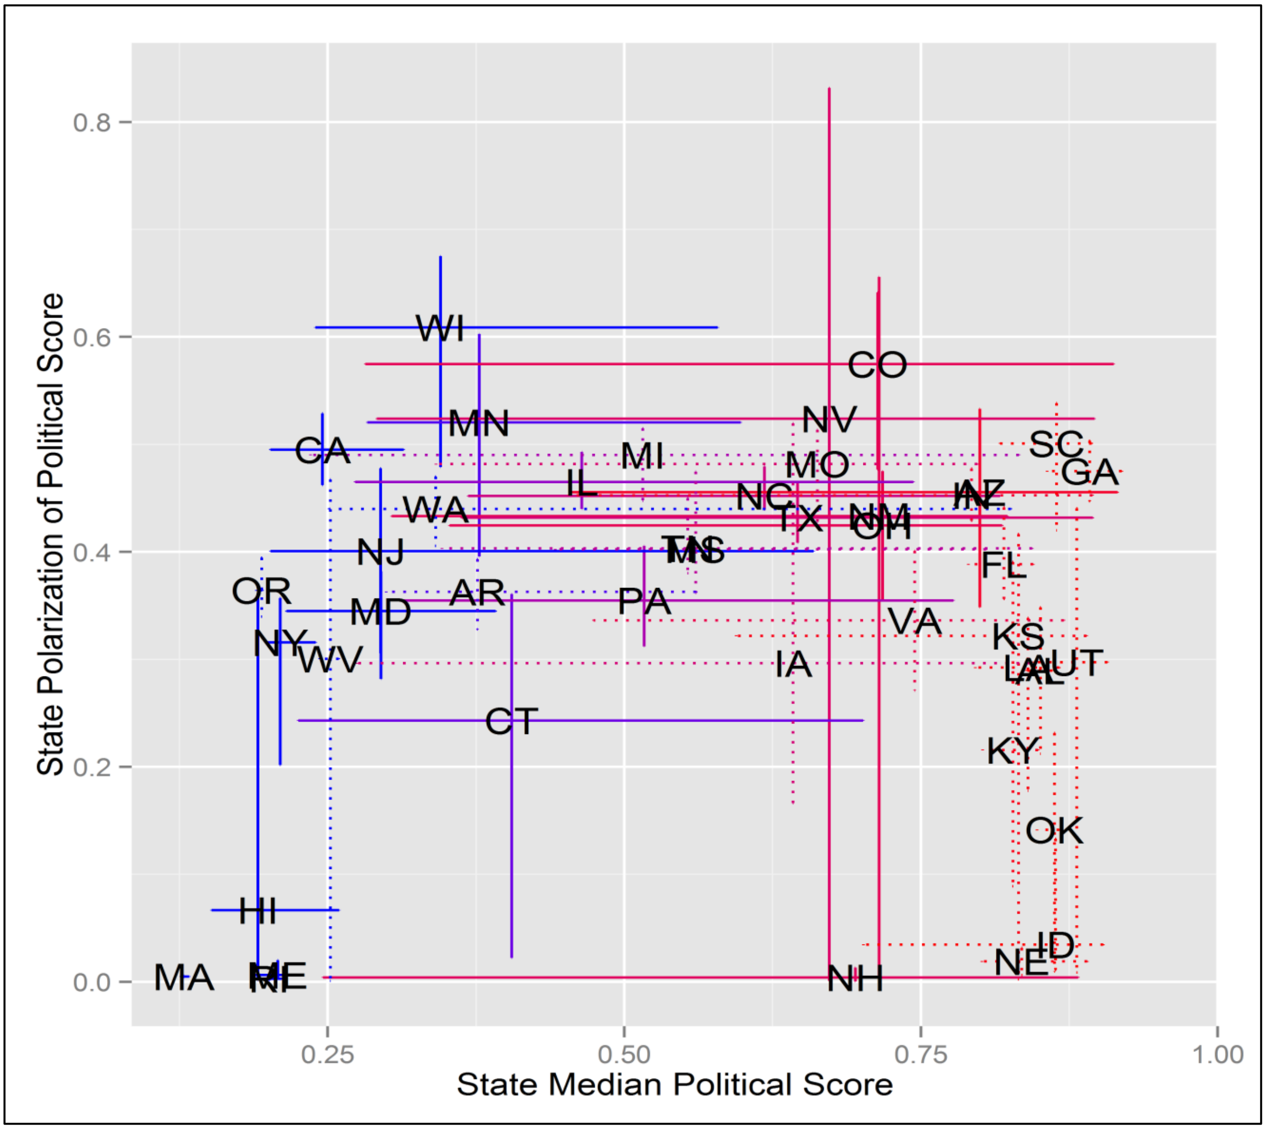

Figure 5. For states with well-defined diversification and a clear observation on their RPS (as mandatory during our study's period or no RPS), we can compare the range of median political score against the range of the polarization of political score. Blue, red, and purple scoring is based on whether the state tends to elect Democrats to Congress (then blue), conservatives to Congress (then red), or some mixture (then purple).

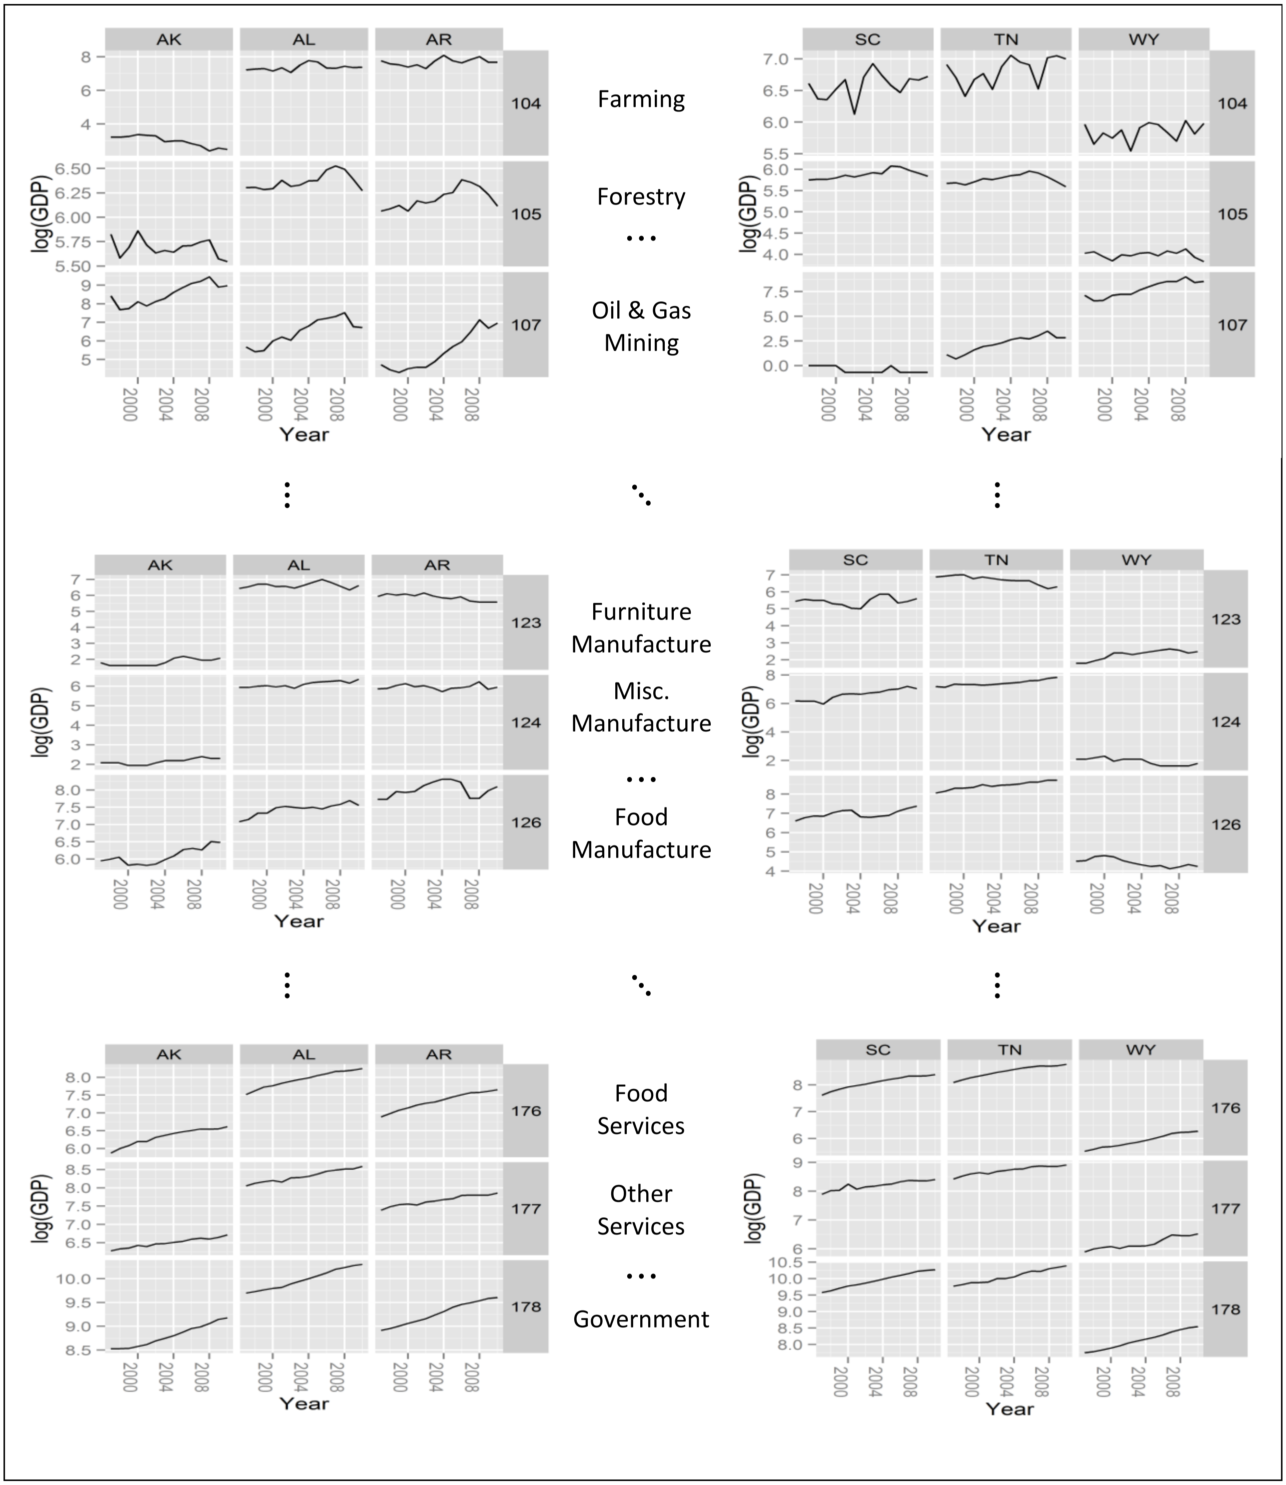

Figure 6. An abridged view of GDP for a 9 industries (selected in order of NAICS code) in 6 non-RPS states (selected alphabetically).

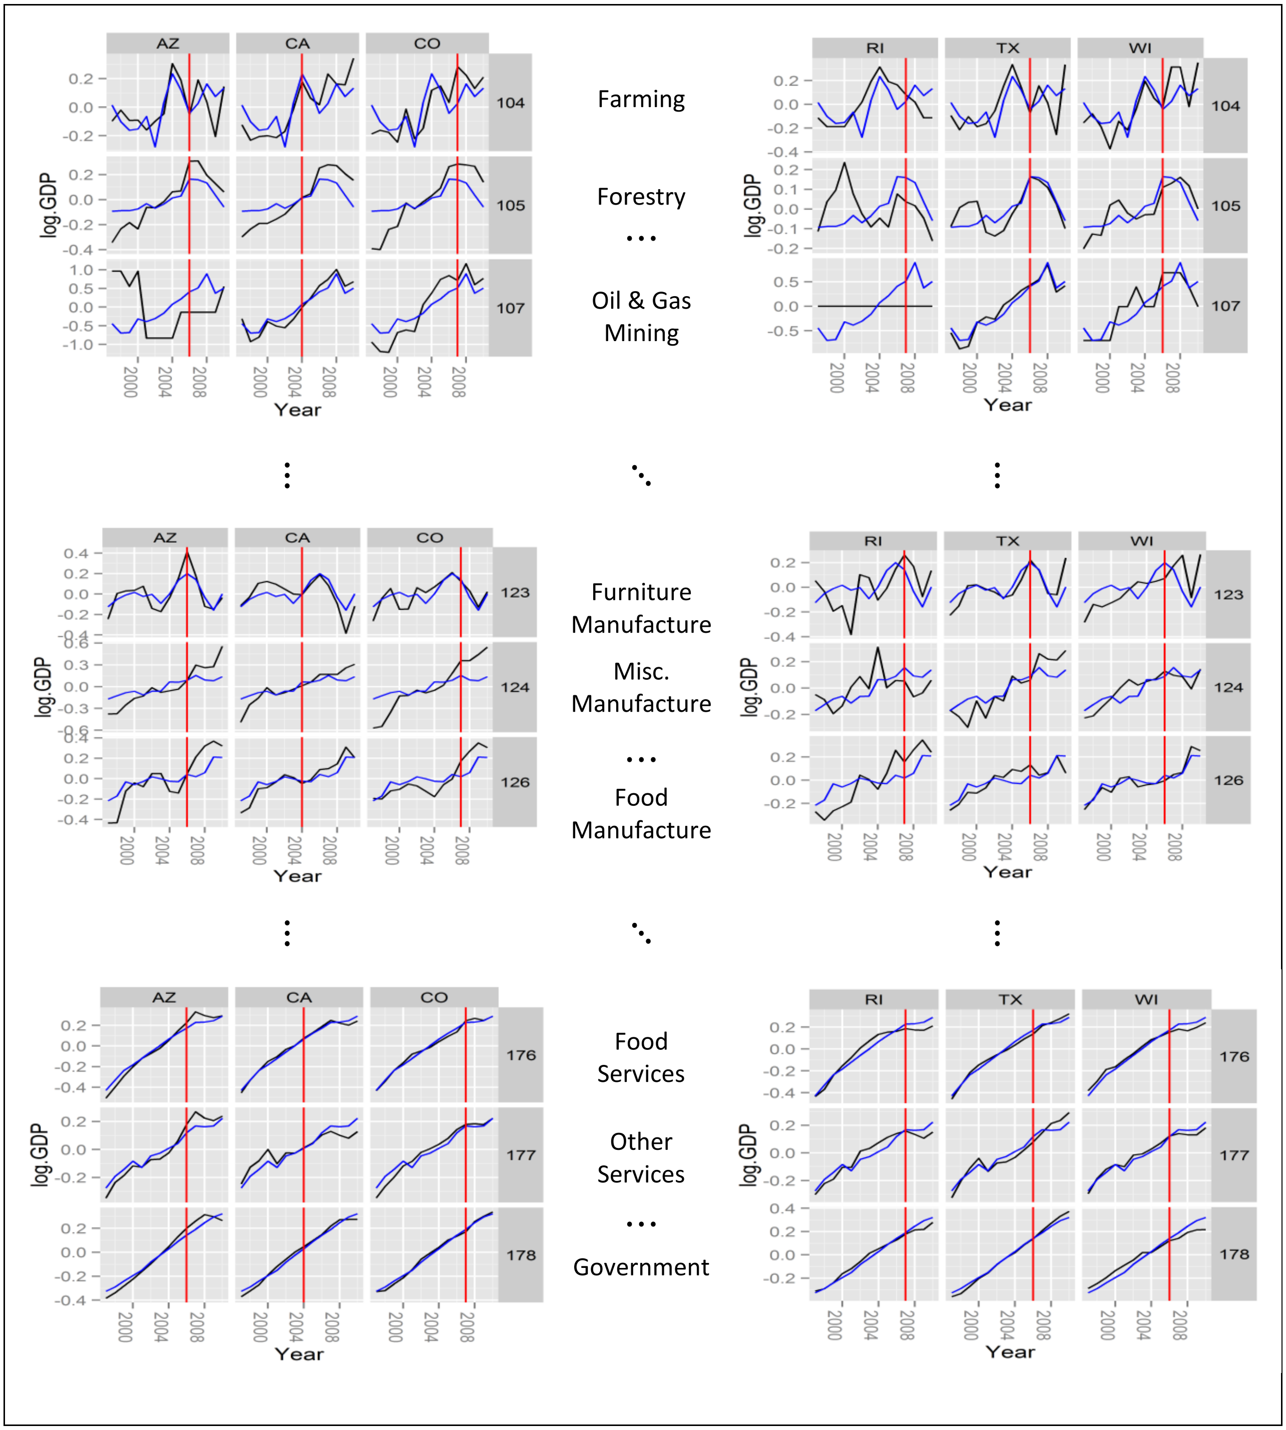

Figure 7. GDP (in log scale) a for select particular RPS States (in black) for select years compared to the average for those same industries in non-RPS states, with a vertical bar indicating the RPS starting date.

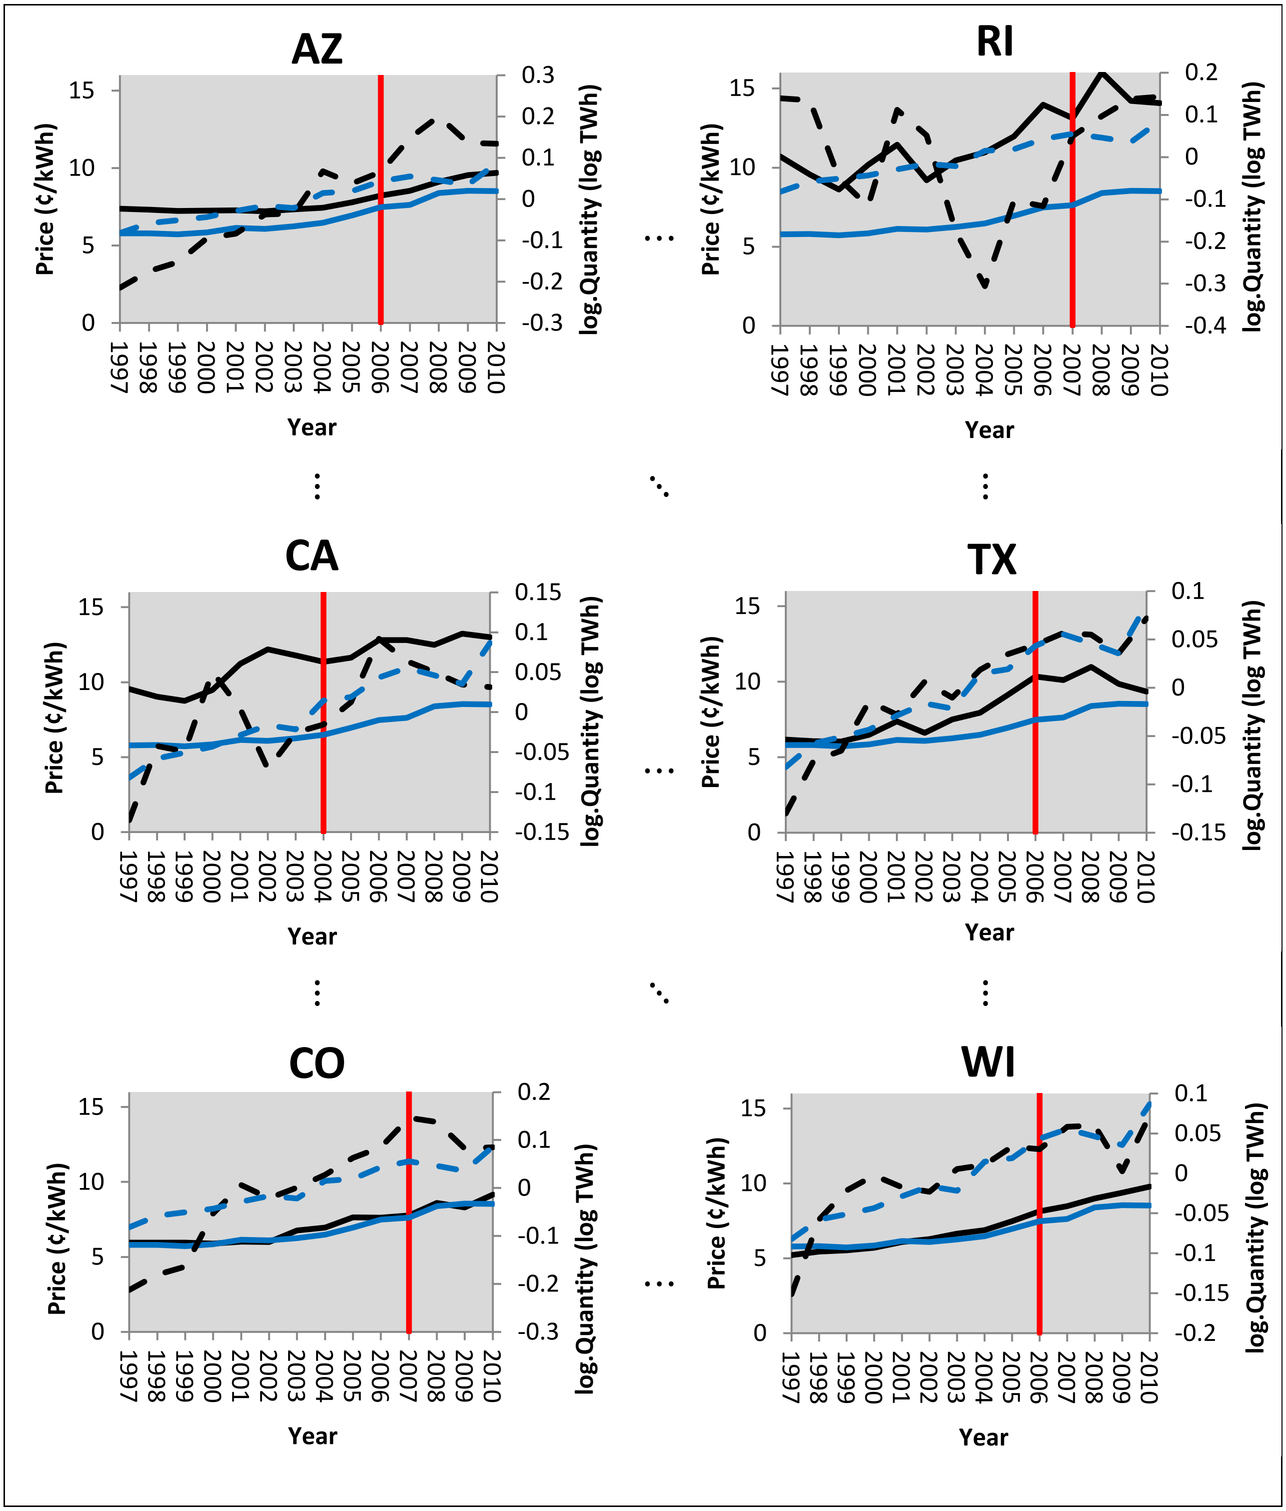

Figure 8. GDP over time for a select few RPS states. The solid lines are the timepath of prices and the dashed indicate output (in log scale).

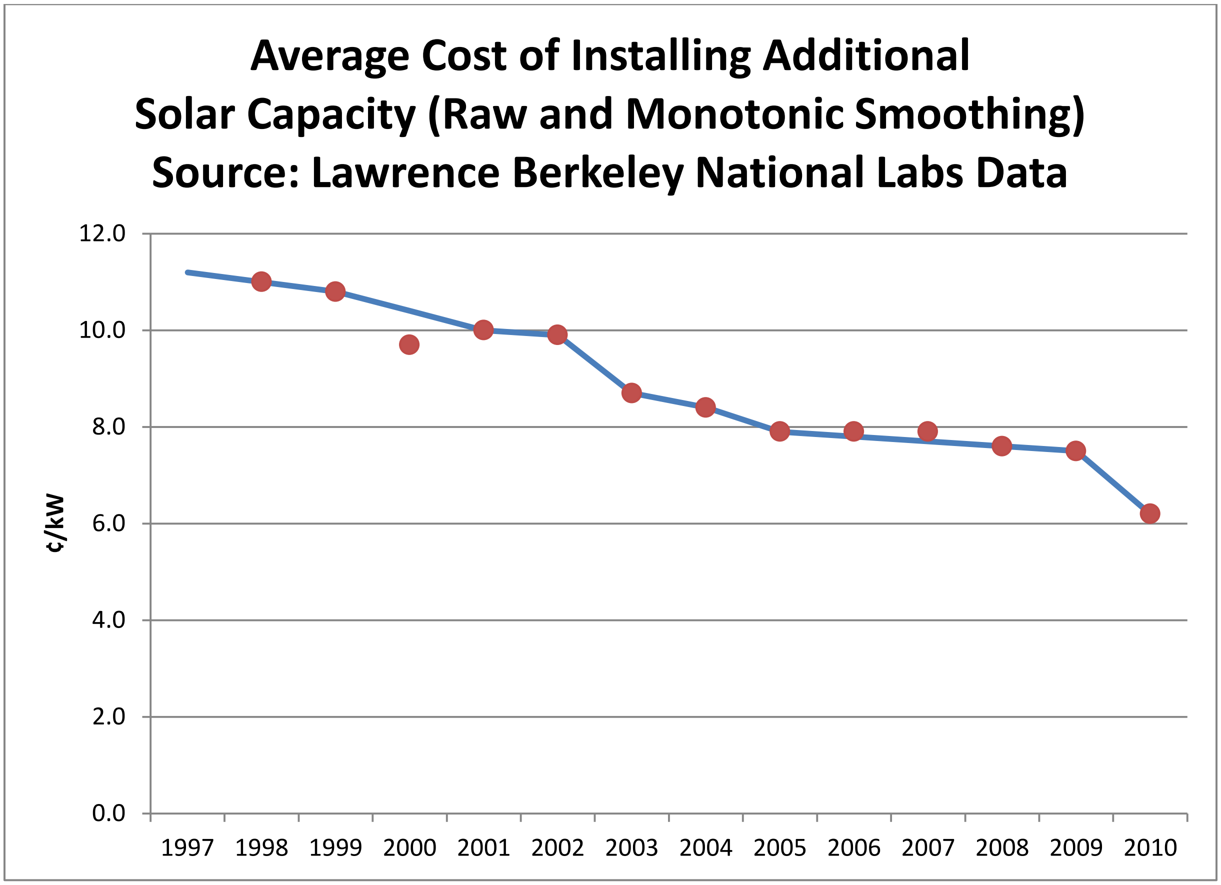

Figure 9. The fall in the average cost of capital that generates renewable energy. Red dots indicate actual data, which was smoothed to a linear trend when non-monotonicities appear or data is missing.

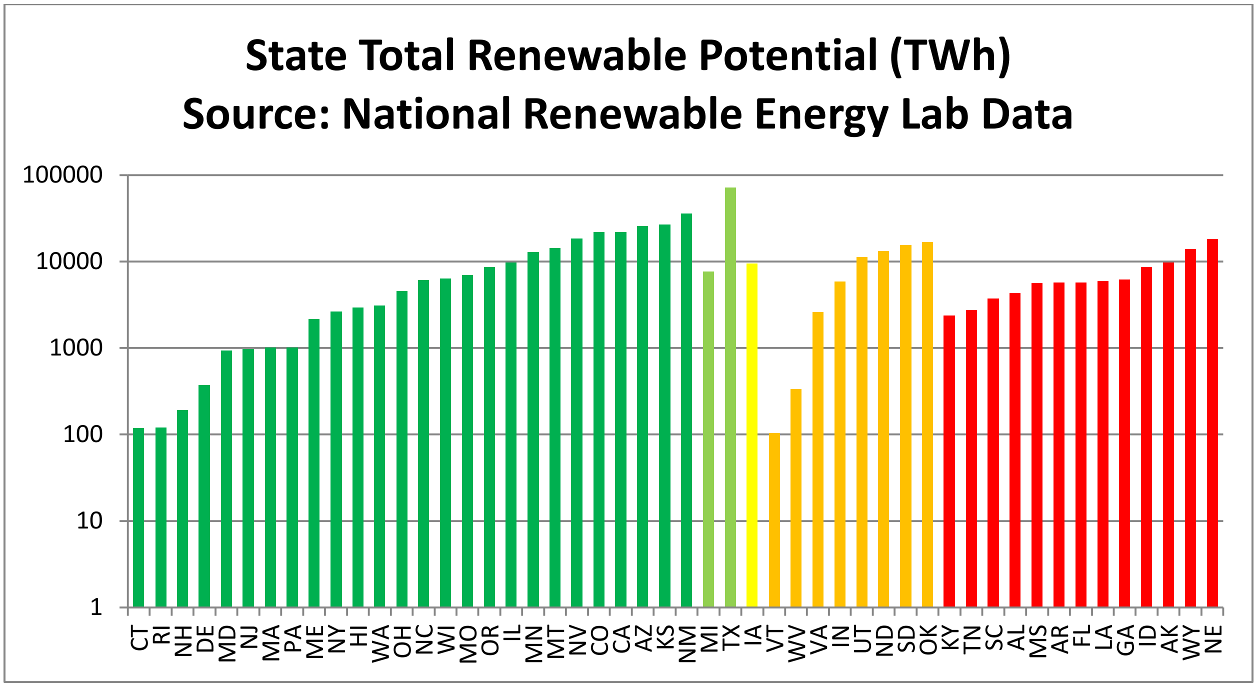

Figure 10. Total Renewable Potential, given current technology (but ignoring whether this capacity is economically efficient, i.e. lowest cost). Green indicates states with a mandatory RPS, the lighter green are the 2 states (Michigan and Texas) where the RPS set a minimum level of capacity, yellow indicates Iowa (mandatory but in name only due its miniscule level and decades old starting date), orange indicates states with a voluntary RPS, and red indicates states that have no RPS.

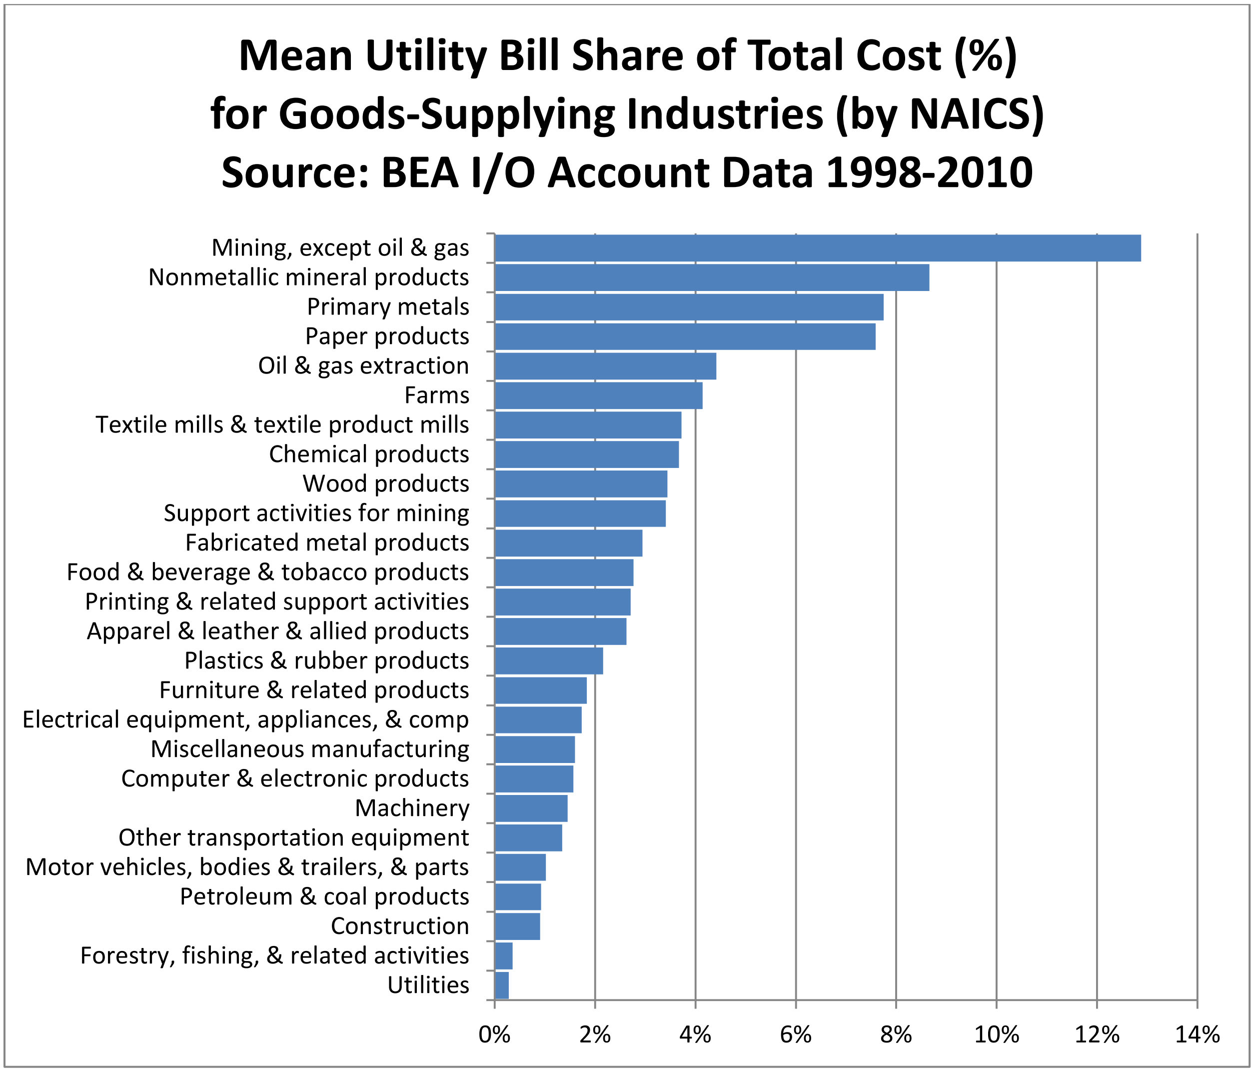

Figure 11. Factor share of utilities to major goods-supplying industries. Utilities, that have the smallest factor shares may well have a large positive elasticity of output (i.e. revenue) to energy prices because that is the price of their output.

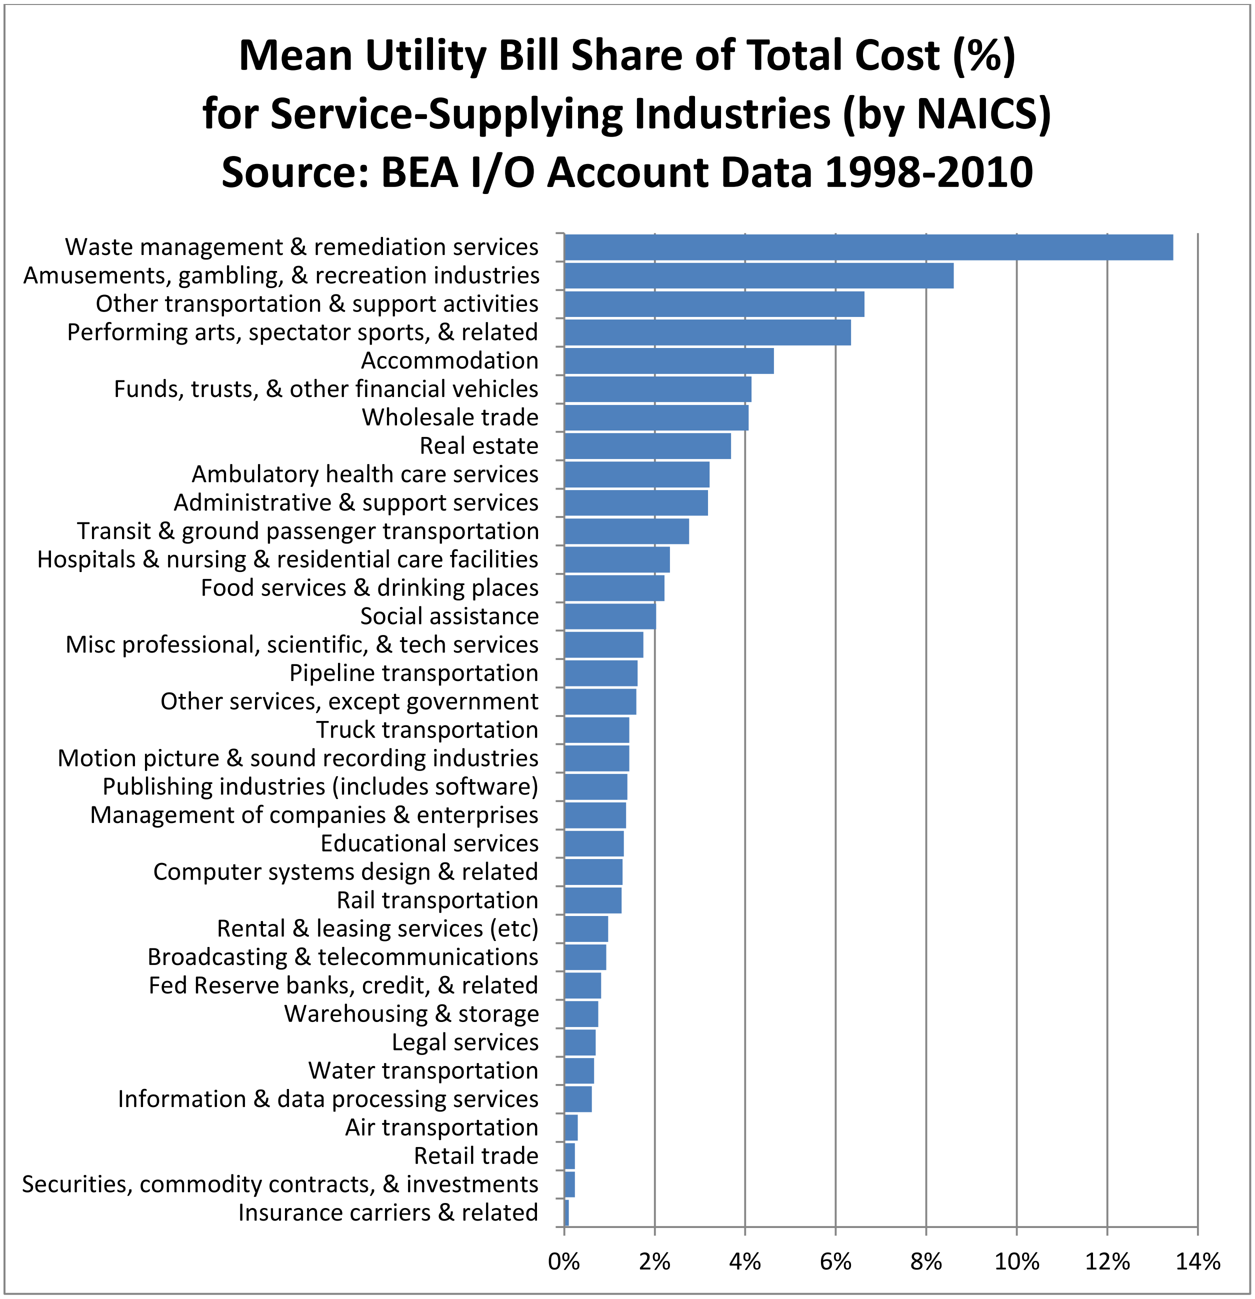

Figure 12. Utilities factor cost shares for major service-supplying industries. Note that government is excluded from this and the previous Figure on principle -- it is usually not treated as just another part of the private economy.

Data Source	Variable(s)	Unit(s) of Measurement	Cross-Sectional Unit(s) of Observation	Years	Number of Observations
DSIRE	Mandatory RPS; RPS Starting Date; RPS Aggressiveness (i.e. Average annual growth rate needed to hit target % of generation from renewable sources by specific year)	Binary; Year; %	50 States	1991-2012	1,100 = 50 ×22
EIA	Generation|Renewable; Total Generation	TWh; TWh	50 States + DC + US Total	1990-2010	1,092 = (5+1+1) ×21
DW-Nominate	Median(Score|State) Min(Score|State) Max(Score|State)	Score (the modeled latent variable from a RUM on voting records)	Members of U.S. House of Representatives for each of the 50 States	1789-2010	36,634
BEA	log(GDP)	$	62 Industries × (50 States + DC + US Total)	1997-2010	45,136 = 62×21× (5+1+1)
EIA	Average State Electricity Price	¢ / kWh	50 States + DC + US Total	1990-2010	1,092 = (5+1+1) ×21
Lawrence Berkeley National Lab	Smoothed (Average Cost of Adding Solar PV Capacity)	$ / W	None	1998-2010	13
National Renewable Energy Lab	Potential Generation from all Renewable Sources	Sum (TWh)	50 States + DC + US Total	None	52
BEA : I/O ("Use") Tables	Cost Share = Utilities/ Sum(All)	% = $ / sum($)	62 Industries × 62 Industries	1998-2010	49,972 = 13×62^2
Census	Population	People	50 States + DC + US Total	1991-2000 & 2001-2010	1,040 = (5+1+1)×20
IEA	Nation-wide Average Operating Costs for Renewable and Non-renewable capacity	M$ / TWh	None	1999-2010	12 12

Table 1. Summary of Data Sources

	Order of Chebyshev Polynomial	
	1	2
Cheby(C_t;1)×Cheby(K_{RMax};0)	- 0.072 ***	0.421
	(0.016)	(0.308)
Cheby(C_t;1)×Cheby(K_{RMax};1)	- 0.072 ***	0.772
	(0.019)	(0.474)
Cheby(C_t;0)×Cheby(K_{RMax};1)	0.015	- 0.092 *
	(0.014)	(0.020)
Cheby(C_t;2)×Cheby(K_{RMax};0)		- 0.093 ***
		(0.038)
Cheby(C_t;2)×Cheby(K_{RMax};1)		- 0.131 *
		(0.057)
Cheby(C_t;2)×Cheby(K_{RMax};2)		- 0.078 **
		(0.025)
Cheby(C_t;1)×Cheby(K_{RMax};2)		0.269
		(0.170)
Cheby(C_t;0)×Cheby(K_{RMax};2)		- 0.060 ***
		(0.014)
Time Fixed Effects	✓	✓
State Fixed Effects	✓	✓
R^2	0.954	0.956
States	50	50
Years	14	14
Observations	700	700
Degrees of Freedom	634	629
Monotonic	✓	✗

Table 2. Results of First Stage regression of log energy price on orthogonal polynomials of instruments.

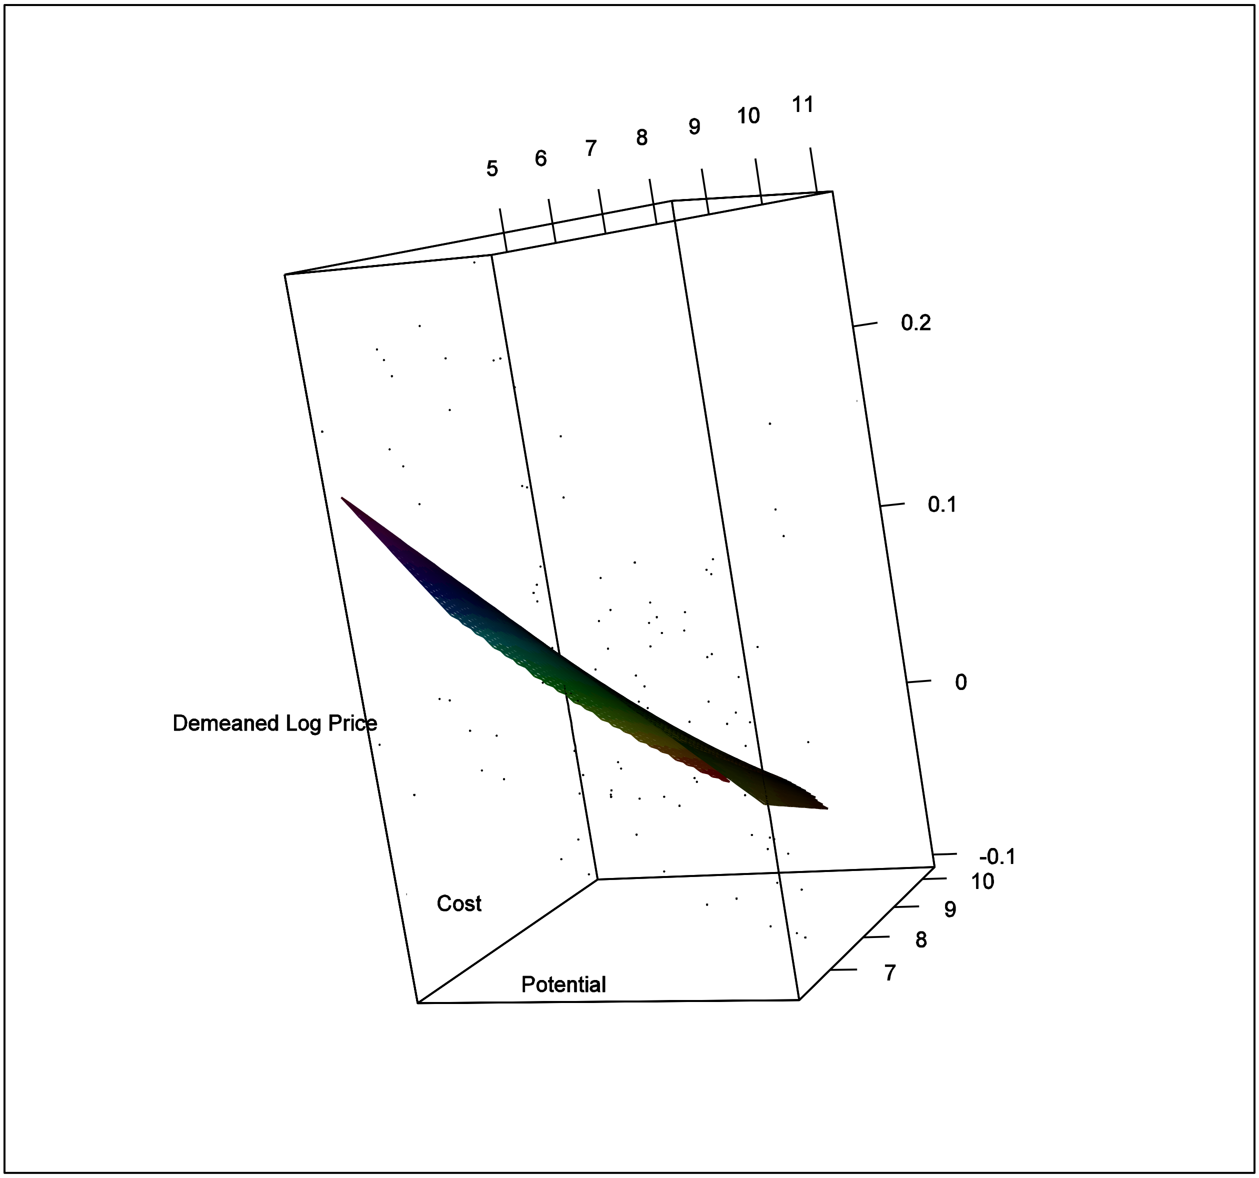

Figure 13. Orthogonal polynomial surface, at selected order, monotonically fitting over instruments.

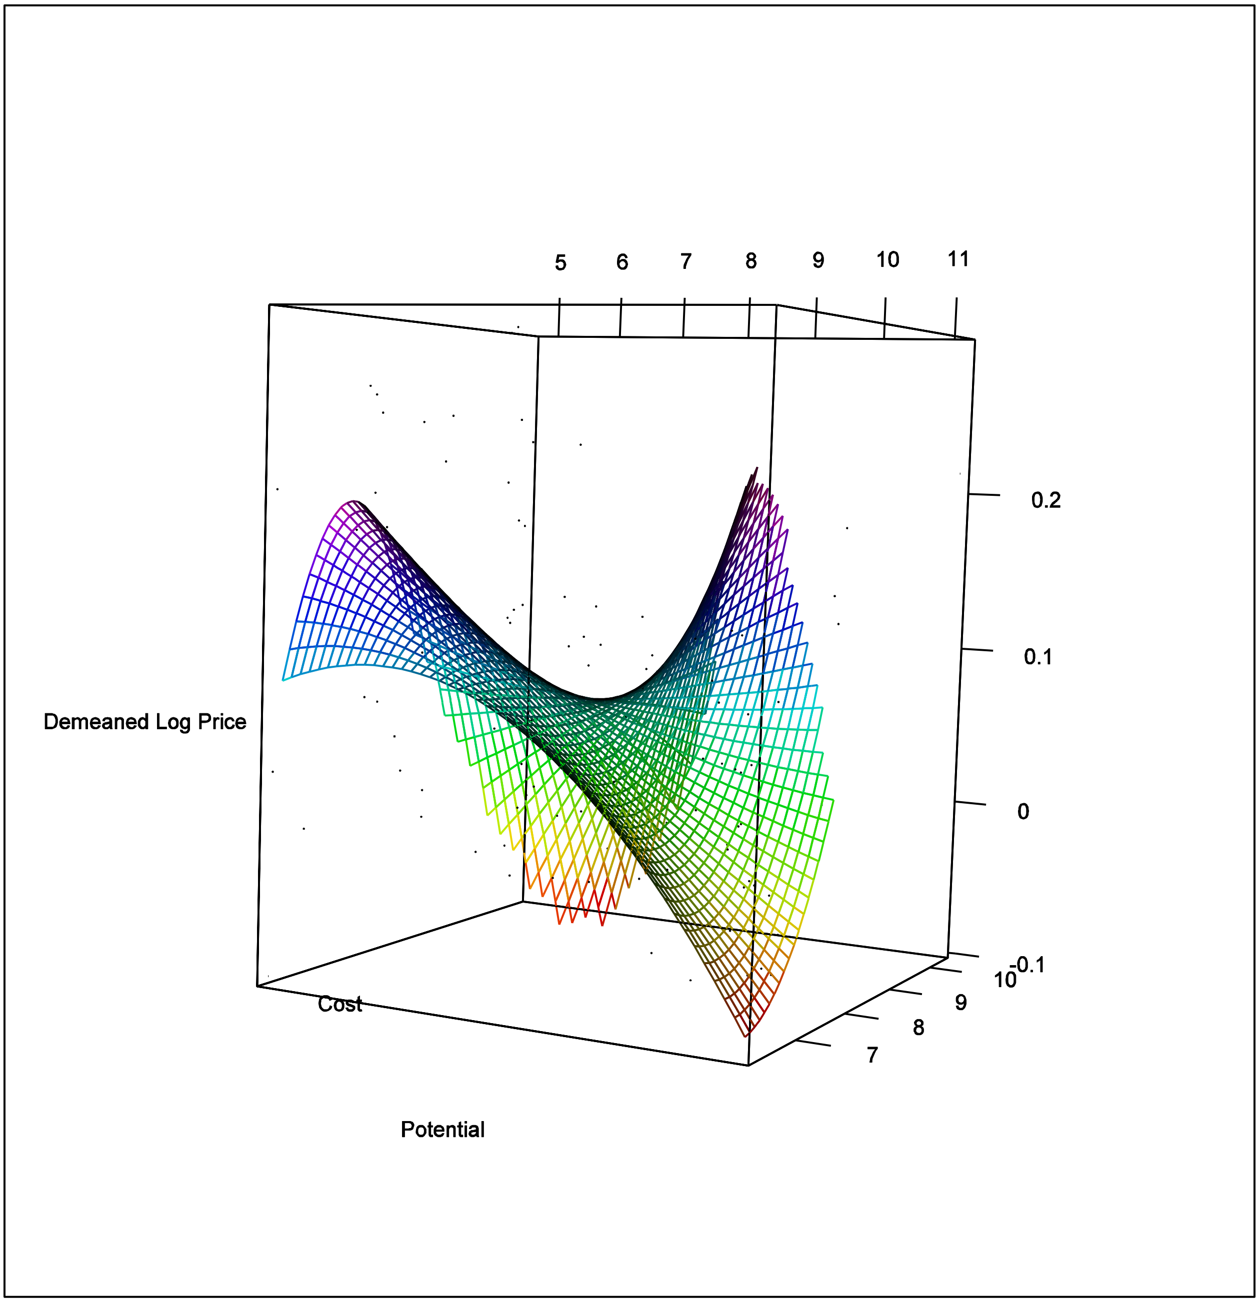

Figure 14. Orthogonal polynomial surface, at next higher order, non-monotonically fitting over instruments.

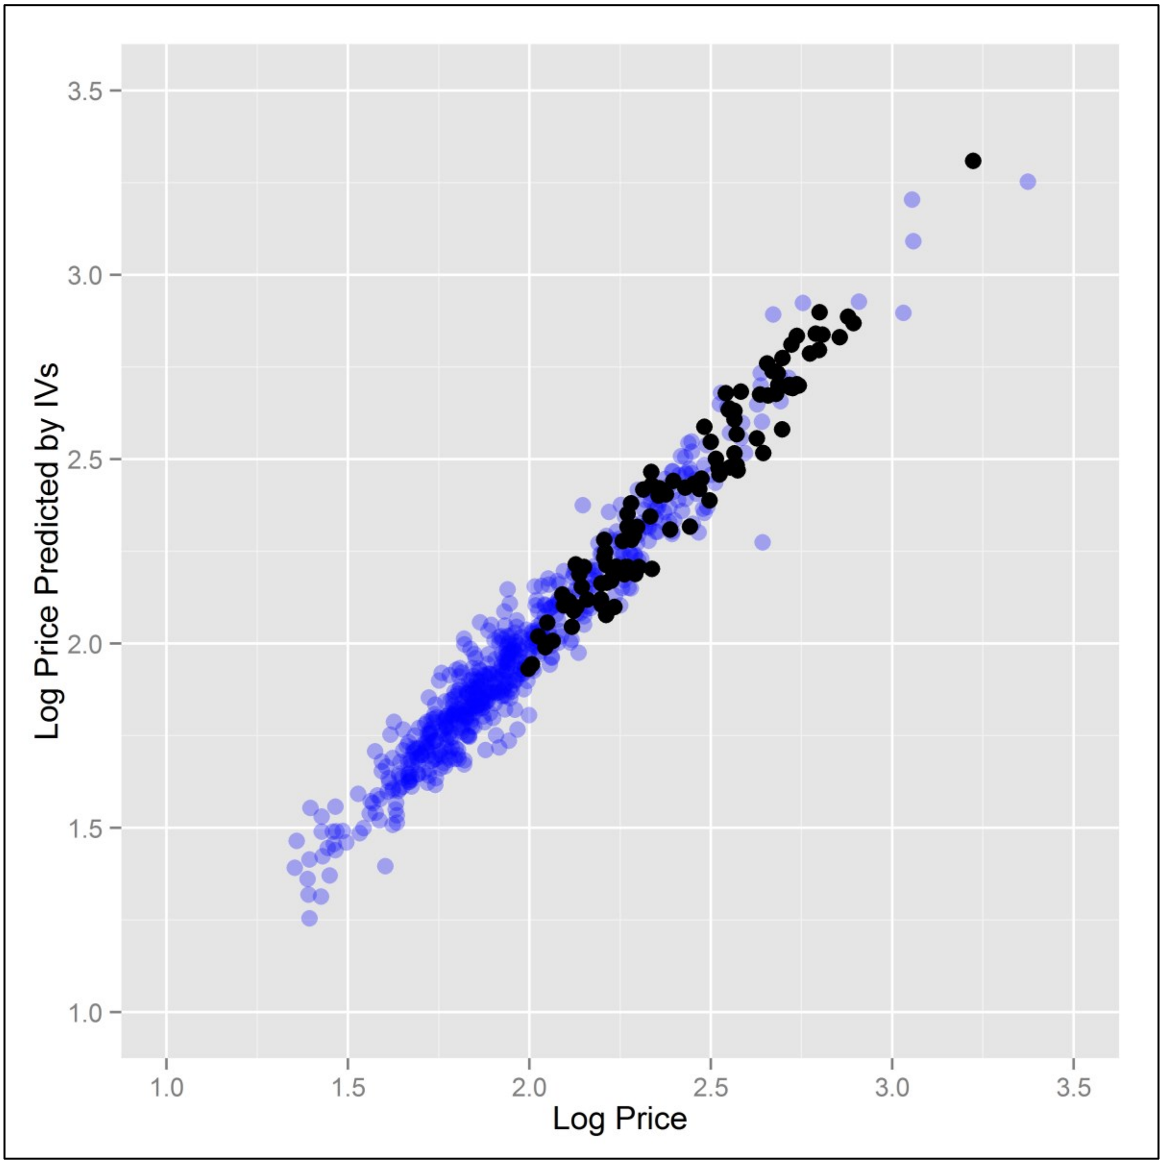

Figure 15. Observed log price versus log price as predicted by the instruments. States without an RPS (yet) are plotted in a semi-transparent blue and states with an RPS are plotted in an opaque black.

Industry Name	3-digit NAICS Code	GDP - Energy Price Elasticity	Elasticity's Standard Error	Time Trend	Trend's Standard Error
Motor vehicles, trailers, & parts	121	-0.752***	(0.160)	-83.654***	(13.768)
Paper products	129	-0.308**	(0.123)	13.185	(10.615)
Oil & gas extraction	107	-0.297	(0.206)	182.240***	(17.745)
Apparel & leather & allied products	128	-0.287**	(0.125)	-108.968***	(10.770)
Wood products	114	-0.199***	(0.068)	-7.844	(5.845)
Mining, except oil & gas	108	-0.184*	(0.105)	67.915***	(9.047)
Farms	104	-0.152**	(0.072)	80.112***	(6.232)
Furniture & related products	123	-0.151**	(0.069)	25.078***	(5.916)
Nonmetallic mineral products	115	-0.104*	(0.063)	10.163*	(5.403)
Printing & related support activities	130	-0.104**	(0.044)	-14.108***	(3.773)
Construction	111	-0.098**	(0.047)	78.276***	(4.056)
Primary metals	116	-0.086	(0.123)	11.814	(10.617)
Forestry, fishing, & related activities	105	-0.073*	(0.044)	39.771***	(3.800)
Chemical products	132	-0.070	(0.095)	73.130***	(8.193)
Food & beverage & tobacco products	126	-0.047	(0.048)	64.528***	(4.101)
Fabricated metal products	117	-0.029	(0.052)	38.230***	(4.451)
Other transportation equipment	122	-0.028	(0.101)	98.339***	(8.704)
Plastics & rubber products	133	-0.004	(0.066)	34.294***	(5.707)
Miscellaneous manufacturing	124	0.042	(0.073)	76.473***	(6.262)
Petroleum & coal products	131	0.078	(0.164)	196.590***	(14.176)
Machinery	118	0.091	(0.074)	33.743***	(6.393)
Textile mills & textile product mills	127	0.154	(0.094)	-31.217***	(8.126)
Support activities for mining	109	0.246	(0.174)	190.038***	(15.022)
Utilities	110	0.280***	(0.024)	61.878***	(2.075)
Computer & electronic products	119	0.349***	(0.126)	69.562***	(10.863)
Electrical equip, appliances, & comp	120	0.370***	(0.094)	22.628***	(8.087)
Government	178	0.033**	(0.013)	99.506***	(1.134)

Table 3. Parameter estimates of output's elasticity to energy price and the industry-wide trend in output from second stage (i.e. IV with a SUR FGLS weighting, hence called 3SLS by some) for goods producing industries.

Industry Name	3-digit NAICS Code	GDP - Energy Price Elasticity	Elasticity's Standard Error	Time Trend	Trend's Standard Error
Funds, trusts, & other financial vehicles	154	-0.643***	(0.131)	245.812***	(11.305)
Pipeline transportation	142	-0.451***	(0.124)	72.341***	(10.683)
Fed Reserve banks, credit, & related	151	-0.325***	(0.051)	134.974***	(4.427)
Publishing industries (includes software)	146	-0.269***	(0.055)	101.187***	(4.716)
Perform arts, spectator sports, & related	172	-0.235***	(0.057)	124.159***	(4.882)
Rental & leasing services (etc)	157	-0.191***	(0.059)	126.670***	(5.122)
Accommodation	175	-0.153***	(0.035)	87.514***	(3.008)
Computer systems design & related	160	-0.128***	(0.049)	150.610***	(4.209)
Information & data processing services	149	-0.104	(0.115)	197.356***	(9.955)
Truck transportation	140	-0.103***	(0.030)	68.447***	(2.627)
Administrative & support services	164	-0.078***	(0.028)	118.355***	(2.399)
Food services & drinking places	176	-0.075***	(0.019)	108.098***	(1.671)
Waste management & remed services	165	-0.075*	(0.045)	126.042***	(3.868)
Retail trade	135	-0.070***	(0.019)	69.577***	(1.676)
Social assistance	170	-0.059**	(0.027)	144.433***	(2.347)
Securities, commodities, & investments	152	-0.056	(0.055)	72.015***	(4.743)
Legal services	159	-0.046	(0.031)	106.523***	(2.708)
Other transportation & support activities	143	-0.040	(0.031)	101.790***	(2.642)
Ambulatory health care services	168	-0.035**	(0.017)	129.987***	(1.428)
Educational services	166	-0.035	(0.029)	143.283***	(2.493)
Insurance carriers & related	153	-0.035	(0.038)	109.678***	(3.318)
Misc prof, scientific, & tech services	161	-0.034	(0.027)	133.406***	(2.356)
Transit & ground passenger transp	141	-0.024	(0.047)	92.285***	(4.021)
Rail transportation	138	-0.020	(0.077)	59.831***	(6.632)
Motion picture & sound record indus	147	-0.018	(0.080)	85.245***	(6.907)
Other services, except government	177	-0.011	(0.017)	67.623***	(1.469)
Wholesale trade	134	-0.002	(0.021)	73.758***	(1.799)
Amusements, gambling, & rec industries	173	-0.001	(0.043)	59.910***	(3.728)
Hospitals & nursing & res care facilities	169	0.002	(0.018)	130.898***	(1.514)
Real estate	156	0.006	(0.024)	95.671***	(2.056)
Broadcasting & telecommunications	148	0.033	(0.034)	71.555***	(2.968)
Management of comp & enterprises	162	0.103	(0.074)	121.936***	(6.403)
Air transportation	137	0.125*	(0.069)	43.943***	(5.930)
Warehousing & storage	144	0.138**	(0.069)	100.554***	(5.970)
Water transportation	139	0.752***	(0.144)	27.282**	(12.436)

Table 4. Parameter estimates of output's elasticity to energy price and the industry-wide trend in output from second stage (i.e. IV with a SUR FGLS weighting, hence called 3SLS by some) for service industries.

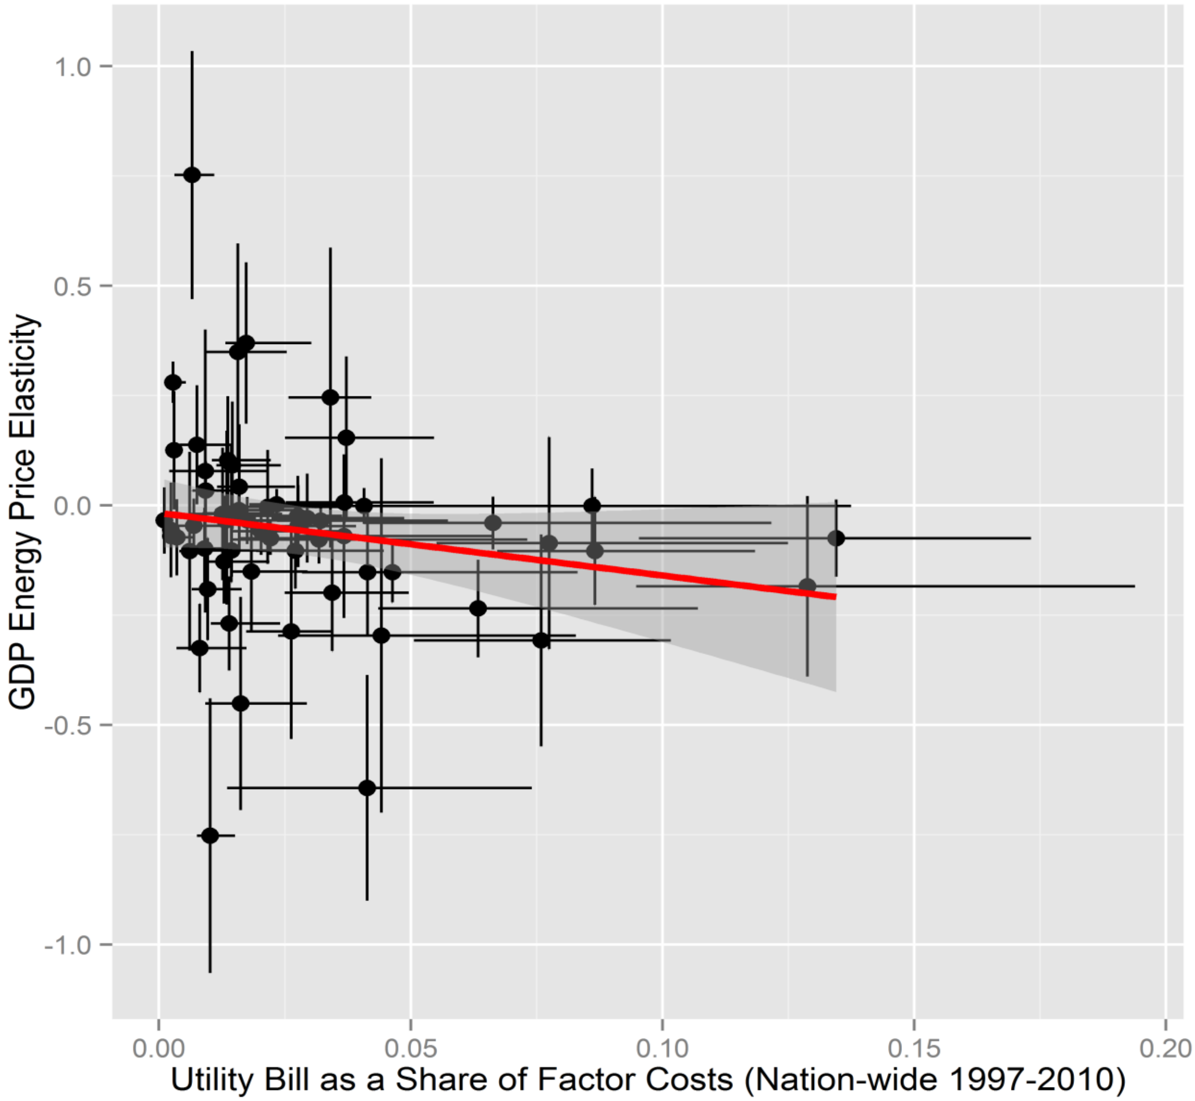

Figure 16. Relationship across 62 industries between estimated elasticities and factor cost shares with uncertainty bars in the verticle direction and variability bars (in the form of observed range) in the horizontal direction.

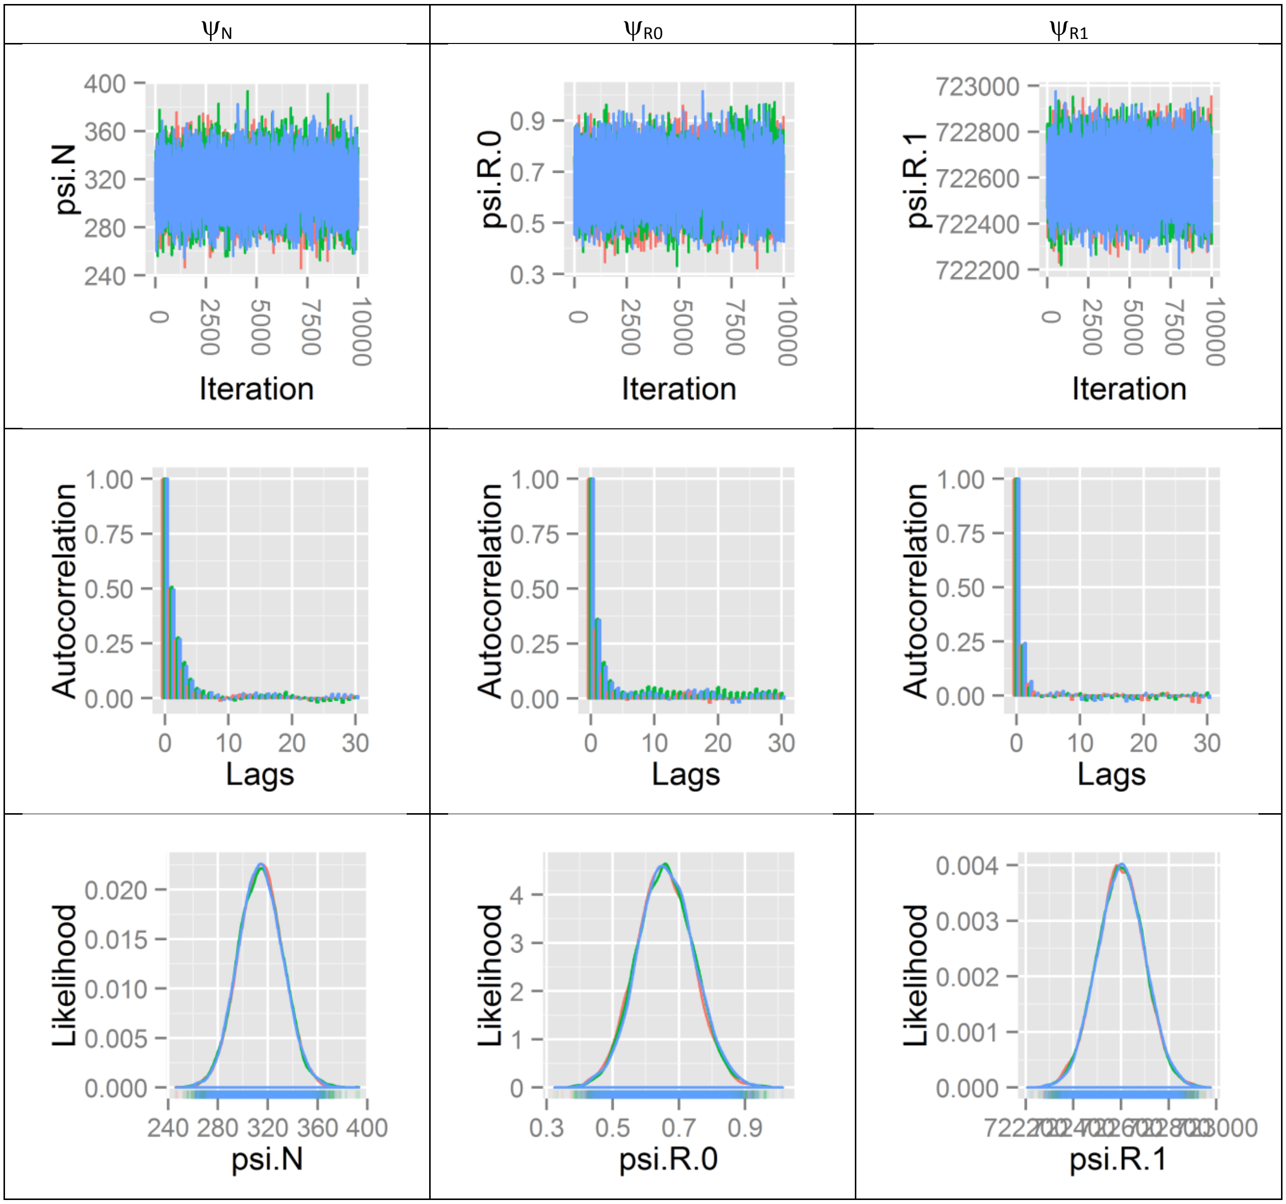

Figure 17. Convergence diagnostics for the model's 3 key parameters.

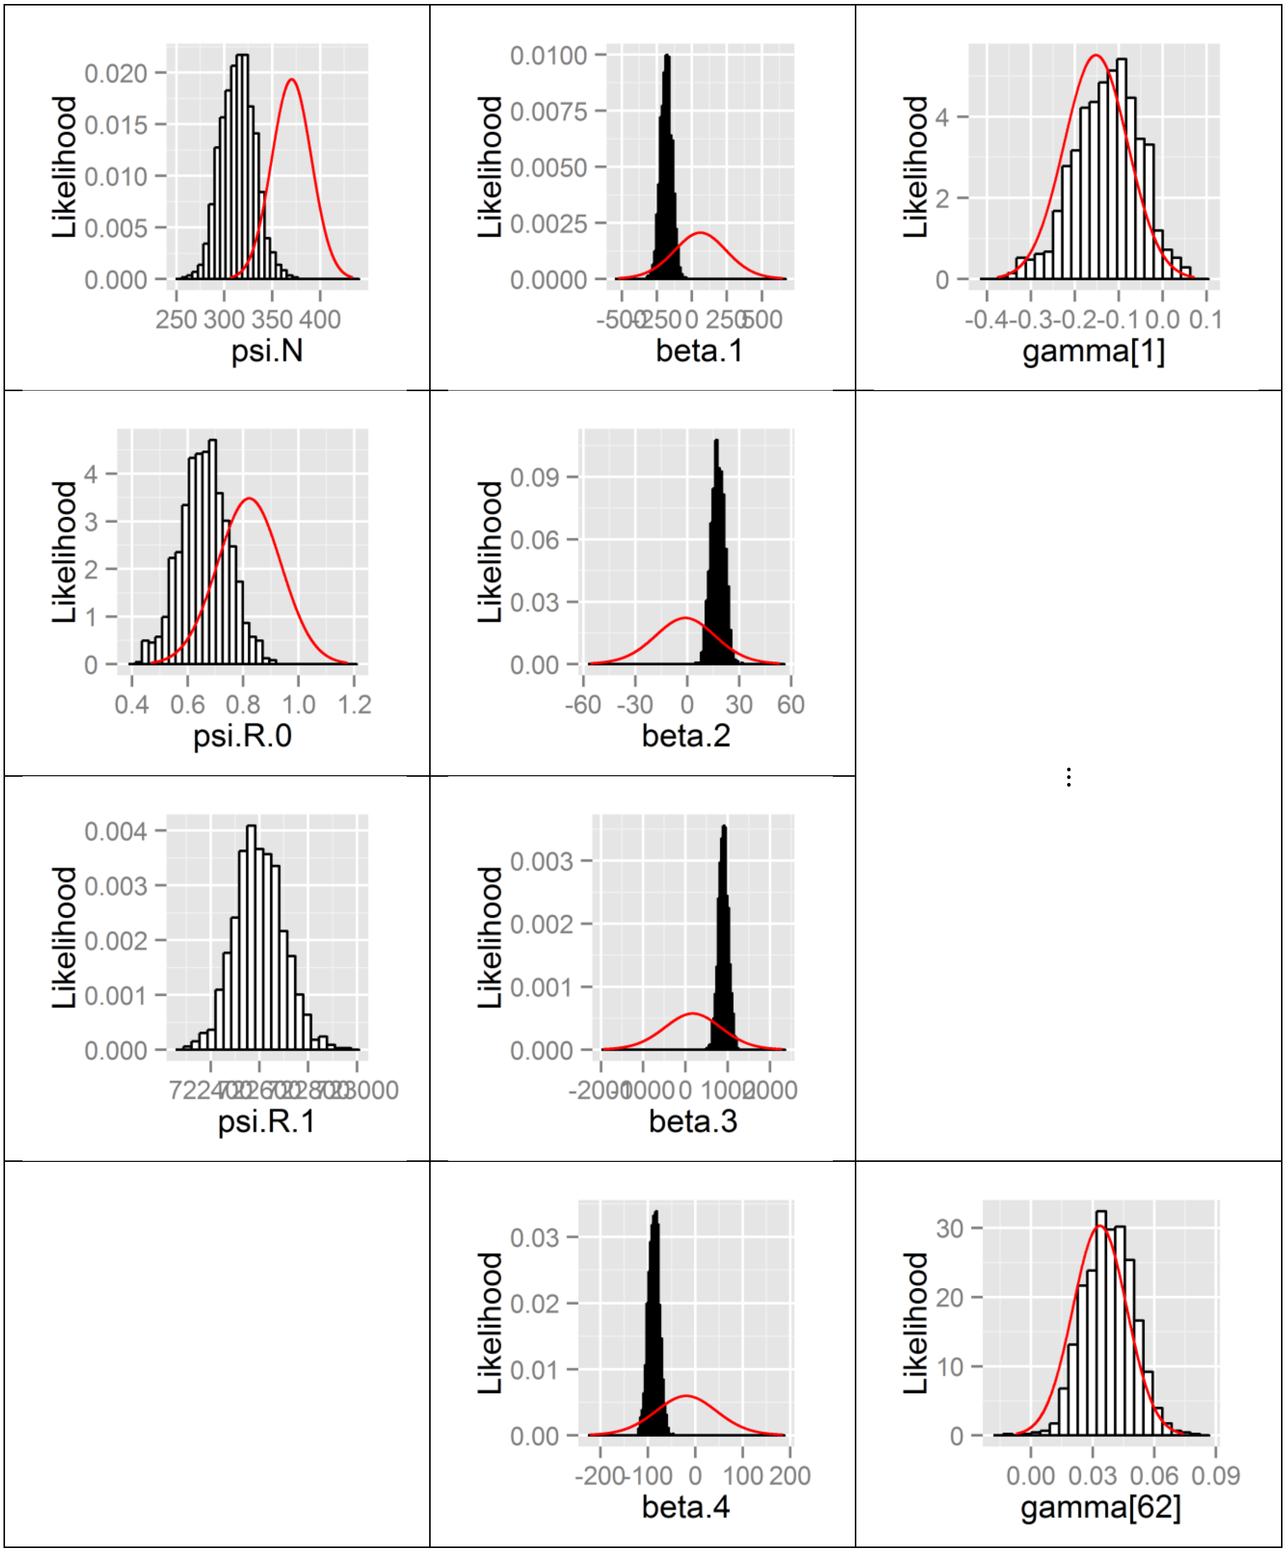

Figure 18. Histograms of [virtually] independent MCMC draws on parameter estimates and red densities for reduced form uncertainty distributions.

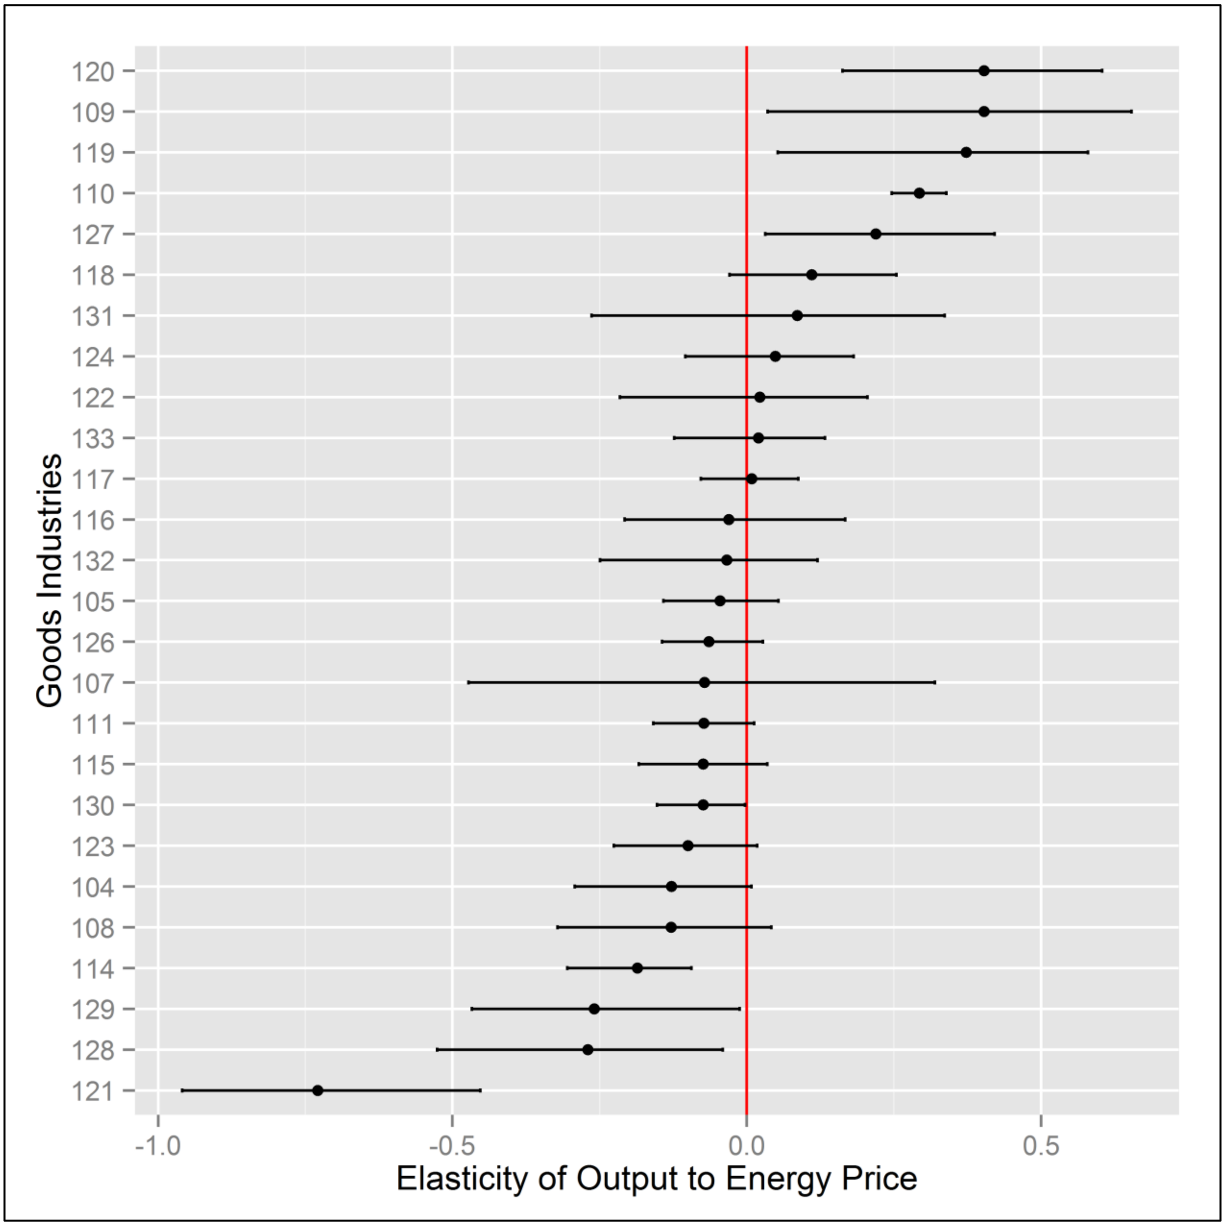

Figure 19. Bayesian posterior estimates of good industries' elsaticities as updated priors from the classical 3SLS estimation.

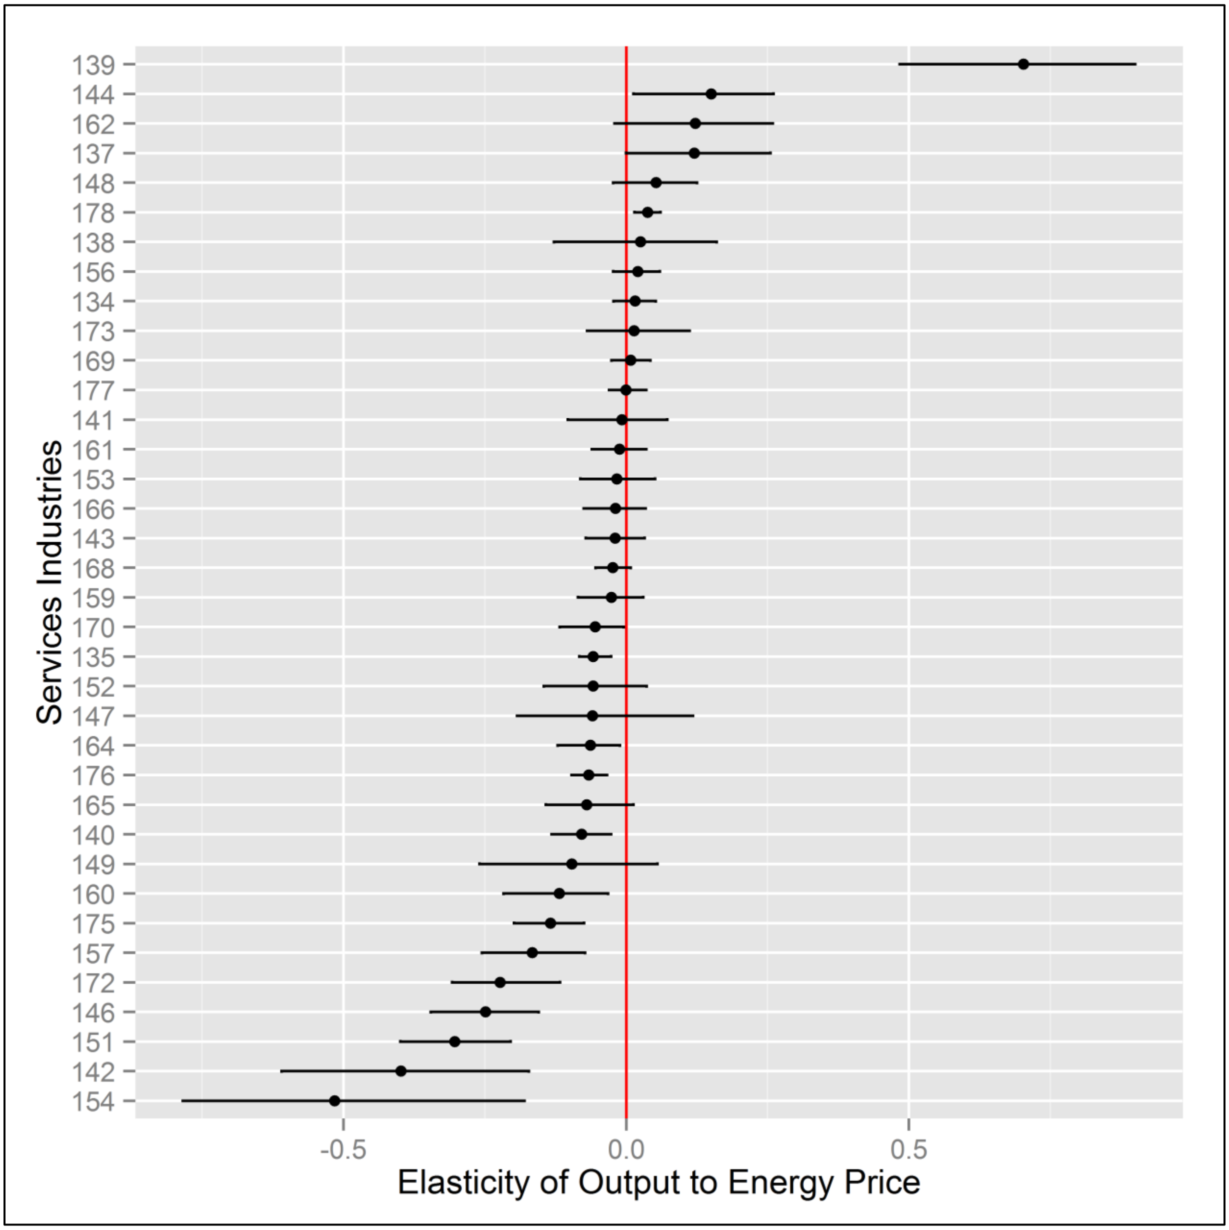

Figure 20. Bayesian posterior estimates of service industries' elsaticities as updated priors from the classical 3SLS estimation.

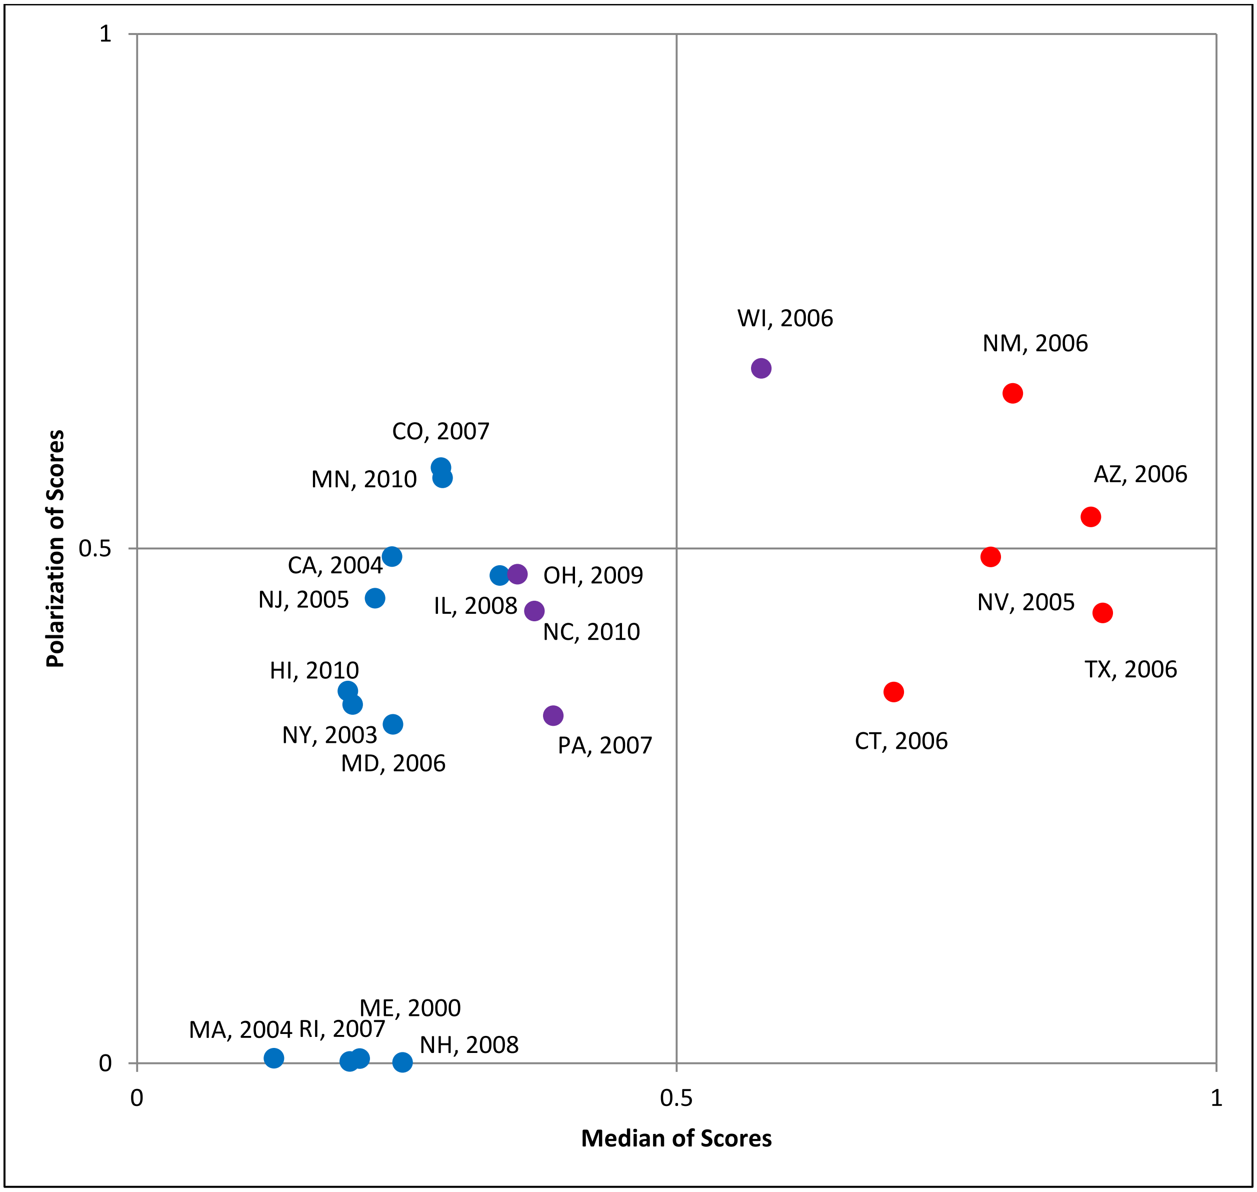

Figure 21. State of political scores at the time that a state adopted an RPS mandate (1997-2010). Delaware and Montana are not pictured because the polarization score was unavailable for those states due to the fact that each of those states has only a single representative in Congress. Nevertheless, their median scores at the time of adopting an RPS mandate was right around a strongly conservative 0.85 (just to the left of Arizona).