Measurement of renewable–energy efficiency and convergence across the US states: evidence from a non–parametric frontier analysis

Lei Pan*[1] and Richard Adjei Dwumfour[2]

[1] *Curtin University and Monash University*
[2] *Curtin University*

Abstract

This paper examines whether the clean-energy transition is producing convergence in renewable-energy efficiency across the United States. Using a balanced panel of the 50 states and the District of Columbia over 2010–2022, we measure renewable-energy efficiency with three non-parametric frontier estimators: conventional DEA under variable returns to scale, robust order-m, and convexified order-m. Renewable electricity generation is treated as the desirable output, while fossil-fuel generation, electric power-sector CO_2 emissions, and total nameplate generating capacity are treated as inputs to be minimised. We then apply the Phillips–Sul log-t test and stochastic-kernel methods to examine convergence dynamics. The results reject full-sample convergence under all three efficiency indicators. Instead, US states form persistent convergence clubs, with high-performing hydro, wind-belt, and low-load states separated from fossil-reliant laggards. Stochastic kernels reveal strong intra-distribution persistence: states largely preserve their relative positions between 2010 and 2022. The findings imply that uniform clean-energy incentives are unlikely to close cross-state renewable-efficiency gaps. Federal and state policy should instead target club-specific constraints, including transmission bottlenecks, fossil lock-in, grid integration, and uneven renewable-resource deployment.

Keywords: Renewable-energy efficiency; Order-m frontier; Convergence clubs; US states; Clean-energy policy
JEL Codes: Q42; C14; C23; R11

1 Introduction

In the public imagination, the clean-energy transition is often told as a story of heroic expansion: more wind turbines on the plains, more solar panels across rooftops and deserts, and more renewable electricity entering the grid each year. This narrative is powerful, but it is incomplete. A state can produce a large volume of renewable electricity and still remain heavily dependent on fossil generation, carbon-intensive power-sector emissions, and an oversized or under-utilised generation fleet. Conversely, a smaller state may produce fewer renewable gigawatt-hours in absolute terms but do so with a much lighter fossil burden and a leaner power system. For energy policy, therefore, the central question is not only how much renewable electricity a jurisdiction produces, but how efficiently it transforms its power system toward renewable generation.

*Corresponding author. Contact: School of Accounting, Economics and Finance, Curtin University, Perth, Australia, WA-6102. Email: lei.pan@curtin.edu.au

This distinction matters especially in the United States. The US electricity transition has never been a single national experiment implemented uniformly across space. It is a patchwork of state-level renewable portfolio standards, regional transmission constraints, heterogeneous resource endowments, different degrees of fossil-fuel lock-in, and uneven political commitments to decarbonisation. A wind-rich state in the Great Plains, a hydro-dependent state in the Pacific Northwest, a solar-abundant state in the Southwest, and a fossil-reliant state in the Southeast all face the same national decarbonisation imperative, but they do not face the same production possibility set. The passage of the Inflation Reduction Act has sharpened this issue by making federal clean-energy incentives more generous and durable, yet even large federal subsidies do not automatically imply that all states will converge toward the same renewable-energy performance frontier (Bistline et al., 2023). The policy question is therefore fundamentally spatial: are US states becoming more similar in renewable-energy efficiency, or is the transition creating persistent clubs of leaders, followers, and laggards?

This paper answers that question by measuring renewable-energy efficiency across the 50 US states and the District of Columbia over 2010–2022 and by testing whether states converge toward a common efficiency path. We define renewable-energy efficiency as the ability of a state to generate renewable electricity while minimising three system burdens: fossil-fuel generation, power-sector carbon dioxide emissions, and total installed generating capacity. This definition follows the eco-efficiency logic that desirable outputs should be expanded while environmental and resource burdens should be reduced (Kuosmanen and Kortelainen, 2005). In this setting, renewable electricity is the desirable output, while fossil generation, electric power-sector CO_2, and total nameplate capacity are inputs to be economised. The resulting measure is not a conventional renewable-share indicator, nor is it a simple emissions intensity measure. It is a frontier-based indicator of how close each state is to the best observed practice in transforming generation resources into renewable electricity.

The empirical setting is particularly well suited to this question. US states operate within the same federal system and are exposed to broadly similar national tax incentives, but they differ sharply in policy design, resource potential, grid structure, industrial composition, and legacy fuel dependence. Previous studies show that state-level policy instruments have played an important role in renewable deployment. Bird et al. (2005) document that wind development in the United States has been shaped by a combination of federal incentives, state policies, and market conditions. Menz and Vachon (2006) show that renewable portfolio standards and other state-level policies are associated with wind-power development. Carley (2009) evaluates state renewable electricity policies and finds evidence that policy adoption matters for renewable generation outcomes. Yin and Powers (2010) further show that the design and stringency of renewable portfolio standards influence in-state renewable generation, while Delmas and Montes-Sancho (2011) emphasise that policy effectiveness depends on state context. Similarly, Shrimali and Kniefel (2011) assess the effects of state policies on renewable electricity deployment and conclude that policy impacts differ by technology and institutional setting. Beyond the US, Johnstone et al. (2010) show that renewable-energy policies stimulate technological innovation, Marques et al. (2010) identify energy security and carbon-reduction motives as drivers of renewable-energy use in Europe, and Fischer and Newell (2008) demonstrate that technology and environmental policies create different incentives for emissions reduction and renewable innovation. This literature explains why renewable deployment differs across jurisdictions, but it has paid less attention to whether jurisdictions are becoming more similar in frontier-based renewable efficiency.

A second strand of literature provides the methodological foundation for measuring such efficiency. The starting point is Farrell (1957), who formalised productive efficiency as a distance from the best-practice frontier. Charnes et al. (1978) introduced data envelopment analysis (DEA) as a non-parametric method for evaluating decision-making units with multiple inputs and outputs, and Banker et al. (1984) extended DEA to variable returns to scale. Environ-

mental and energy applications then adapted this frontier logic to joint production systems in which desirable outputs coexist with undesirable outputs. Färe et al. (1989) introduced a non-parametric approach to productivity comparisons when some outputs are undesirable, Chambers et al. (1996a) developed benefit and distance functions, and Chung et al. (1997) used directional distance functions to handle productivity measurement with undesirable outputs. In energy applications, Hu and Wang (2006) introduced the total-factor energy-efficiency framework for Chinese regions, Zhou et al. (2008b) surveyed DEA applications in energy and environmental studies, Zhou et al. (2008a) examined environmental performance under different DEA technologies, and Zhou and Ang (2008) developed linear-programming models for measuring economy-wide energy efficiency. Broader reviews by Sueyoshi and Goto (2012), Sueyoshi and Goto (2017), and Mardani et al. (2017) confirm that DEA has become a central tool for evaluating energy and environmental performance.

Despite this progress, two methodological challenges remain. First, conventional DEA can be highly sensitive to outliers because the estimated frontier is determined by extreme observations. This is particularly important in the US renewable-energy setting, where small-load states with very low fossil generation may appear unusually efficient relative to large fossil-intensive states. Robust and probabilistic frontier estimators provide a way forward. Cazals et al. (2002) introduce the order-m estimator as a robust non-parametric frontier method, while Daraio and Simar (2007a) develop conditional non-parametric frontier methods that allow richer frontier comparisons. Inference issues in non-parametric frontier analysis are further addressed by Simar and Wilson (1998) and Simar and Wilson (2007). Recent renewable-energy performance studies also emphasise the need for robust frontier methods under data uncertainty (Li et al., 2024). Second, even when efficiency can be measured, it is not obvious whether the cross-sectional distribution is converging. A declining average gap is not the same as convergence, and a stable mean can hide the formation of separate convergence clubs. The convergence literature shows that distributional dynamics can be persistent and multi-peaked (Quah, 1996), while the Phillips–Sul framework provides a formal method for testing whether units share a common transition path and for identifying convergence clubs when full convergence fails (Phillips and Sul, 2007, 2009). Kounetas et al. (2021) apply a related non-parametric frontier and convergence approach to US eco-efficiency, showing that environmental performance may converge within clubs rather than across the full sample.

This paper contributes to the literature in three ways. First, it shifts the empirical focus from renewable-energy deployment to renewable-energy efficiency. Many policy discussions rank states by renewable generation or renewable shares, but these indicators do not account for the fossil generation, emissions, and installed capacity required to support the observed renewable output. By treating fossil generation, electric-sector CO_2, and total capacity as inputs to be minimised, the paper provides a policy-relevant measure of how efficiently each state converts its power-sector structure into renewable electricity. This is especially useful for policy makers because it links renewable performance directly to the practical burdens faced by system planners: fossil dependence, carbon emissions, and capacity overhang.

Second, the paper combines three complementary empirical tools: conventional DEA under variable returns to scale, robust order-m frontier estimation, and convexified order-m estimation. This triangulation matters because the renewable transition contains both genuine frontier leaders and potential outliers. A single deterministic DEA score may overstate the role of extreme observations, while a robust probabilistic frontier can reveal whether high-performing states remain efficient when compared against randomly drawn peer groups. Our paper therefore offers a more credible measurement framework for renewable-energy efficiency than a single index or a simple ratio measure.

Third, the paper connects renewable-energy efficiency measurement with convergence-club analysis. Instead of asking only which states are efficient in a given year, it asks whether the entire cross-section is moving toward a common long-run efficiency path. The Phillips–Sul log-t

test is used to test full-sample convergence and to recover convergence clubs, while stochastic kernels are used to visualise distributional persistence and long-run dynamics. This allows the analysis to distinguish between three very different policy worlds: a world of national catch-up, a world of stable inequality, and a world of club convergence in which states improve within groups but remain separated across groups.

The main results point strongly toward the third world. The full-sample convergence hypothesis is rejected, indicating that US states do not share a common renewable-efficiency transition path. Instead, the data reveal a set of convergence clubs. High-performing states include hydro-rich, wind-belt, and low-load states that generate renewable electricity with comparatively small fossil and emissions burdens. Mid-ranked states display stable within-club transition paths. At the bottom, fossil-reliant states with large legacy generation fleets remain far from the frontier. Stochastic-kernel evidence reinforces this interpretation: the distribution of renewable efficiency is highly persistent, with probability mass concentrated around the diagonal rather than rotating toward a common mean. In policy terms, the transition is not simply a national race in which all states gradually close the same gap. It is a stratified process in which initial conditions, resource endowments, and policy environments continue to shape long-run performance.

These findings have direct implications for clean-energy policy design. If renewable-efficiency convergence were occurring nationally, uniform federal incentives might be sufficient to close cross-state gaps. The evidence here suggests otherwise. Stable convergence clubs imply that marginal policy needs differ across the distribution. Frontier states may require transmission expansion, storage integration, and market-design reforms to absorb additional renewable generation efficiently. Middle-club states may benefit from policies that accelerate technology diffusion and reduce interconnection delays. Laggard states may need more targeted support: grid investment, coal and gas transition planning, stronger renewable procurement rules, and federal financing instruments designed to overcome fossil lock-in. The key policy message is therefore not that all states should be treated identically, but that federal and state policy should recognise where each state sits in the renewable-efficiency distribution and which club-specific constraint prevents further catch-up.

The rest of the paper is organised as follows. Section 2 describes the data and the construction of the renewable-energy efficiency indicators, and presents the non-parametric frontier methods, the Phillips–Sul convergence test, and the stochastic-kernel approach. Section 3 reports the efficiency scores and distributional dynamics. Section 4 presents the convergence-club results and robustness checks. Section 5 concludes with policy implications for state and federal renewable-energy strategy.

2 Data and Empirical Methodology

2.1 Data

The empirical analysis is based on a balanced panel of $51 \times 13 = 663$ state–year observations covering the 50 US states plus the District of Columbia (henceforth "DC") over the period 2010–2022. We rely entirely on publicly available data published by two US federal statistical agencies: the U.S. Energy Information Administration (EIA) and the U.S. Bureau of Economic Analysis (BEA). All five variables, their units, their primary source tables, and the corresponding URLs are listed in Table 1; pooled summary statistics are reported in Table 2.

The desirable output is annual net renewable electricity generation, constructed as the sum of utility–scale wind, solar (photovoltaic and solar–thermal), conventional hydroelectric, geothermal, and biomass net generation, expressed in gigawatt–hours. The figures are taken from the EIA *State Electricity Profiles* and consolidated through the EIA's State Energy Data System (SEDS) renewable–electricity series. We use net rather than gross generation, so plant–level station service is already netted out. Distributed (small–scale) solar PV is excluded from the

Table 1: Variables and data sources.

Variable (unit)	Role in DEA	Source	Specific dataset / table
Renewable electricity generation (GWh)	Output (v)	EIA (U.S. Energy Information Administration, 2023a)	*State Electricity Profiles*, Table 5 "Electric power industry generation by primary energy source", summed across wind, utility–scale solar PV/thermal, conventional hydroelectric, geothermal, and biomass categories. Series code: SEDS `REPRB.[STATE].A` (renewable energy production, electric power sector).
Fossil–fuel net generation (GWh)	Input (d_1)	EIA (U.S. Energy Information Administration, 2023a)	Same Table 5, summed across coal, petroleum, natural gas, and other gases categories.
CO_2 emissions from the electric power sector (thousand metric tonnes)	Input (d_2)	EIA (U.S. Energy Information Administration, 2023b)	*State Energy–Related Carbon Dioxide Emissions*, electric power sector total. Available at `https://www.eia.gov/environment/emissions/state/`.
Total nameplate generating capacity (MW)	Input (d_3)	EIA (U.S. Energy Information Administration, 2023a)	*State Electricity Profiles*, Table 4 "Existing generation capacity", sum across all primary energy sources.
Real state GDP (millions of chained 2017 USD)	Auxiliary	BEA (U.S. Bureau of Economic Analysis, 2023)	Regional Economic Accounts, "GDP by State" release, Table SAGDP9N (real GDP by state, all industries).

Table 2: Summary statistics, pooled 2010–2022 (N=663).

Variable	Mean	Std. dev.	Min	Max
Renewable generation (GWh)	11,722	18,493	36	117,438
Fossil–fuel generation (GWh)	47,659	45,932	95	268,510
CO_2 from power sector (kt)	33,710	35,018	26	191,804
Total nameplate capacity (MW)	22,710	22,170	51	134,885
Real state GDP (m. 2017 USD)	376,221	471,832	31,892	2,932,910

headline series to avoid the structural break introduced when EIA began reporting it separately in 2014; including small–scale solar modifies the levels for California, Hawaii, and Arizona but does not change the convergence patterns we report.[1]

Following the eco–efficiency convention of Kuosmanen and Kortelainen (2005), three burdens are reinterpreted as inputs whose use should be minimised given the renewable output. First, *net fossil–fuel generation* (GWh) is the sum of coal, natural gas, petroleum, and other–gas generation from the same EIA Table 5. Second, CO_2 *emissions from the electric power sector* (thousand metric tonnes) is taken from the EIA's *State Energy–Related Carbon Dioxide Emissions* series; we restrict attention to the electric–power sector total to match the scope of the output and the other input variables. Third, *total nameplate generating capacity* (MW) is the sum of installed nameplate capacity across all primary energy sources from EIA Table 4. Capacity is included to penalise states holding large under–utilised fleets and to capture the intuition that, holding renewable output constant, a leaner overall fleet is more efficient.

Real state GDP (millions of chained 2017 USD) is from BEA's Regional Economic Accounts, specifically the "GDP by State" release (Table SAGDP9N, all industries). It is not an input to the DEA programme but is used in the discussion of policy implications and to anchor the state–size dimension of the choropleth in Figure 1.

[1] Results available on request.

The panel is balanced: every variable is observed annually for every state–year cell from 2010 to 2022, yielding $51 \times 13 = 663$ observations. The starting year (2010) is dictated by the desire to fully cover the period of accelerated wind and utility–scale solar deployment following the 2009 American Recovery and Reinvestment Act, and the ending year (2022) is the latest year for which all five series were available in their final, non–preliminary form at the time of the analysis.

2.2 Probabilistic frontier analysis

Let v denote the desired output (renewable generation) and let $d = (d_1, \ldots, d_p)$ denote the vector of inputs that should be minimised ($p = 3$ here). The renewable–generating technology set is:

$$\Psi = \{(d, v) \in \mathbb{R}_+^{p+1} \mid \text{output } v \text{ can be produced with input vector } d\}. \tag{1}$$

A state's renewable efficiency in the input direction is, following Farrell (1957),

$$\varepsilon(d, v) = \inf\{\theta \mid (\theta d, v) \in \Psi\}. \tag{2}$$

Conventional DEA under variable returns to scale (VRS) estimates the boundary of Ψ as the convex hull of the observed input–output combinations (Kuosmanen and Kortelainen, 2005):

$$\hat{\varepsilon}_{\mathrm{DEA}}(d_o, v_o) = \min\big\{\theta \mid \textstyle\sum_i \omega_i d_i \le \theta d_o,\ \sum_i \omega_i v_i \ge v_o,\ \sum_i \omega_i = 1,\ \omega_i \ge 0\big\}. \tag{3}$$

The order–m estimator (Cazals et al., 2002; Daraio and Simar, 2007b) replaces the deterministic envelope of Equation (3) with a probabilistic frontier. For a fixed integer m, draw m peer states with replacement from the conditional distribution of D given $V \ge v_o$, and compute the deterministic free–disposal hull (FDH) input efficiency of the unit against this random reference. Repeating B times and averaging yields:

$$\hat{\varepsilon}_m(d_o, v_o) = \frac{1}{B} \sum_{b=1}^{B} \inf\big\{\theta \mid d_o \theta \ge d_i^{(b)} \text{ for some peer } i\big\}, \tag{4}$$

where $d_i^{(b)}$ are the m peer inputs in draw b. The estimator can exceed unity for super–efficient units that dominate most of their random peers, which is exactly what we expect for very low–load states such as DC. The convexified order–m estimator $\hat{\varepsilon}_m^C$ is defined analogously but uses the DEA–VRS programme (3) on each random peer set rather than the FDH hull, restoring convexity at the cost of slightly more sensitivity to the random draws (Daraio and Simar, 2007b). By construction $\hat{\varepsilon}_{\mathrm{DEA}} \le \hat{\varepsilon}_m^C \le \hat{\varepsilon}_m$. We set m=20 and B=200 throughout (with B=100 for the more costly $\hat{\varepsilon}_m^C$). All efficiency scores are computed year–by–year so that the reference set is the contemporaneous frontier.

2.3 Phillips–Sul convergence test

Let X_{it} denote any of the three efficiency indicators for state i in year t. Phillips and Sul (2007) model the panel through a non–linear time–varying factor structure $X_{it} = \delta_{it}\mu_t$, where μ_t is the common factor and δ_{it} a state–specific loading. Convergence is the property $\delta_{it} \to \delta$ as $t \to \infty$ for all i, which they show can be tested through the log–t regression:

$$\log\big(H_1/H_t\big) - 2\log\log(t+1) = \hat{a} + \hat{b}\,\log t + \hat{u}_t, \qquad t = [rT]+1, \ldots, T, \tag{5}$$

where $H_t = N^{-1}\sum_i (h_{it}-1)^2$, $h_{it} = X_{it}/\bar{X}_{\cdot t}$, and $\hat{b} = 2\hat{a}$ in the asymptotic limit. We use $r = 0.3$ as recommended by Phillips and Sul. The null of convergence is rejected at the 5% level if the HAC t–statistic on $\hat{b}$ falls below -1.65.

When the full–sample test rejects, the iterative club algorithm of Phillips and Sul (2009) can recover convergence sub–groups. Briefly: i) sort units by their last–period value of X_{iT} in descending order; ii) form the largest core group at the top of the sort that still passes the log–t test; iii) sieve the remaining units one–by–one into the core if doing so does not break the test; iv) recurse on the leftover units. After this primary classification, pairs of consecutive clubs are merged whenever the joint log–t test does not reject convergence, yielding the final classification.

2.4 Distributional dynamics

Following Quah (1996), we treat the cross–section of efficiency scores as a stochastic process and estimate, by Gaussian kernels with plug–in bandwidth, the conditional density $f_\tau(y \mid x) = f(x,y)/f(x)$, where $x = \log \hat{\varepsilon}_{i,t}$ and $y = \log \hat{\varepsilon}_{i,t+\tau}$. The conditional density evaluated on a common grid yields a transition probability matrix P, and the long–run ergodic distribution is the left eigenvector of P associated with the unit eigenvalue.

3 Results and Discussion

3.1 Renewable–energy efficiency scores

The map in Figure 1 reports the order–m renewable efficiency for 2022. Bubble size is proportional to renewable generation; colour encodes the efficiency score. Because the score is unbounded above and Vermont and DC are extreme outliers, the colour scale is clipped at the 95th percentile so that the bulk of the distribution remains visually distinguishable; the actual scores for the four super–efficient units are reported as annotations next to their bubbles. The picture is striking. Vermont ($\hat{\varepsilon}_m \approx 20.2$) and DC ($\approx 46.1$) are extreme super–efficient outliers because they generate small but non–negligible renewable volumes against very small fossil and capacity bases. Idaho (≈ 3.1), South Dakota (≈ 2.5), and Maine (≈ 2.0) sit comfortably above unity. The mainstream of the distribution – including the wind–belt and Pacific giants Texas, California, Iowa, Oregon, and Washington – clusters in the 0.6–1.0 range. Florida, with its very large fossil fleet and a comparatively modest renewable rollout, is the laggard at $\hat{\varepsilon}_m \approx 0.08$, joined by Ohio, Indiana, and Georgia.

The aggregate dynamics, summarised in Figure 2, plot two time series jointly. The blue solid line (left axis, scale 1.80–2.20) is the cross–sectional mean of the order–m efficiency, $\bar{\hat{\varepsilon}}_{m,t} = N^{-1}\sum_i \hat{\varepsilon}_{m,it}$. The red dashed line (right axis, scale 9.5–15) is the Phillips–Sul cross–sectional dispersion measure $H_t = N^{-1}\sum_i \big(h_{it} - 1\big)^2$, where $h_{it} = \hat{\varepsilon}_{m,it}/\bar{\hat{\varepsilon}}_{m,t}$ is the relative transition path used in the log–t regression (5). By construction $H_t = 0$ would correspond to perfect cross–sectional equality and $H_t \to 0$ as $t \to \infty$ is the sufficient condition for full–sample convergence in the Phillips and Sul (2007) framework.

Three features of the figure are worth emphasising. First, neither series exhibits a clear trend. The cross–sectional mean rises modestly from 1.97 in 2010 to 2.01 in 2022, an increase of less than 2% over thirteen years, and dispersion falls only slightly, from $H_{2010} = 12.06$ to $H_{2022} = 11.48$, a drop of roughly 5%. There is, in short, no visible σ–convergence at the aggregate level. This is the first piece of descriptive evidence consistent with the formal log–t test, which rejects full–sample convergence in Section 3.3.

Second, the two series move in tight lockstep: the contemporaneous correlation between $\bar{\hat{\varepsilon}}_{m,t}$ and H_t is 0.87. Both peak in 2011 (mean 2.17, $H = 14.81$), both spike again in 2014 (mean 2.12, $H = 13.37$), both bottom out in 2018 (mean 1.82, $H = 9.86$). This positive co–movement is initially surprising, because under standard σ–convergence one expects dispersion to fall when the mean rises (laggards catch up faster than leaders). The opposite pattern here has a clear mechanical explanation: the four super–efficient states (Vermont, DC, Idaho, South Dakota) carry enormous weight in both moments. When the random–peer draws used by the order–m estimator happen to match these states with weaker peers, their efficiency scores spike,

Figure 1: Order–m renewable–energy efficiency by state, 2022

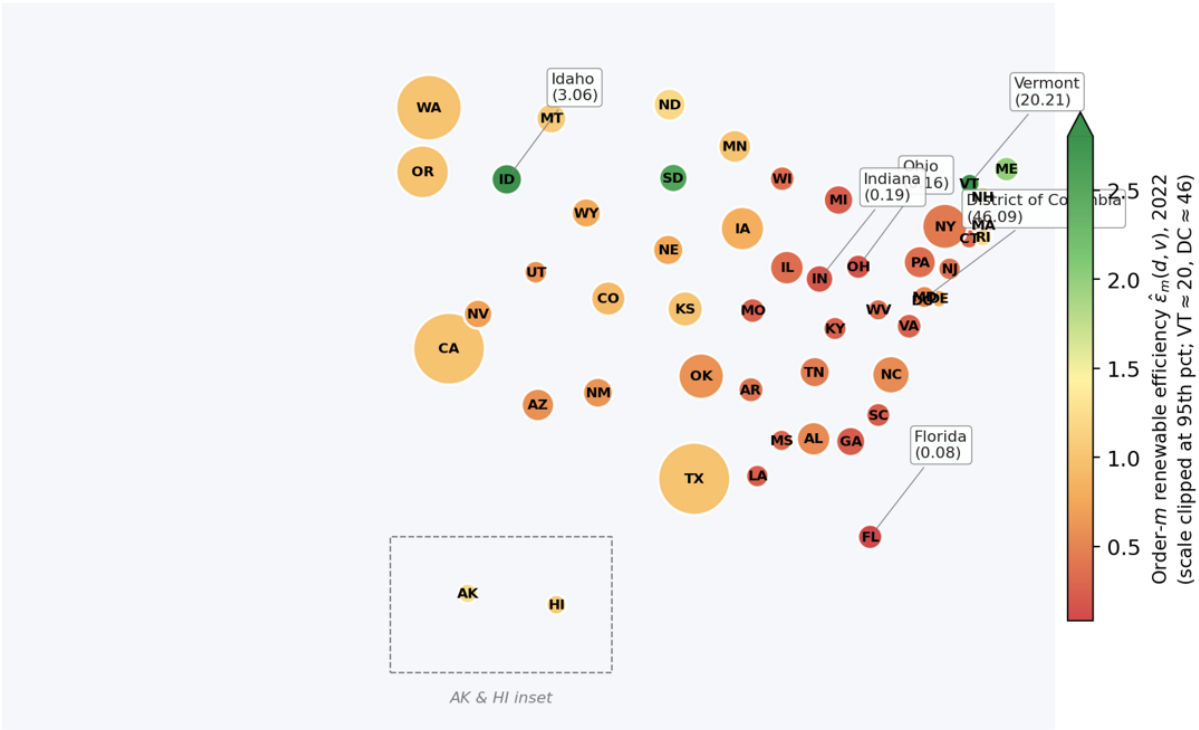

Note: Bubble size is proportional to total renewable generation; bubble colour is the order–m efficiency score, with the colour scale clipped at the 95th percentile (the saturated dark–green colour at the top of the bar therefore covers all values above approximately 2.5). The four super–efficient outliers and the laggard Florida are annotated with their actual $\hat{\varepsilon}_m$ values.

Figure 2: Average renewable efficiency $\bar{\hat{\varepsilon}}_{m,t}$ (blue solid line, left axis) and the Phillips–Sul cross–sectional dispersion $H_t = N^{-1}\sum_i (h_{it}-1)^2$ (red dashed line, right axis), 2010–2022

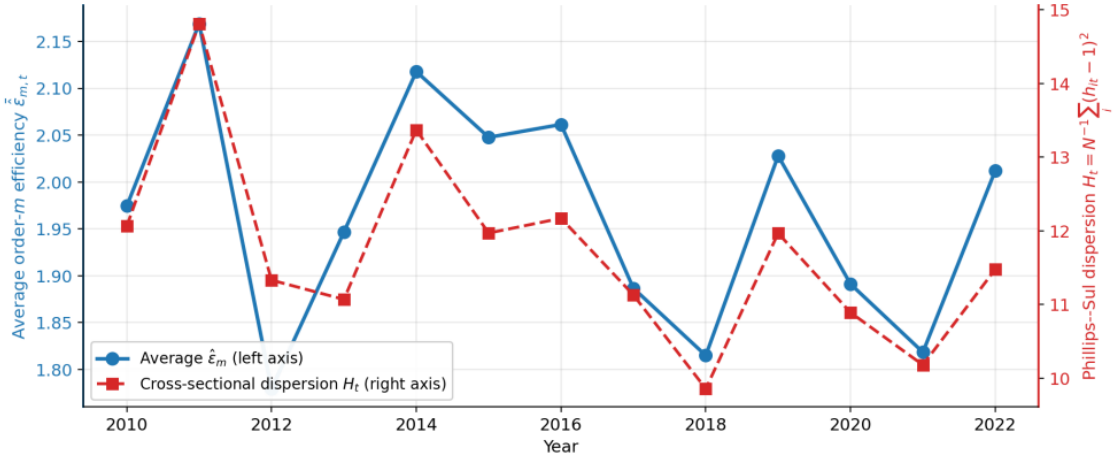

Note: Both axes use independent scales chosen to fill the plotting area; the coloured spines and tick labels match the colour of the corresponding series. The two series have a contemporaneous correlation of 0.87.

pulling both the mean and the dispersion up together. The co–movement therefore tells us that the cross–sectional distribution is dominated by a small number of outliers whose year–to–year trajectories drive aggregate dynamics, rather than by gradual catch–up among the bulk of the sample.

Third, the year–on–year volatility in H_t is substantial, the ratio of maximum to minimum dispersion across the sample is roughly 1.5 (14.81/9.86), and dwarfs the net long–run change. Any inference about "trends" in cross–sectional inequality of renewable efficiency over a sample of this length is therefore fragile. This motivates the formal log–t test in Section 3.3: rather than reading a trend off Figure 2, we let the data tell us, through the asymptotic theory of Phillips and Sul (2007), whether the dispersion is shrinking at a rate consistent with convergence. The answer is that it is not, at least not for the full sample, but several sub–groups of states do display the within–group convergence that the aggregate hides.

3.2 Distributional dynamics

Figure 3 reports the stochastic kernel of the order–m renewable efficiency indicator estimated between the start and the end of the panel, i.e. on the joint distribution of $(\log \hat{\varepsilon}_{m,i,2010}, \log \hat{\varepsilon}_{m,i,2022})$ for $i = 1, \ldots, 51$. The figure should be read as the empirical analogue of the transition operator $P_\tau(x, A)$ defined in Section 2. The height (3D, left) or colour brightness (contour, right) at coordinate (x, y) is the conditional density $\hat{f}_\tau(y \mid x)$, that is, the probability density that a state whose log–efficiency was x in 2010 is observed at y in 2022. The dashed white 45° line marks the locus $y = x$, where 2022 log–efficiency exactly equals 2010 log–efficiency: a state on this line has not changed its relative position. The far left and far right of the support are masked because the marginal density of 2010 log–efficiency falls below 5 % of its mode there and the conditional density is uninformative.

Figure 3: Stochastic kernel of the order–m renewable–efficiency indicator, 2010 vs. 2022

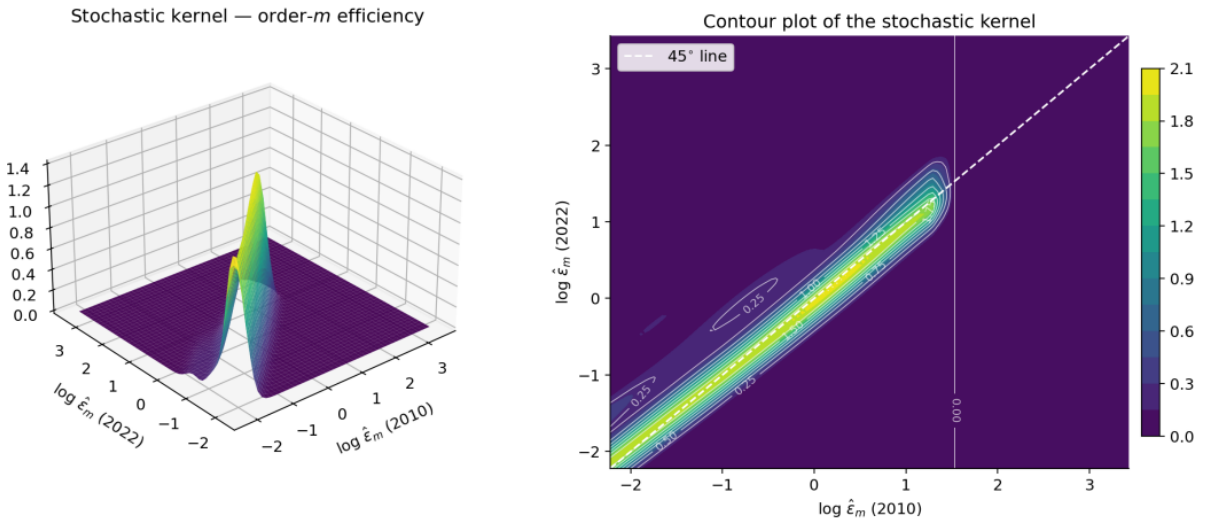

Note: **Left panel:** 3D conditional–density surface; the vertical axis is $\hat{f}_\tau(y \mid x)$. **Right panel:** contour plot of the same surface; brighter colours denote higher density. The dashed white line is the 45° identity line $y = x$, on which a state's 2022 log–efficiency equals its 2010 log–efficiency. Concentration of probability mass along this line indicates that states preserve their relative rank over the sample period.

Three substantive features of the kernel matter for the convergence question. First, the bright yellow ridge, the locus of maximum conditional density, sits exactly on top of the 45° line. There is no rotation away from the diagonal, no tilting, no "twisting" of the ridge into the off–diagonal region. In Quah (1996)'s language, this is a sign of strong intra–distribution

persistence: states tend to keep their relative rank over a 13–year horizon. The contour plot makes this even clearer; the 0.50 and 0.75 contours form narrow tubes hugging the diagonal, with no appreciable mass in the upper–left (laggards becoming leaders) or lower–right (leaders becoming laggards) regions.

Second, the ridge is not of uniform height along the diagonal: it displays a distinct primary peak at $(\log \hat{\varepsilon}_m \approx 1,\ 1)$ corresponding to mid–ranked states such as Texas, Iowa, and Oregon, and a smaller secondary mode at $(\log \hat{\varepsilon}_m \approx -1,\ -1)$ corresponding to the cluster of fossil–reliant laggards (Florida, Ohio, Indiana). This bimodality on the diagonal is the empirical signature of the "twin peaks" that Quah (1996) originally documented for cross–country output and that Kounetas et al. (2021) found for US eco–efficiency. Substantively, it says that the cross–sectional distribution of renewable efficiency is not unimodal: there are two distinct "types" of US states, and the two types are reproduced over time rather than merging.

Third, what we do not see in Figure 3 is equally informative. There is no probability mass at the upper–left or lower–right corners: no laggard from 2010 has become a leader by 2022, and no leader has fallen to laggard status. There is no ridge running parallel to the x–axis, which would have indicated mean–reversion (every state being pulled to a common long–run value). And there is no ridge running flatter than 45°, which would have indicated σ–convergence (the spread of the distribution narrowing over time). The picture is therefore one of stable cross–sectional inequality: the same gap that separated leaders from laggards in 2010 still separates them in 2022, and the ergodic distribution discussed below confirms that this gap is unlikely to close in the long run absent a structural break.

The same pattern, in attenuated form, shows up for the two robustness indicators (Figure 7, discussed in Section 4). The bimodality is sharpest under the order–m estimator because that estimator preserves the high tail occupied by Vermont, DC, Idaho, and South Dakota; the conventional DEA estimator, bounded above by unity, compresses this tail and therefore displays a single dominant mode with a smaller secondary shoulder.

The long–run ergodic densities, obtained by iterating the discretised transition kernel until convergence, are plotted in Figure 4 for all three indicators. In each case the ergodic distribution is concentrated to the right of the 2010 distribution, indicating an upward drift in average efficiency. For the conventional DEA indicator the long–run distribution is sharply unimodal at high values; for the two order–m indicators it remains slightly bimodal, with a small secondary mode in the high tail occupied by Vermont and DC.

Figure 4: Initial (2010) and long–run ergodic densities of the three efficiency indicators

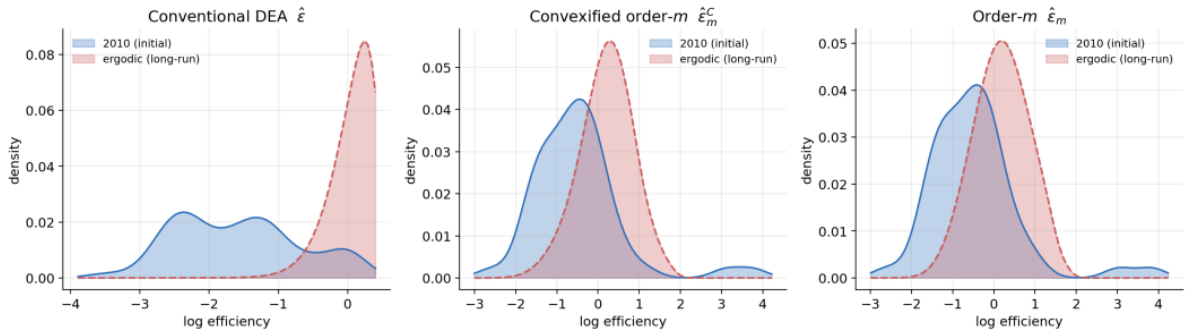

Note: The ergodic distribution lies to the right of the initial distribution in every case but retains a residual bimodality for the order–m indicators.

3.3 Phillips–Sul convergence clubs

Table 3 reports the log–t test for the full sample and the final club classification. The full–sample t–statistic is -13.36, well below the 5% critical value of -1.65, so the null of universal convergence is decisively rejected. The club algorithm finds five robust sub–groups. Six states are classified as "diverging": Vermont, DC, Idaho, and South Dakota at the top of the sort, where their efficiency scores are too far above their nearest peers to share a club, and Maryland and Florida at the bottom, where Florida's large fossil fleet and Maryland's idiosyncratic mix prevent assignment.

Table 3: Phillips–Sul club classification, order–m indicator

Group	$\hat{b}$	$\hat{\alpha}$	t–stat	Members
Club 1	-0.133	-0.067	-1.378	Maine, New Hampshire, Rhode Island, Alaska, North Dakota, Hawaii, Montana, Minnesota, Kansas, Oregon, California, Texas, Washington, Colorado, Iowa, Delaware, Wyoming, Nebraska, Nevada, Tennessee
Club 2	-0.937	-0.469	-1.616	Massachusetts, New Mexico, Oklahoma, Utah, Arizona, North Carolina, Alabama
Club 3	-0.459	-0.230	-1.598	New York, Arkansas, New Jersey, Illinois, West Virginia, Wisconsin, Connecticut, Mississippi, Michigan, Pennsylvania
Club 4	0.026	0.013	0.080	Missouri, Virginia, Kentucky, Louisiana, South Carolina, Georgia
Club 5	0.173	0.087	0.245	Indiana, Ohio
Diverging	—	—	—	Vermont, District of Columbia, Idaho, South Dakota, Maryland, Florida

Note: The full–sample log–t test gives $\hat{b} = -0.730$, $t = -13.36$, rejecting the null of full–sample convergence at the 5% level. The within–club speed–of–convergence parameter is $\hat{\alpha} = \hat{b}/2$ (Phillips and Sul, 2007).

Figure 5 plots the Phillips–Sul relative transition paths $h_{it} = \hat{\varepsilon}_{m,it}/\bar{\hat{\varepsilon}}_{m,t}$ for all 51 units, separated into two panels because the cross–sectional distribution of h spans more than two orders of magnitude. Mechanically, $h_{it} = 1$ corresponds to a state at the contemporaneous sample mean, $h_{it} < 1$ to a below–mean state, and $h_{it} > 1$ to an above–mean state. The horizontal grey line at $h = 1$ in both panels is therefore a visual reference for "average performance". Phillips–Sul convergence within a club requires the within–club h_{it} paths to collapse onto a common time path; divergence shows up as paths that fan out, drift apart, or sit on permanently different levels.

The five merged convergence clubs of Table 3 appear as five tightly stacked horizontal bands of thick coloured lines, in descending order of efficiency. Club 1 ($n = 20$) sits highest at a club mean of $h \approx 0.50$ (range 0.43–0.59 over the period), well below the unit reference line because Club 1 is the bulk of the "mainstream" distribution and the two extreme divergers DC and Vermont, shown in the right panel, lift the cross–sectional mean above all of these states. Club 2 ($n = 7$) sits at $h \approx 0.31$, Club 3 ($n = 10$) at $h \approx 0.18$, Club 4 ($n = 6$) at $h \approx 0.13$, and Club 5 (Indiana and Ohio) at $h \approx 0.09$. Three features of this stack are worth emphasising. First, the bands do not cross: the ordering established in 2010 is preserved throughout the sample, which is the visual analogue of the strong intra–distribution persistence we documented from the stochastic kernel (Figure 3). Second, the bands are flat: the year–on–year fluctuations within each club average out, so a state once classified into Club 3 stays in Club 3 thirteen years later. Third, the within–club spread (the thin faint lines around each thick line) tightens visibly for Clubs 1 and 3, the within–Club 1 coefficient of variation falls from 0.43 in 2010 to 0.32 in

Figure 5: Phillips–Sul relative transition paths $h_{it} = \hat{\varepsilon}_{m,it}/\bar{\hat{\varepsilon}}_{m,t}$ by club

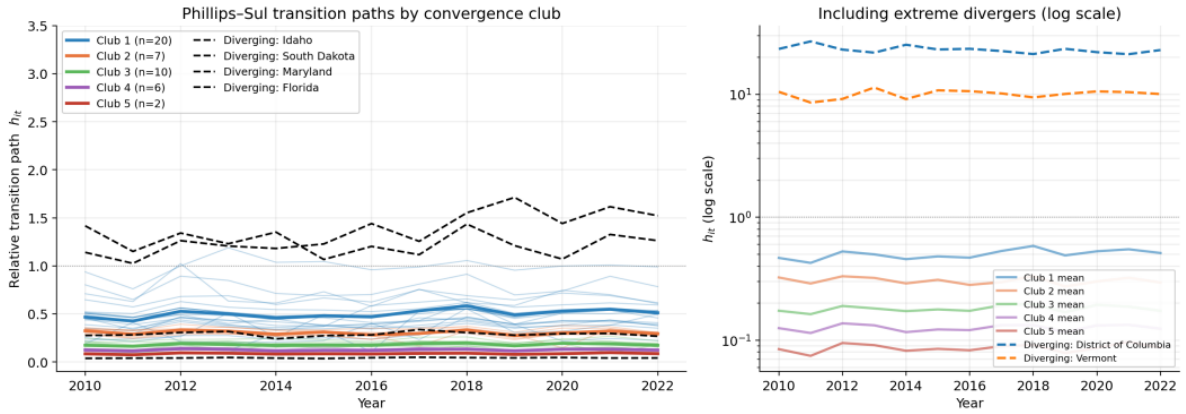

Note: **Left panel**: linear scale, with thick lines showing club means and thin lines individual states. **Right panel**: log scale, including the two extreme divergers DC and Vermont.

2022, which is the within–club σ–convergence that the formal log–t test detects.

The four black dashed lines in the left panel are four of the six diverging states. From top to bottom, they are Idaho (oscillating between $h \approx 1.03$ and 1.71), South Dakota (1.07–1.43), Maryland (0.24–0.34), and Florida (0.04–0.05). Idaho and South Dakota sit just above the unit reference line, well above Club 1's mean of 0.50, and their year–on–year oscillation is large enough that no club algorithm can absorb them into a stable neighbouring group. Maryland's path is squeezed between Clubs 2 and 3 without belonging to either. Florida sits below Club 5 (its 2022 $h \approx 0.04$ is less than half of Club 5's mean $h \approx 0.09$); the three orders of magnitude between Florida and DC, both classified as "diverging", illustrate that this category is heterogeneous, some divergers are super–performers too far above the herd, and som are laggards too far below.

The right panel re–plots the five club means together with the two extreme divergers, DC and Vermont, on a log scale spanning 10^{-1} to $10^{1.5}$. Two facts that the linear panel cannot show become visible here. First, DC's relative transition path sits between $h \approx 21$ and $h \approx 27$ throughout the sample, with a mean of 23.1. Its log–efficiency therefore exceeds the cross–section mean by roughly a factor of 20–25, not by "an order of magnitude" but by an order and a half. Vermont sits an order of magnitude lower than DC and an order of magnitude higher than the rest, oscillating between $h \approx 8.6$ and $h \approx 11.3$ with a mean of 10.1. Second, both DC and Vermont follow flat paths: their relative position is as stable across the 13 years as that of any club member, but their levels are too far above the bulk of the distribution to merge with any club under the log–t test (the test rejects merging when joint t–statistics fall below -1.65, and any attempt to put DC or VT into Club 1 produces t–statistics far below this threshold). On the log scale the five club means stack up between $h \approx 0.09$ (Club 5) and $h \approx 0.50$ (Club 1), preserving the same descending order as in the linear panel and confirming that the linear–scale ordering is not an artefact of how we scale the axis.

Overall, the two panels of Figure 5 make the formal output of the Phillips–Sul algorithm visually transparent. Convergence in the strong, full–sample sense ($h_{it} \to 1$ for every state) is plainly absent: no path approaches the unit line as t grows. Within–club convergence, different states converging to different but well–defined common levels, is the dominant pattern, with five tight bands that match the five clubs identified by the algorithm. The diverging units split into two visually distinct types: "too–good–to–share" divergers (DC, VT, ID, SD) whose levels are above all club means, and "too–bad–to–share" divergers (MD, FL) whose levels are below all club means. This typology is exactly what generates the bimodal stochastic kernel of Figure 3: the long–run distribution inherits the high tail of the super–efficient divergers and the low tail of Florida, neither of which is being absorbed into the mainstream by 2022.

The geographic distribution of the clubs in Figure 6 reveals interesting spatial regularities. Club 1 combines the Pacific hydro states (WA, OR, CA), the wind–belt (TX, KS, NE, ND, IA), and the small New England low–load states (NH, RI, ME), all of which share the property that renewable generation is large relative to their fossil and capacity bases. Club 2 collects sun–belt states with fast solar growth but still–large fossil fleets. Clubs 3 and 4 are predominantly Mid–Atlantic and Southeast states with mid–range performance, and Club 5 (IN, OH) is a tight pair of Rust–Belt fossil holdouts. The diverging Florida belongs to no group: its trajectory is neither close to Club 5 (where its 2022 value would put it) nor to any other.

Figure 6: Geographic distribution of the five Phillips–Sul convergence clubs (and the six diverging states, in black) for the order–m renewable–efficiency indicator, 2010–2022

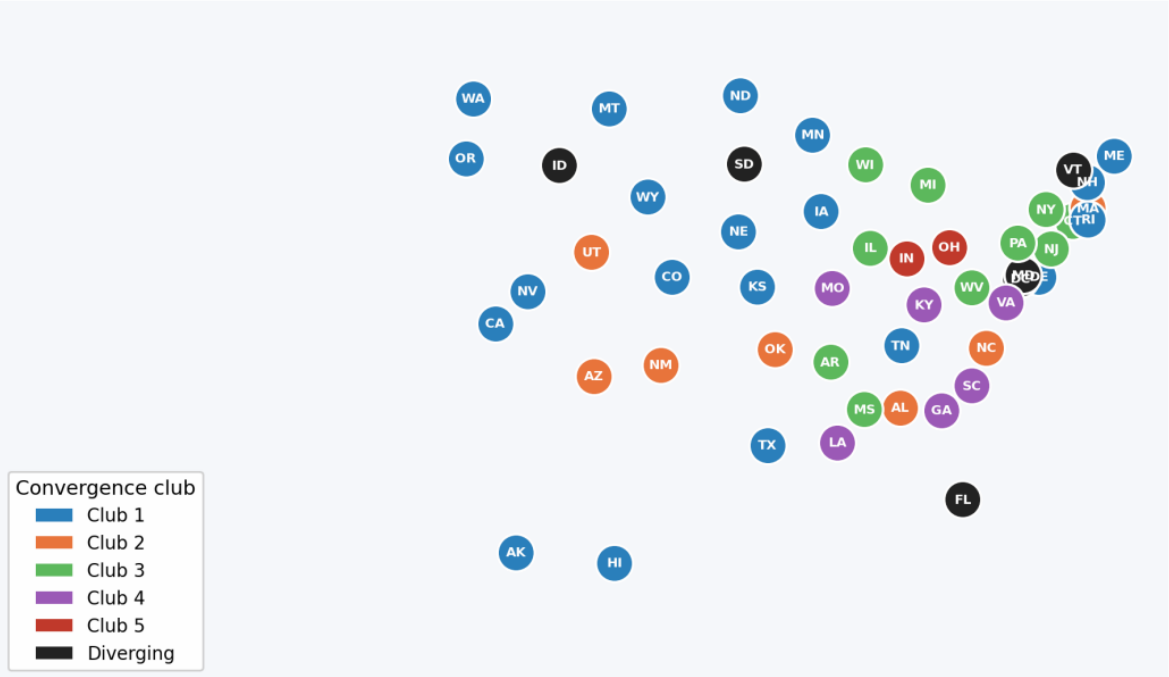

The within–club speed–of–convergence parameter is the $\hat{\alpha}$ of Phillips and Sul (2007), recovered as $\hat{\alpha} = \hat{b}/2$ from the slope coefficients reported in Table 3. Reading those coefficients directly, the within–club speeds are $\hat{\alpha} = -0.067$ for Club 1, -0.469 for Club 2, -0.230 for Club 3, $+0.013$ for Club 4, and $+0.087$ for Club 5. Two caveats apply before drawing any substantive comparison from these numbers. First, under the Phillips–Sul one–sided test the relevant statistic is the HAC t–ratio on $\hat{b}$, not $\hat{b}$ itself, and the test only distinguishes "passes the convergence null" from "rejects". The t–ratios in Table 3 (ranging from -1.62 for Club 2 to $+0.24$ for Club 5) all sit above the -1.65 threshold and so are all classified as "not rejected", but they do so by very different margins: Clubs 2 and 3 only barely clear the threshold, while Clubs 4 and 5 do so comfortably. Second, the positive $\hat{b}$ estimates for Clubs 4 and 5 are statistically indistinguishable from zero ($t = 0.08$ and $t = 0.24$ respectively); pointwise speed comparisons that rely on the magnitude of these estimates therefore have very wide confidence intervals and we do not over–interpret them.

What the numbers do support is a qualitative observation. The four laggard and mid–range clubs (2, 3, 4, 5) display tighter within–club co–movement than Club 1: the club with the most heterogeneous membership, which spans Pacific hydro states, the wind belt, and the small New England low–load states. Club 1 also has the smallest absolute $\hat{b}$ (-0.133, $\hat{\alpha} = -0.067$) and the lowest t–ratio in absolute value (-1.38): it is the largest club, the most heterogeneous club, and the closest to failing the convergence test. This is the opposite of the pattern reported by Koune-

tas et al. (2021) for US eco–efficiency in CO_2, SO_2 and NO_x, where the high–performing club showed the strongest within–club convergence; the difference is intuitive given that renewable–energy generation is the output of a sector whose technology frontier is itself moving rapidly (utility–scale solar costs fell by roughly 80% over 2010–2022), so the relatively heterogeneous Club 1 contains states at very different stages of frontier adoption.

4 Robustness

We re–estimate the convergence patterns under the conventional DEA–VRS estimator $\hat{\varepsilon}$ and the convexified order–m estimator $\hat{\varepsilon}^C_m$. Figure 7 shows the stochastic kernel of each, in the same 3D–surface and contour–plot format as Figure 3. Each contour panel includes the 45° identity line $y = x$ as a dashed white reference (and the 3D panels include the same line in dashed red on the $z = 0$ floor).

Figure 7: Robustness check: stochastic kernels for the conventional DEA indicator $\hat{\varepsilon}$ (left column) and the convexified order–m indicator $\hat{\varepsilon}^C_m$ (right column), 2010 vs. 2022

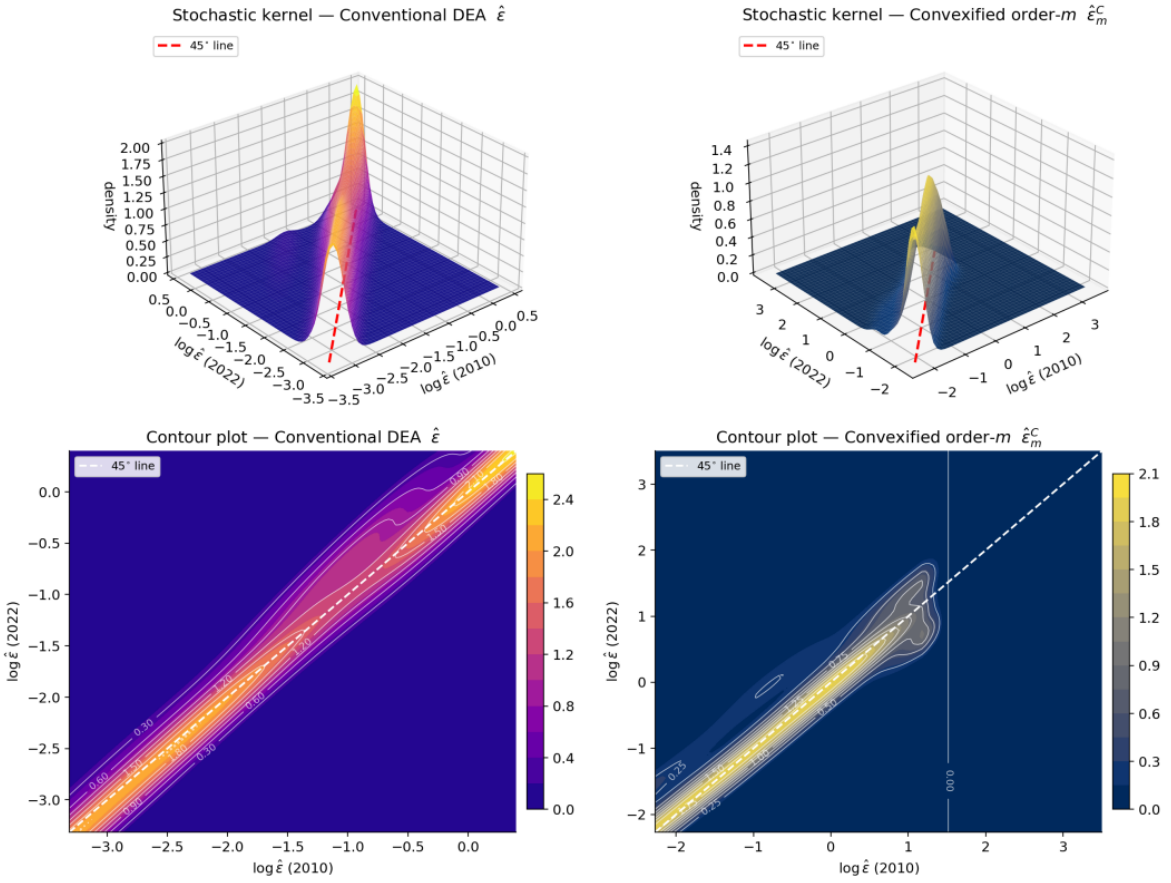

Note: **Top row:** 3D conditional–density surfaces; the dashed red line on the floor is the 45° identity line $y = x$. **Bottom row:** contour plots of the same surfaces; brighter colours denote higher density. The dashed white line is the 45° line. Concentration of probability mass on this line indicates intra–distribution persistence: states keeping their relative rank from 2010 to 2022.

The two robustness kernels tell two slightly different stories. The contour for the conventional DEA estimator (bottom–left) shows a clean ridge of probability mass aligned with the 45° line, with a single dominant peak near ($\log\hat{\varepsilon} \approx -1, -1$) and no visible secondary mode. The DEA ridge is somewhat thicker (perpendicular to the diagonal) than the order–m ridge in Figure 3, indicating a small amount of additional short–range mobility. The convexified order–m contour

(bottom–right) also displays a single ridge on the 45° line, peaking at ($\log \hat{\varepsilon} \approx 1, 1$), with a faint shoulder visible in the lower–left tail of the support but no clearly separated secondary mode. Both robustness kernels therefore confirm the intra–distribution–persistence finding from Section 3: states keep their relative rank under any of the three estimators. They do not, however, replicate the sharp bimodality of the order–m kernel. The reason is mechanical: the DEA estimator is bounded above by unity and so compresses the high tail where the four super–efficient states (VT, DC, ID, SD) sit; the convexified order–m allows the tail to extend but its convexification step pulls the secondary mode towards the main ridge. The bimodality finding is therefore a feature of the unrestricted order–m estimator, not a generic property of the underlying data.

Because the three estimators map the same underlying state–year inputs and outputs onto different numerical efficiency scores, with different boundary properties, different scales, and different sensitivity to outliers, it is natural for the Phillips–Sul algorithm, applied panel–by–panel, to identify a different number of clubs under each. This is exactly the design of the robustness check used by Kounetas et al. (2021), whose original study reports six, nine, and ten initial clubs across their three estimators. What we look for is consistency in the qualitative findings, whether the full–sample test still rejects, whether the laggard and champion groupings remain recognisably the same, not identity in the integer count of clubs.

In our application, the algorithm yields 5 final (merged) clubs under the order–m indicator, 8 under the conventional DEA estimator, and 5 under the convexified order–m estimator. Despite this difference in counts, the qualitative classification is highly consistent. The full–sample log–t test rejects convergence under all three indicators with t–statistics far below the -1.65 threshold: $t = -13.36$ for the order–m indicator, $t = -31.54$ for the conventional DEA, and $t = -17.34$ for the convexified order–m. The four super–efficient states (VT, DC, ID, SD) join the top club in the conventional DEA case (because the DEA estimator is bounded above by unity, they lose their outlier status), but the bottom of the distribution, IN, OH, KY, WV, MS plus Florida, is identified as a laggard cluster under all three estimators. The substantive conclusion that there is club convergence rather than US–wide convergence is therefore robust to the choice of frontier estimator, even though the exact number of clubs into which the mid–range states are partitioned is not.

5 Conclusions and Policy Implications

This paper has applied a probabilistic frontier analysis to the question of whether the US states are converging in renewable–energy efficiency. Three estimators (DEA–VRS, order–m, convexified order–m), a parametric convergence test (Phillips–Sul log–t), and a non–parametric convergence diagnostic (stochastic kernels) all point to the same conclusion. The full sample of states is not converging in renewable efficiency, but five well–defined sub–groups are. Two clear clusters emerge in the cross–section: a high–efficiency cluster that includes Pacific hydro, wind–belt, and small New England states, and a low–efficiency cluster of fossil–reliant Rust–Belt and Sunshine–State laggards. Four states (VT, DC, ID, SD) operate as super–efficient outliers, and Florida diverges from below.

The policy implications are direct. First, the absence of full–sample convergence is consistent with the patchwork structure of US clean–energy policy: states with stringent renewable–portfolio standards are pulling away from states without them, and the federal production tax credit and investment tax credit are not, by themselves, sufficient to close the gap. Second, the existence of stable convergence clubs implies that "one–size–fits–all" federal policy will not be efficient: the marginal abatement cost of further renewable deployment is very different between Club 1 (mostly already on the frontier) and Club 4–5 (with large fossil fleets and limited high–quality wind or solar resources). Targeted transmission investment, regional capacity markets, and federal financing geared toward the laggard clubs would close the gap faster than uniform

tax credits. Third, the persistence revealed by the stochastic kernel suggests that, absent a structural break, the cross–state gap between the high–performing club and the laggard club is unlikely to close on its own. The long–run ergodic distribution implied by the kernel is in fact unimodal, the dispersion narrows as the laggards slowly catch up to the mainstream of the distribution, but the catch–up is slow enough that the cross–sectional ranking established in 2010 remains effectively in place over a 13–year horizon.

There are potentially two limitations in this study. First, the analysis treats fossil–fuel generation, CO_2 from the power sector, and total nameplate capacity as inputs to be minimised given the renewable output. This is a deliberately sharp normative stance and other modelling choices (e.g. joint production with directional distance functions in the spirit of Chambers et al. (1996b)) would yield different efficiency rankings. Second, the panel ends in 2022 and therefore predates the bulk of the deployment expected under the Inflation Reduction Act, which makes the convergence patterns documented here a useful baseline for future evaluation.

References

Banker, R. D., Charnes, A., and Cooper, W. W. (1984). Some models for estimating technical and scale inefficiencies in data envelopment analysis. *Management Science*, 30(9):1078–1092.

Bird, L., Bolinger, M., Gagliano, T., Wiser, R., Brown, M., and Parsons, B. (2005). Policies and market factors driving wind power development in the united states. *Energy Policy*, 33(11):1397–1407.

Bistline, J., Blanford, G., Brown, M., Burtraw, D., Domeshek, M., Farbes, J., Fawcett, A., Hamilton, A., Jenkins, J., Jones, R., King, B., Kolus, H., Larsen, J., Levin, A., Mahajan, M., Marcy, C., Mayfield, E., McFarland, J., McJeon, H., Orvis, R., Patankar, N., Rennert, K., Roney, C., Roy, N., Schivley, G., Steinberg, D., Victor, N., Wenzel, S., Weyant, J., Wiser, R., Yuan, M., and Zhao, A. (2023). Emissions and energy impacts of the inflation reduction act. *Science*, 380(6652):1324–1327.

Carley, S. (2009). State renewable energy electricity policies: An empirical evaluation of effectiveness. *Energy Policy*, 37(8):3071–3081.

Cazals, C., Florens, J.-P., and Simar, L. (2002). Nonparametric frontier estimation: A robust approach. *Journal of Econometrics*, 106(1):1–25.

Chambers, R. G., Chung, Y., and Färe, R. (1996a). Benefit and distance functions. *Journal of Economic Theory*, 70(2):407–419.

Chambers, R. G., Chung, Y., and Färe, R. (1996b). Benefit and distance functions. *Journal of Economic Theory*, 70(2):407–419.

Charnes, A., Cooper, W. W., and Rhodes, E. (1978). Measuring the efficiency of decision making units. *European Journal of Operational Research*, 2(6):429–444.

Chung, Y. H., Färe, R., and Grosskopf, S. (1997). Productivity and undesirable outputs: A directional distance function approach. *Journal of Environmental Management*, 51(3):229–240.

Daraio, C. and Simar, L. (2007a). Conditional nonparametric frontier models for convex and nonconvex technologies: A unifying approach. *Journal of Productivity Analysis*, 28(1–2):13–32.

Daraio, C. and Simar, L. (2007b). Conditional nonparametric frontier models for convex and nonconvex technologies: A unifying approach. *Journal of Productivity Analysis*, 28(1–2):13–32.

Delmas, M. A. and Montes-Sancho, M. J. (2011). U.s. state policies for renewable energy: Context and effectiveness. *Energy Policy*, 39(5):2273–2288.

Färe, R., Grosskopf, S., Lovell, C. A. K., and Pasurka, C. A. (1989). Multilateral productivity comparisons when some outputs are undesirable: A nonparametric approach. *The Review of Economics and Statistics*, 71(1):90–98.

Farrell, M. J. (1957). The measurement of productive efficiency. *Journal of the Royal Statistical Society: Series A*, 120(3):253–281.

Fischer, C. and Newell, R. G. (2008). Environmental and technology policies for climate mitigation. *Journal of Environmental Economics and Management*, 55(2):142–162.

Hu, J.-L. and Wang, S.-C. (2006). Total-factor energy efficiency of regions in china. *Energy Policy*, 34(17):3206–3217.

Johnstone, N., Haščič, I., and Popp, D. (2010). Renewable energy policies and technological innovation: Evidence based on patent counts. *Environmental and Resource Economics*, 45(1):133–155.

Kounetas, K. E., Polemis, M. L., and Tzeremes, N. G. (2021). Measurement of eco–efficiency and convergence: Evidence from a non–parametric frontier analysis. *European Journal of Operational Research*, 291(1):365–378.

Kuosmanen, T. and Kortelainen, M. (2005). Measuring eco–efficiency of production with data envelopment analysis. *Journal of Industrial Ecology*, 9(4):59–72.

Li, J., Wu, H., Zhu, C., and Goh, M. (2024). Evaluating and analyzing renewable energy performance in oecd countries under uncertainty: A robust dea approach with common weights. *Applied Energy*, 375:124115.

Mardani, A., Zavadskas, E. K., Streimikiene, D., Jusoh, A., and Khoshnoudi, M. (2017). A comprehensive review of data envelopment analysis (dea) approach in energy efficiency. *Renewable and Sustainable Energy Reviews*, 70:1298–1322.

Marques, A. C., Fuinhas, J. A., and Manso, J. R. P. (2010). Motivations driving renewable energy in european countries: A panel data approach. *Energy Policy*, 38(11):6877–6885.

Menz, F. C. and Vachon, S. (2006). The effectiveness of different policy regimes for promoting wind power: Experiences from the states. *Energy Policy*, 34(14):1786–1796.

Phillips, P. C. B. and Sul, D. (2007). Transition modeling and econometric convergence tests. *Econometrica*, 75(6):1771–1855.

Phillips, P. C. B. and Sul, D. (2009). Economic transition and growth. *Journal of Applied Econometrics*, 24(7):1153–1185.

Quah, D. T. (1996). Twin peaks: Growth and convergence in models of distribution dynamics. *The Economic Journal*, 106(437):1045–1055.

Shrimali, G. and Kniefel, J. (2011). Are government policies effective in promoting deployment of renewable electricity resources? *Energy Policy*, 39(9):4726–4741.

Simar, L. and Wilson, P. W. (1998). Sensitivity analysis of efficiency scores: How to bootstrap in nonparametric frontier models. *Management Science*, 44(1):49–61.

Simar, L. and Wilson, P. W. (2007). Estimation and inference in two-stage, semi-parametric models of production processes. *Journal of Econometrics*, 136(1):31–64.

Sueyoshi, T. and Goto, M. (2012). Data envelopment analysis for environmental assessment: Comparison between public and private ownership in petroleum industry. *European Journal of Operational Research*, 216(3):668–678.

Sueyoshi, T. and Goto, M. (2017). A literature study for dea applied to energy and environment. *Energy Economics*, 62:104–124.

U.S. Bureau of Economic Analysis (2023). Regional economic accounts: GDP by state. https://www.bea.gov/data/gdp/gdp-state.

U.S. Energy Information Administration (2023a). State electricity profiles. https://www.eia.gov/electricity/state/.

U.S. Energy Information Administration (2023b). State energy–related carbon dioxide emissions. https://www.eia.gov/environment/emissions/state/.

Yin, H. and Powers, N. (2010). Do state renewable portfolio standards promote in-state renewable generation? *Energy Policy*, 38(2):1140–1149.

Zhou, P. and Ang, B. W. (2008). Linear programming models for measuring economy-wide energy efficiency performance. *Energy Policy*, 36(8):2911–2916.

Zhou, P., Ang, B. W., and Poh, K. L. (2008a). Measuring environmental performance under different environmental dea technologies. *Energy Economics*, 30(1):1–14.

Zhou, P., Ang, B. W., and Poh, K. L. (2008b). A survey of data envelopment analysis in energy and environmental studies. *European Journal of Operational Research*, 189(1):1–18.