When Green Investment Meets the Grid: Complementarity, Deployment, and Macroeconomic Effects

Tai Young-Taft, Ph.D.*

October 20, 2026

Abstract

Investment in renewable generating capacity can raise the value of transmission services, but technical complementarity does not by itself imply an aggregate output gain or a rationale for public subsidies. We distinguish three forms of evidence. First, a two-state constant-elasticity-of-substitution framework establishes conditions for technical and final-output complementarity, optimal resource allocation, and the sensitivity of allocation to inherited stocks and commissioning delays. Second, an analysis of the end-2025 U.S. interconnection-queue snapshot reports five-year commissioning bounds under explicitly stated missing-date and status-history assumptions. Third, state-level electricity, output, and federal renewable-program obligation data yield observational associations that are imprecise or sensitive to weighting, influential states, and sample selection. We also consider illustrative equal-resource-budget simulations. We neither observe a geographically validated grid-deliverability stock nor identify the macroeconomic return to matched generation and transmission investment. The contribution is a conditional allocation benchmark and a transparent account of what queue measurement and aggregate state data can establish about complementary green investment.

Keywords: renewable energy; transmission; interconnection; investment composition; missing data.
JEL: Q42, Q48, H54, O41.

1 Introduction

A renewable investment program can purchase generating equipment without securing timely access to the electricity network. This distinction matters for interpreting green industrial policy: expenditure, proposed megawatts, operating capacity, delivered electricity, and aggregate value added are different economic objects. Transmission can make generating capital more productive, but that mechanism alone does not establish either a positive fiscal multiplier or an optimal public subsidy. Cost, timing, incumbent assets, externalities, and the institutional division of benefits all matter.

This paper asks when investment in delivery capability raises the productive return to clean generation and what available national data can establish about that mechanism. The theoretical baseline contains only two endogenous states, clean generating capacity and effective grid capability. A constant-elasticity-of-substitution technology yields analytic conditions for technical and

*Research Scholar, Levy Economics Institute of Bard College. Correspondence: taiyoungtaft@gmail.com. *Preliminary and incomplete. Not for distribution or citation without the author's permission.*

final-output complementarity, dynamic user-cost conditions, and a closed-form one-period budget allocation. An optional knowledge stock is discussed but not estimated. This deliberate restriction keeps the main mechanism tractable without equating a stylized knowledge equation with a fully specified endogenous-innovation economy, which we leave to future work.

The empirical analysis combines EIA generation and capacity histories for 1990–2025, EIA emissions through 2024, Berkeley Lab's complete end-2025 queue snapshot, and BEA real state GDP for 1997–2025, annual population and personal income, NOAA-derived weather controls through 2024, and historical RPS/CES targets. ARPA-E award metadata are audited separately and are not treated as annual research spending. The datasets are not joined at project level: the EIA and BEA state-year panels and the queue request panel answer different measurement questions. We do not infer that every interconnection request corresponds to a unique EIA generator.

The argument is organized around three distinct objects. Section 3 derives a conditional investment-allocation mechanism; Sections 4–5 describe actual requests, generating capacity, electricity output, and state-level economic associations; and Section 6 illustrates how a fixed resource envelope performs under assumed stocks, costs, and delays. Queue-status timing bounds summarize part of the implementation problem, but they are not observations of the model's effective grid-capability stock. Nor are realized generation shares or grant obligations measures of new generation and transmission investment.

The incremental empirical exercise makes the role of missing or inconsistent actual-operation dates explicit when reporting fixed-horizon interconnection outcomes. That measurement analysis complements, rather than supersedes, the established literature on nationwide queue delays, withdrawals, interconnection costs, and network upgrades. The state macroeconomic specifications indicate empirical associations consistent with our theory, subject to qualifications concerning identification and robustness.

2 Related literature and the contribution boundary

The public-economics question is distinct from a claim that technologically complementary capital must receive subsidies. Gerarden et al. [2026] assess intervention rationales and evidence in renewable-energy industrial policy. Romer [1990] and Acemoglu et al. [2012] provide foundational accounts of nonrival knowledge and directed technical change. Our two-state benchmark holds knowledge fixed and asks how generating capital and effective grid capability interact under a common resource constraint. A public-investment rationale depends on a documented private–social wedge, such as network spillovers, regulation, coordination, or carbon damages, not on complementarity alone. Fullerton and Ta [2025] analyze renewable standards in general equilibrium, and Barrage and Nordhaus [2024] illustrate the additional damage accounting needed for welfare evaluation.

Transmission economics and infrastructure policy supply the central application. Davis et al. [2023] survey the institutional and economic obstacles to expanding the U.S. transmission grid. Hausman [2025] examines the effects of transmission and power flows on efficiency and producer profits; Botterud et al. [2024] model interregional transmission, emissions, and reliability in a physically richer setting. Doshi [2026] uses a Texas transmission-expansion setting to study longer-run wind siting and investment responses. In contrast, our stylized stock model does not resolve line-by-line power flows. Its comparative statics clarify assumptions and allocation choices rather than replacing project-level or engineering evidence.

Research on generator interconnection is especially close to our request-level analysis. Johnston

et al. [2023] investigate interconnection queues. Gorman et al. [2025] analyze more than 38,000 requests and document delays, withdrawals, and connection costs, so those patterns are not new findings here. Johnston et al. [2026] associate PJM network-upgrade charges and cancellations with prior transmission spending. Our distinct, narrower calculation uses the end-2025 queue snapshot to bound *operation within five years* when current status is known but actual operation dates are missing or inconsistent. Reporting both extrema avoids a point estimate conditional on complete event dates, under transparent assumptions about terminal statuses and historical coverage. It does not establish precedence over existing completion statistics: vintages, inclusion rules, project identities, follow-up horizons, and denominators differ and have not been harmonized to a common estimand.

The leap from project deployment to macroeconomic returns is another, separate research question. Ramey [2020] reviews why the timing, targeting, and financing of infrastructure investment complicate output-return estimation. Renewable electricity and growth studies—including Salari et al. [2021], Jeon [2022], and Atems and Hotaling [2018]—use different outcomes, geographic panels, and empirical strategies. Changes in a state's renewable-generation fraction combine installation, dispatch, weather, demand, and interregional trade, and therefore do not constitute exogenous green-investment shocks. Similarly, federal award obligations are not annual clean R&D expenditures and cannot stand in for matched generation and grid capital. Pollin et al. [2014] provide an investment-and-employment policy motivation, not a multiplier that can be imported as an estimate for this model.

The contribution is thus an explicit conditional allocation benchmark, fixed-horizon queue measurement robust to missing event timing, and a transparent sensitivity analysis of aggregate correlations. Existing transmission studies provide causal and engineering evidence about specific assets. Our contribution clarifies the inference between those studies, broad queue summaries, and aggregate macroeconomic claims.

3 A two-state investment framework

The mechanism is a supply-side allocation problem, not a claim about an aggregate fiscal multiplier. Nonenergy capital and labor are fixed conditioning inputs absorbed in $Z_t > 0$. Clean capacity $C_t > 0$ and effective grid capability $G_t > 0$ are the only endogenous states. Grid capability denotes useful delivery services, not transmission book value, queue capacity, or line miles. Baseline knowledge $A > 0$ is fixed. The CES technology function gives

$$E(C, G) = A[\alpha C^{\rho} + (1 - \alpha)G^{\rho}]^{1/\rho}, \quad \rho = 1 - \sigma^{-1} < 1, \quad 0 < \alpha < 1, \quad \sigma > 0, \tag{1}$$

with the continuous Cobb–Douglas limit at $\rho = 0$. This is a reduced form for useful services and is not an electrical power-flow model. In particular, it does not reproduce congestion at individual nodes, reliability constraints, or the temporal matching of variable generation and demand. Conditional final output is

$$Y_t = Z_t(D_t + E(C_t, G_t))^{\eta}, \qquad 0 < \eta < 1, \tag{2}$$

where $D_t \geq 0$ is an exogenous outside electricity-service supply. The perfect substitution of these *delivered services* is a simplifying assumption, not a statement that intermittent nameplate capacity and firm power are interchangeable. This macroeconomic bridge holds nonenergy factors fixed and excludes demand stimulus, financing frictions, and endogenous fossil dispatch.

3.1 Technical complementarity and its macroeconomic qualification

Let $S = \alpha C^{\rho} + (1-\alpha)G^{\rho}$. Direct differentiation gives

$$E_C = A\alpha C^{\rho-1} S^{1/\rho-1}, \tag{3}$$

$$E_{CG} = A\alpha(1-\alpha)(1-\rho)C^{\rho-1}G^{\rho-1}S^{1/\rho-2} > 0, \tag{4}$$

$$\frac{EE_{CG}}{E_C E_G} = 1-\rho = \frac{1}{\sigma}. \tag{5}$$

These identities also hold by continuity at $\rho = 0$. The CES aggregator is increasing, concave, and homogeneous of degree one. Thus $E_{CC} < 0$, $E_{GG} < 0$, and the Hessian is negative semidefinite, not negative definite.

Proposition 1 (bottleneck and output complementarity). An increase in effective grid capability raises the marginal clean-electricity service product of capacity. If $\sigma < 1$, this marginal product tends to zero as $G \downarrow 0$ at fixed $C > 0$. Nevertheless, the output cross-partial is positive if and only if

$$\frac{1}{\sigma} > (1-\eta)\frac{E}{D+E}. \tag{6}$$

In particular, $\sigma \leq 1$ is sufficient. When $D = 0$, the necessary and sufficient condition is $\eta > \rho$. Technical complementarity alone therefore does not establish output complementarity for all parameter values.

Proof. The first claim follows from the displayed cross-partial. For $\rho < 0$, S diverges as $G \downarrow 0$ while $1/\rho - 1 < 0$, implying $E_C \to 0$. Applying the chain rule to (2) gives

$$Y_{CG} = Y_E \frac{E_C E_G}{E}\left[\frac{1}{\sigma} - (1-\eta)\frac{E}{D+E}\right], \tag{7}$$

which proves the remaining statements. □

The corresponding local output elasticity is

$$\frac{\partial \log Y}{\partial \log C} = \eta \frac{E}{D+E}\frac{\alpha C^{\rho}}{S}. \tag{8}$$

The factor $E/(D+E)$ matters: a large percentage change in a small clean-service sector need not imply a large aggregate output change. Neither this elasticity nor σ is identified by a regression of state GDP on renewable nameplate capacity alone.

3.2 Dynamics, decentralization, and the planner

Time is continuous and investment is installed immediately in the baseline:

$$\dot{C} = I_C - \delta_C C, \qquad \dot{G} = I_G - \delta_G G, \quad I_C, I_G \geq 0. \tag{9}$$

A representative owner of the consolidated production system maximizes discounted output net of investment resources,

$$\max_{I_C, I_G} \int_0^{\infty} e^{-rt}\left[Y(C,G) - \sum_{j\in\{C,G\}}\left(p_j I_j + \frac{\kappa_j}{2} I_j^2\right)\right] dt, \qquad r, p_j, \kappa_j > 0. \tag{10}$$

Quadratic installation costs prevent instantaneous adjustment. The specification is a quasilinear surplus criterion; it is not a calibrated representative-household welfare function. An exogenous numeraire sector can absorb transfers and resource purchases. Assume stationary Z, D, A for the equilibrium statements below. With current-value shadow prices q_C, q_G, necessary conditions are

$$I_j = \max\{0, (q_j - p_j)/\kappa_j\}, \tag{11}$$

$$\dot{q}_j = (r + \delta_j)q_j - Y_j, \qquad j = C, G, \tag{12}$$

$$\lim_{t\to\infty} e^{-rt} q_j(t) X_j(t) = 0, \qquad (X_C, X_G) = (C, G). \tag{13}$$

Under concavity, a feasible path satisfying these conditions and the transversality conditions is optimal. At an interior stationary solution,

$$I_j^* = \delta_j X_j^*, \qquad Y_j^* = (r + \delta_j)(p_j + \kappa_j \delta_j X_j^*). \tag{14}$$

These conditions equate marginal output with the user cost of installed capital. A competitive final-output producer paying marginal-product rental rates Y_C and Y_G to the respective capital owners decentralizes this benchmark. Because it contains no externality or missing market, complementarity by itself creates no wedge between private and social investment and supplies no general argument for subsidies. A public-investment rationale requires an additional, documented wedge: for example, transmission benefits not captured by the investor, a regulated return below the social return, or coordination failures. A subsidy is a transfer; treating a dollar of subsidy as a dollar of installed resources requires an additionality and incidence model.

If installation costs vanish and an interior steady state exists, define $u_j = (r + \delta_j)p_j$. The ratio condition becomes

$$\frac{G^*}{C^*} = \left[\frac{1-\alpha}{\alpha}\frac{u_C}{u_G}\right]^{\sigma}. \tag{15}$$

A higher grid user cost lowers the optimal grid-capacity ratio; a higher generation user cost raises it. Convex adjustment costs alter the ratio through endogenous marginal installation costs in (14). These equilibrium comparisons do not identify either the speed of adjustment or an observed policy effect.

3.3 Exact equal-budget allocation and corner solutions

Consider a one-time, immediately productive resource budget $B > 0$, with fixed unit costs and no adjustment costs for this separate allocation exercise. Put a share $x \in [0, 1]$ into generation:

$$C_1 = C_0 + xB/p_C, \qquad G_1 = G_0 + (1 - x)B/p_G. \tag{16}$$

Solar-only, grid-only, and equal-share coordination are $x = 1, 0, 1/2$. Maximizing either E_1 or Y_1 yields the same allocation because $Y_E > 0$. Let

$$R = \left[\frac{1-\alpha}{\alpha}\frac{p_C}{p_G}\right]^{\sigma}. \tag{17}$$

Proposition 2 (optimal mix). The unique maximizing share is

$$x^* = \left[\frac{G_0 + B/p_G - RC_0}{B/p_G + RB/p_C}\right]_0^1, \tag{18}$$

where $[z]_0^1 = \min(1, \max(0, z))$. An interior solution equates marginal services per dollar and satisfies $G_1/C_1 = R$. If $G_0 + B/p_G \leq RC_0$, all funds optimally go to grid capability; if $G_0 \geq R(C_0 + B/p_C)$, all funds go to generation. An interior optimized mix weakly dominates both corners and strictly dominates them when the optimum is strictly interior; a preset fifty-fifty mix has no such guarantee.

Proof. Strict concavity along the downward-sloping budget line ensures a unique solution. The first-order condition is $E_C/p_C = E_G/p_G$. Since $E_C/E_G = [\alpha/(1-\alpha)](G/C)^{1/\sigma}$, this implies $G/C = R$. Substitution in the budget gives (18); the Kuhn–Tucker inequalities give the corners. □

For finite-horizon simulations the appropriate comparison instead uses discounted resource costs and the entire service path. Under a constant marginal service value v_j after commissioning, a unit installed after an externally imposed delay d_j and depreciating at δ_j has marginal present value

$$\frac{v_j}{p_j}\int_{d_j}^{T} e^{-rt}e^{-\delta_j(t-d_j)}dt = \frac{v_j}{p_j}\frac{e^{-rd_j}[1-e^{-(r+\delta_j)(T-d_j)}]}{r+\delta_j}, \quad d_j < T. \tag{19}$$

This is a local timing diagnostic, not a global solution because marginal values change with both stocks. Delays can overturn a static policy ranking. Endogenous construction pipelines would add states; to respect the two-state baseline, any timing exercises must use externally specified installation schedules and identify them as such. Queue duration is not identical to transmission construction time.

3.4 Innovation and interpretation

A third-state extension can restore the prior draft's law $\dot{A} = \phi R^{\zeta} A^{\lambda} - \delta_A A$, with $0 < \zeta < 1$ and $0 \leq \lambda < 1$. Romer's nonrival knowledge and the directed-technical-change framework motivate considering innovation spillovers, but this equation is neither a derivation nor an empirical implementation of those full models. Research spending must compete for the same resource budget, and knowledge externalities require explicit private and social returns. We exclude this extension from estimated policy claims. Likewise, Nordhaus-style climate welfare would require emissions and damage accounting absent from this two-state model. Pollin-style investment and employment accounting and the infrastructure macroeconomics literature motivate the policy question without providing estimates of the transmission interaction here.

4 Original data and measurement

4.1 EIA electricity panel

We use the Total Electric Power Industry aggregate and exclude producer subcategories to prevent double counting. The analytical technology categories are Wind and Solar Thermal and Photovoltaic; solar is therefore not a PV-only variable. Generation is measured in MWh and nameplate capacity in MW. State codes are normalized to the 50 states and District of Columbia.

Absent technology records remain missing; they are not assigned zero. This is an unbalanced panel, and the 51 jurisdictions represented in the combined file are not all present in each estimation sample.

The annual ratio $Q_{it}/(h_t C_{it})$ uses year-end nameplate capacity, with $h_t = 8{,}784$ in leap years and 8,760 otherwise. It is not a conventional engineering capacity factor for a fixed fleet. Commissioning

late in the year, retirements, resource quality, curtailment, maintenance, and reporting scope all affect it. Two merged observations exceed one; they are retained and flagged rather than automatically winsorized because annual generation may reflect capacity retired before the year-end capacity measurement, permitting the ratio to exceed one even without a reporting error. A robustness specification uses $(C_{it} + C_{i,t-1})/2$ only when both adjacent observations exist. This is a constructed exposure proxy.

EIA estimated emissions are in metric tons. Dividing state power-sector CO_2 by generation gives production-based carbon intensity, not consumption-based emissions. Interstate imports and exports preclude interpreting a state's generation mix as its residents' electricity consumption.

4.2 Queue requests, status, and event dates

The Queued Up 2026 dataset is the source for the project analysis [Rand et al., 2026]. Its 38,201 request rows comprise 24,221 withdrawn, 8,513 active, 4,789 operational, 668 suspended, and 10 unknown-status records. The sample covers transmission interconnection and does not represent all distribution-connected or behind-the-meter generation. Historical reporting completeness varies across providers. An application can be speculative, revised, or resubmitted; counts are requests rather than a census of unique physical plants.

There are no exact duplicated rows, but 270 rows share an entity/request-ID pair with another row. We retain those records with row identifiers and rerun the comparison excluding all ambiguous pairs. Missing component capacities in hybrid requests remain missing. We never add storage MW or storage energy to generator MW to estimate operating generation. The primary regressions compare standalone solar and wind, and descriptive bounds report hybrid and storage categories separately. These are counts of the full microdata sheet, not a claimed reconstruction of every filtered headline total in the source report; for example, our 8,513 active rows exceed the publisher landing page's approximately 8,200 active projects. The exact difference in filtering or vintage has not been reconciled.

Only 3,289 of 4,789 operational requests have an actual online date. Requiring valid request dates, nonnegative durations, and operation by December 31, 2025 leaves 3,098 durations. Forty-eight operational rows have negative recorded durations; these are excluded from duration summaries, not repaired. The current status regressions do not use event dates, so those rows are not discarded from that separate analysis.

4.3 BEA, demographics and unavailable investment inputs

BEA SAGDP1 LineCode 1 provides all-industry real GDP in millions of chained 2017 dollars. The 50 states plus DC contribute 1,479 observations over 1997–2025, with no missing GDP in that extraction. We do not add chained-dollar industry components. Joining the EIA renewable shares leaves 931 missing-share state-years and 504 complete observations for the contemporaneous growth specification.

A geographically validated state-year transmission-investment or physical grid-deliverability series is unavailable for the present estimations. FERC Form 1 archives contain 571 annual 2014–2016 utility filings and 465 observed transmission-plant-addition facts, but these are utility accounting capitalizations rather than state-allocated cash investment or delivered grid services. An asset or project-level geographic crosswalk and reconciliation of company accounts are prerequisites for state aggregation. The main results also lack a validated annual clean-specific research-expenditure

and actual generation-investment panel; ARPA-E project award amounts are multiyear metadata. Annual demographic and income data are available for the observational controls in Section 5.6. Crucially, no estimated equation observes the model's grid state G.

5 Empirical designs and results

5.1 Fixed-horizon completion with missing event dates

For requests entering in 2000–2020 with an observed request date and a known snapshot status, we require that the fifth calendar anniversary precede the end-2025 snapshot. This yields 22,344 requests. Define O_i as recorded operational status, V_i as a valid actual operation date, and T_i as elapsed time from request to operation. For technology group k, the lower and upper shares are

$$L_k = n_k^{-1} \sum_{i \in k} O_i V_i \mathbf{1}\{T_i \leq 5\}, \tag{20}$$

$$U_k = L_k + n_k^{-1} \sum_{i \in k} O_i (1 - V_i). \tag{21}$$

Thus unknown timing among recorded operational requests can either precede or follow the horizon. A date is valid only if it is not before the request or after the snapshot. The implementation uses calendar anniversaries rather than rounded elapsed years. These are worst-case bounds on unobserved event timing in the spirit of partial identification [Manski, 2003], conditional on recorded outcomes. Analogous bounds use withdrawal status and withdrawal dates. Withdrawal is a competing terminal outcome, not independent censoring.

These bounds condition on the supplied current statuses being mutually exclusive absorbing paths for each record. They do not account for missing historical records, incorrect status classifications, or plants that operated and were later recategorized. Active and suspended records receive zero realized operation by five years under that maintained interpretation, even when requests are older than five years; this is a statement about snapshot histories, not a claim that these projects will never operate. The intervals are missing-data bounds; completion and withdrawal upper bounds are separate marginal bounds rather than numbers that should be added.

Table 1: Five-year queue outcomes: bounds in percent

Technology	Requests	Operation lower	Operation upper	Withdrawal lower	Withdrawal upper
Solar	8614	5.526	9.310	46.947	69.143
Wind	5431	10.127	15.632	37.231	67.244
Solar+Battery	1260	1.270	4.444	35.079	58.333
Battery	1433	4.117	7.188	56.106	67.760
Gas	2677	18.491	26.261	37.579	61.412

Request cohorts 2000–2020; exact-date five-year follow-up. Bounds reflect missing or invalid event timing conditional on observed status. Technology groups have different cohort and geographic composition.

For standalone solar, 476 of 8,614 requests are known to have operated within five years; another 326 operational requests have uncertain event timing. The resulting interval is 5.53–9.31 percent. Wind's corresponding figures are 550 known completions and 299 timing-unknown operational requests among 5,431, giving 10.13–15.63 percent. Gas has higher observed bounds, while storage and solar-plus-battery have lower ones. Nonoverlap of these raw intervals is not a causal technology effect: cohort composition, provider coverage, location, financing, and project selection differ. The

principal measurement result is the width of the intervals and the dependence of a point estimate on assumptions about missing timing. Unlike Gorman et al. [2025], we do not estimate comparative connection costs or attribute withdrawals to a mechanism; our horizon-specific bounds also use a different data vintage and denominator, precluding direct comparison of headline rates.

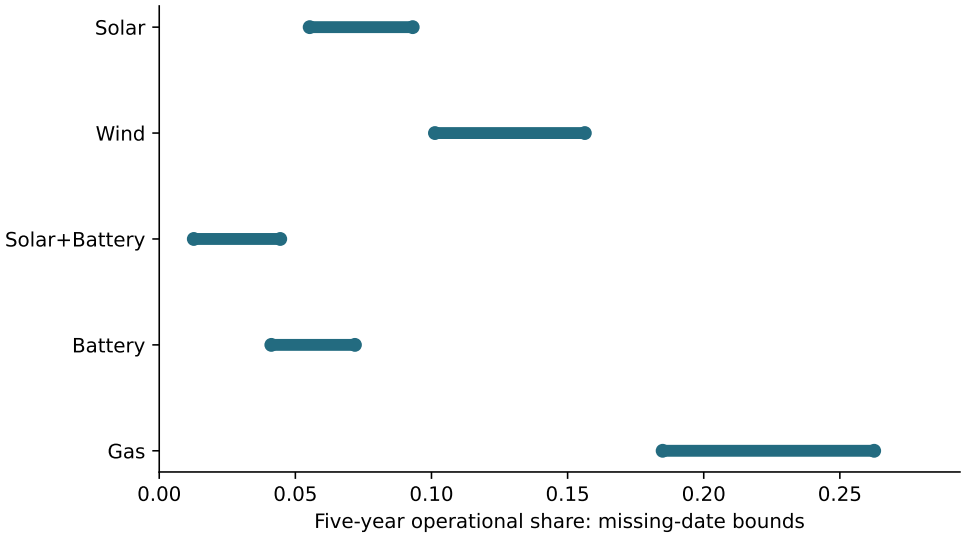


Figure 1: Five-year operation bounds by requested technology. Horizontal spans reflect unknown event timing.

5.2 Mature-cohort operational associations

The 2000–2018 cohorts have at least seven calendar years of potential follow-up, but older cohorts have longer exposure. Among 17,446 requests of all technologies in these cohorts, the recorded operation rates are 15.59 percent for standalone solar, 20.21 percent for wind, and 30.45 percent for gas. Unresolved and withdrawn requests remain in the denominator. These are snapshot status shares, not common-duration completion probabilities.

For solar and wind requests with positive primary MW and available controls, we estimate

$$\Pr(O_i = 1 \mid X_i) = \Lambda\left(\beta S_i + \gamma \log MW_i + \lambda_{q(i)} + \mu_{r(i)}\right), \tag{22}$$

where S_i indicates standalone solar, $q(i)$ is entry year, and $r(i)$ is region. The main sample contains 11,333 requests and 48 transmission providers. Provider-clustered sandwich covariance allows within-provider dependence. The model describes operational status conditional on application and recorded covariates; it is not a treatment-effects or structural entry model.

The main odds ratio is 0.711, with a 95 percent interval of 0.487–1.038 ($p = 0.078$). Standardizing fitted probabilities over the same covariate sample gives a solar-minus-wind difference of -4.80 percentage points; this is a model-based point contrast, without an independently calculated interval. Older cohorts and exclusion of ambiguous request IDs give odds ratios of 0.649 and 0.719, both with intervals including one. Substituting provider for region fixed effects gives 0.660 with an upper interval endpoint of 0.997 ($p = 0.048$).

Differences in siting, permitting, private profitability, reporting, and selection into the queue remain unobserved. Providers can share markets and shocks, so provider clustering does not remove

Table 2: Solar relative to wind: operational odds ratios

Specification	N	Clusters	OR	95% lower	95% upper
Region FE, 2000–2018	11333	48	0.711	0.487	1.038
Region FE, 2000–2015	7036	39	0.649	0.399	1.056
Unambiguous IDs	11237	48	0.719	0.484	1.067
Provider FE, 2000–2018	11333	48	0.660	0.437	0.997

Binomial logit, entry-year effects and log primary-component MW. Region effects except in the provider-effect alternative. Provider-clustered Wald intervals; no multiple-testing adjustment.

all spatial dependence; uneven cluster sizes also complicate asymptotic inference. No variable in this regression measures transmission investment, physical congestion, or effective grid capability. The estimates cannot identify Proposition 1's elasticity.

Among operational requests with credible dates, the median request-to-operation duration is 3.74 years for the 2020 commissioning cohort, 5.11 years for 2024, and 5.44 years for 2025 (188, 257, and 163 observations, respectively). These completed-project cohorts are selected on realization and date availability. Their trend cannot be interpreted as the waiting time of a representative applicant or attributed solely to transmission constraints.

5.3 Realized generation, capacity, and emissions

For each technology, we estimate the descriptive association

$$\log Q_{it} = a_i + b_t + \beta \log C_{it} + u_{it}, \tag{23}$$

using positive generation and at least 10 MW of nameplate capacity. State and year effects absorb invariant state attributes and common year shocks. State-clustered standard errors use a finite-sample correction and cluster-based t inference. Alternative samples start in 2010, require 100 MW, or substitute average adjacent-year capacity. The coefficient is not a return to investment: capacity and generation are jointly determined and share reporting structure.

Table 3: EIA descriptive regressions and denominator sensitivity

Specification	Slope	Cluster SE	N	States
Wind: baseline	0.968	0.051	833	41
Wind: 2010 onward	1.036	0.059	596	41
Wind: >=100 MW	1.048	0.037	689	37
Wind: midpoint MW	0.999	0.035	807	40
Solar: baseline	0.996	0.045	580	48
Solar: 2010 onward	0.995	0.045	552	48
Solar: >=100 MW	1.079	0.063	397	47
Solar: midpoint MW	1.099	0.037	555	48
Emissions: all years	-0.410	0.131	470	40
Emissions: 2010 onward	-0.408	0.129	439	40

Generation models are log–log; emissions models regress changes in metric tons CO_2/MWh on changes in the wind-plus-solar generation fraction. All include state and year effects. The two outcome types have different units.

Baseline slopes are 0.968 for wind and 0.996 for solar. The solar slope rises to 1.099 with midpoint capacity, whereas wind's is 0.999. This denominator sensitivity is consistent with capacity exposure timing affecting annual ratios, but it does not separately estimate that mechanism. The estimates do

not establish absence of local curtailment: annual aggregation can mask binding hourly constraints, and resource quality and geographic expansion can offset those constraints.

National source totals show wind increasing from 39.516 GW and 94.652 TWh in 2010 to 159.732 GW and 461.577 TWh in 2025. The corresponding solar category rises from 0.987 GW and 1.212 TWh to 154.301 GW and 300.303 TWh. These are source-defined electric-industry categories, not an asserted census of all distributed solar. Their combined share of the source's total generation rises from 2.32 to 17.22 percent.

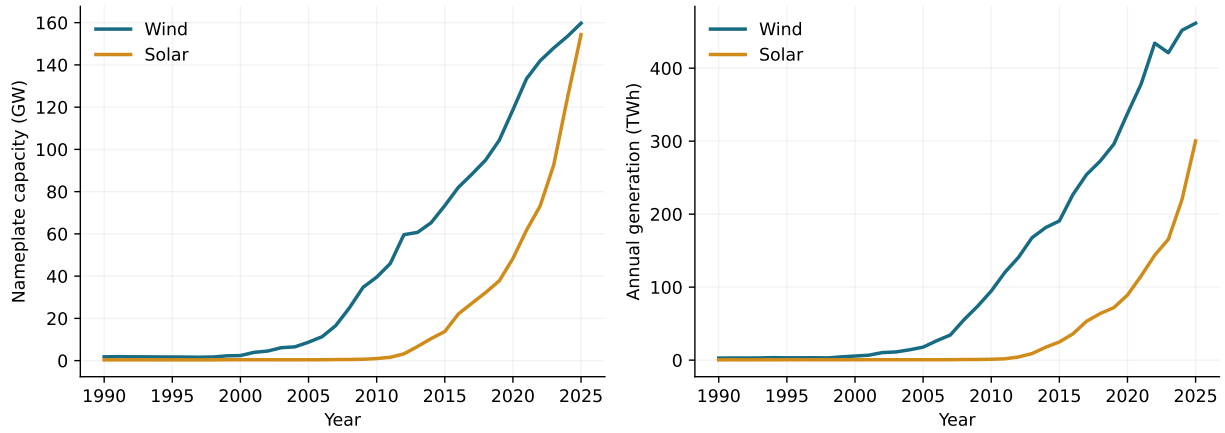


Figure 2: Observed national EIA capacity and generation. Nameplate capacity is a stock; annual MWh is a flow. The historical solar category combines thermal and photovoltaic generation.

A separate first-difference regression relates changes in production-based carbon intensity to changes in the renewable share, with state and year effects. Its slope is -0.410 metric tons/MWh per unit fraction (cluster SE 0.131; 470 observations). A ten-percentage-point share change is associated with approximately 0.041 fewer metric tons/MWh. This is partly an accounting relationship between composition and intensity, with fossil mix, demand, imports, and other contemporaneous changes unresolved. It is not marginal avoided emissions attributable to the investment program.

5.4 Macroeconomic associations and their limits

We estimate

$$\Delta \log GDP_{it} = a_i + b_t + \theta \Delta s_{it} + \epsilon_{it}, \tag{24}$$

where s is the wind-plus-solar share of state generation. Changes are used rather than a regression of two trending levels. Complete cases require actual observations of both technologies in both adjacent years; absent source records are not zero-filled. The main specification has 504 observations across 42 jurisdictions. Selection toward places and periods with reported renewable activity is substantial.

The contemporaneous coefficient is 0.0025 (SE 0.0517; 95 percent interval $[-0.1019, 0.1070]$). For a ten-percentage-point renewable-share change, the associated growth change is about 0.025 percentage points, with an interval approximately $[-1.019, 1.070]$ percentage points. Lagged, post-2010, geographic-exclusion, and pandemic-exclusion specifications remain imprecise. The

Table 4: Real GDP growth and renewable-share changes: descriptive associations

Specification	Slope	SE	95% lower	95% upper	N
contemporaneous	0.003	0.052	-0.102	0.107	504
lagged share change	0.051	0.053	-0.057	0.158	462
future-change diagnostic	0.026	0.037	-0.048	0.101	502
2010 onward	0.002	0.051	-0.102	0.106	481
exclude AK HI DC	0.009	0.048	-0.088	0.106	483
exclude 2020-21	-0.036	0.055	-0.148	0.075	429

BEA all-industry real GDP; EIA generation shares. State/year effects and state-clustered intervals. Share changes are fractions, not percentage points. The future-change specification is a timing diagnostic, not an identification test.

future-share-change coefficient also includes zero. This does not establish exogeneity or "pass" a pretrend test.

5.5 Controls, weighting, and nonlinear sensitivity

To assess sensitivity to the state-year composition of the panel, we also estimate distributed-lag associations for annual changes in the wind-plus-solar generation share:

$$\Delta \log GDP_{it} = \alpha_i + \tau_t + \sum_{\ell=0}^{3} \beta_\ell \Delta s_{i,t-\ell} + \gamma' Z_{i,t-1} + u_{it}, \qquad \Theta = \sum_{\ell=0}^{3} \beta_\ell. \tag{25}$$

The common-window changes begin in 2011, and a complete four-term lag block yields 2014–2025 outcomes. The broad sample has 350 state-years from 40 states; 13 states have 156 continuously reported observations. The latter are selected for reporting completeness, not random assignment. We compare equal-observation weights with weighted least squares using *prior-year real state GDP*. These weights change the descriptive estimand toward larger state economies; they are neither household-income nor population weights.

The full $Z_{i,t-1}$ set includes nominal BEA mining, manufacturing, and utilities shares of state value added, the log of the EIA average electricity price, lagged real GDP growth, and log real GDP. Nominal sector shares avoid the nonadditivity of chained-dollar industry values. These controls are predetermined in calendar time but may respond to previous electricity investment and structural change, so the estimates cannot be interpreted as causally "adjusted." The samples in Table 8 are held fixed within each defined subset. Standard errors are clustered by state; restricted-null Rademacher wild-cluster bootstrap tests, following the clustered-inference approach of Cameron et al. [2008], are reported alongside conventional clustered tests. For the small balanced panels, all 2^{13} (or 2^{12}) sign configurations are enumerated.

Across the broad and selected panels, the estimates are too specification-sensitive for a stable macroeconomic interpretation. The broad unweighted lag sum is essentially zero in both the baseline ($\widehat{\Theta} = 0.004$) and full-control (0.072) specifications. The corresponding GDP-weighted coefficients are positive (0.694 and 0.667); the full-control wild-cluster p-value is 0.043 before any adjustment for the exploratory search across models, but rises to 0.102 after California is removed. The 13-state full-control GDP-weighted model similarly has an exact sign-enumeration p-value of 0.049, which rises to 0.103 without New Mexico. Conventional clustered p-values and wild-cluster results do not

always agree, notably in the small balanced panel. Per-year normalization of GDP weights changes little, while leave-one-state-out exercises indicate substantial leverage.

A quadratic change-in-generation-share term has different signs or significance under alternative samples and weighting; an interaction between share change and prior renewable share is not robust to GDP weighting or exclusion of New Mexico. Neither nonlinear exercise measures an interaction between *investment in generation* and *investment in grid deliverability*, and hence neither tests increasing returns to complementary investment. The controlled results should be read as descriptive specification sensitivity, not multipliers, welfare benefits, or effects of a policy shock.

The sample underlying the contemporaneous GDP regression (504 state-years, 42 jurisdictions) differs from the four-lag panel (350 observations, 40 states), the demographic-control models (430 current-change and 399 lagged-change observations, 40 states), and the federal-funding sample (702 observations, all 51 jurisdictions). Cross-column differences therefore reflect changes in sample and estimand as well as specification.

5.6 Population, personal income, weather and policy sensitivity

BEA SAINC1 supplies annual total personal income (millions of current dollars), population (persons), and per-capita personal income (current dollars). The September 30, 2026 update provides all 1,479 state-years in 1997–2025; 1996 is retained to construct initial demographic lags. SAINC30 lines 10, 100 and 110 match SAINC1 lines 1, 2 and 3 exactly. Reconstructing per-capita income from rounded total income and population differs from the published value by at most \$0.55. State FIPS and USPS joins are one-to-one and preserve every original panel row and missing cell.

Census Vintage 2019 estimates for 2010–2019 and Vintage 2025 estimates for 2020–2025 provide a separately labeled population comparison, not a replacement for BEA's supplied population series. The largest difference is 4.07 percent in New Jersey in 2019. Vintage differences are consequently material. EIA SEDS 2024F supplies NOAA-derived population-weighted heating and cooling degree days, base 65 degrees Fahrenheit, for all 51 jurisdictions in 1997–2024. Weather remains missing in 2025. No observation is interpolated or assigned zero.

For size variable z, the year-t weight is $z_{i,t-1} / \sum_{j=1}^{51} z_{j,t-1}$, defined only when all 51 source values are observed. Lags are constructed on the full annual grid before regression filtering; WLS weights are rescaled to mean one on each estimation sample. Population, nominal GDP, total personal income and per-capita personal income implement different descriptive emphases. They are not sampling weights or a welfare criterion. Real chained-dollar GDP weights are retained only as a sensitivity, given nonadditivity. No product of population and income is used.

We separately regress annual change in log real GDP on current or one-year-lagged change in the wind-plus-solar generation fraction. Both models include state and year effects, lagged log population, lagged log nominal per-capita personal income, population log growth and heating/cooling degree days. Population growth can itself be endogenous; the control set does not establish identification. The current-change common sample contains 430 observations from 40 states over 2011–2024; the lagged-change sample contains 399 observations from 40 states. Earlier Census lags and missing weather restrict this common comparison. Maximal-available samples are also supplied in the replication tables.

Table 5: Demographic and income weighting: controlled GDP-growth associations

Lagged weight	Current share change	One-year-lagged change
Equal state	0.0070 [−0.1085, 0.1225]	0.0489 [−0.0720, 0.1699]
BEA population	−0.0089 [−0.1769, 0.1592]	0.0965 [0.0027, 0.1903]
Census population	−0.0091 [−0.1772, 0.1590]	0.0963 [0.0025, 0.1901]
Total personal income	0.0025 [−0.1694, 0.1744]	0.0911 [−0.0010, 0.1832]
Per-capita income	0.0052 [−0.1116, 0.1219]	0.0428 [−0.0738, 0.1595]
Nominal GDP	0.0055 [−0.1667, 0.1777]	0.0913 [−0.0002, 0.1827]
Real GDP (sensitivity)	0.0072 [−0.1647, 0.1792]	0.0918 [0.0008, 0.1827]

Point estimates and pointwise 95 percent confidence intervals. State/year effects and the demographic/weather controls described in the text. State-clustered CR1 covariance with finite-sample correction and $G - 1$ degrees of freedom; no wild-cluster or multiple-testing adjustment is claimed for this table. Share changes are fractions. A one-percentage-point share change corresponds to coefficient times 0.01 log points, approximately the coefficient in percentage points of GDP growth. This single-lag exercise differs from the earlier four-term lag-sum specifications.

The population-weighted lagged coefficient is 0.0965 ($p = 0.044$), but excluding 2020–2021 gives $p = 0.105$ and excluding CA/NY/TX/FL gives $p = 0.056$. Its leave-one-state-out p-values range from 0.009 to 0.131. Income and nominal-GDP weights produce positive lagged point estimates without a stable five-percent rejection across influence checks.

The June 2026 Berkeley Lab target workbook provides historical RPS/CES targets for 33 jurisdictions, including 32 with RPS targets. Missing states are not coded as no-policy states. Target denominators and applicability vary across utilities and states, and source qualifications are retained. In the RPS-covered sample, adding a lagged RPS target produces a population-weighted lagged coefficient of 0.1528 (95 percent interval 0.0534–0.2522; 349 observations, 31 states, 2001–2024). The same-sample estimate without the target is 0.1524. Excluding 2020–2021 yields $p = 0.062$ in this policy-adjusted specification.

The ARPA-E public index reconciles to 1,721 unique projects, including 118 tagged Renewables. It lacks 286 release dates and 1,280 term-start dates. Displayed award amounts are multiyear metadata, not annual obligations, outlays or research expenses. Reported lead-project locations do not locate all partner research; canceled records and selected-but-not-active records remain explicitly labeled. No spending regression is estimated from these data. FERC accounting plant additions, actual capital expenditure and physical grid service likewise remain distinct concepts. No validated paired clean-investment/grid-investment panel was obtained.

USAspending records provide a separate, narrowly classified funding-exposure robustness exercise. The official transaction collection contains 28,870 unique prime grant records for 2010–2025 across assistance listings 81.087 and 81.135. These are signed calendar-year federal obligations, not annual research expense or cash outlays. Our subsequent models restrict 81.087 to the matching renewable-program title and preserve sensitivity to geography fields and relative-income scaling. Missing country codes and two code/title inconsistencies prevent treating the unfiltered administrative data as a clean research-expenditure series.

5.7 Renewable-program grant exposure and a transmission-accounting audit

The official Federal Program Inventory defines assistance listing 81.087 as renewable-energy research and development spanning solar, biomass, hydrogen/fuel cells, wind, hydropower and geothermal (https://fpi.omb.gov/program/81.087). This program is broader than wind-plus-solar electricity. We retain 21,252 prime transaction records with that code and its matching title, quarantining two inconsistent code/title records. ARPA-E remains separate rather than being labeled uniformly clean-specific. Obligations are signed commitments; deobligations, zero-dollar amendments, cash outlays, annual research expense and research output are not interchangeable.

Define grant-funding intensity $f_{it} = 10^4 B_{it}/PI_{it}$, where B is observed signed calendar-year program obligations and PI is total nominal personal income in the same year. The ratio measures basis points of state income, not real research services or a welfare weight. We estimate

$$\Delta \log GDP_{it} = a_i + b_t + \beta f_{i,t-1} + \gamma' Z_{it} + u_{it}, \tag{26}$$

with prior-year log population, prior-year log nominal per-capita income, prior-year GDP growth and current heating/cooling degree days. Absent funding cells remain missing. The controlled sample contains 702 observations from all 51 jurisdictions over 2011–2024, rather than the selected generation-share sample above.

Table 6: Prior-year renewable-program funding and GDP growth

Weight	Estimate	95% interval	Cluster p	Wild p
Equal state	0.090	[-0.048, 0.229]	0.196	0.220
Population	0.028	[-0.100, 0.157]	0.657	0.643

Estimates and intervals are multiplied by 100: percentage points of GDP growth per one-basis-point increase in prior-year obligations relative to income. They are not causal dollar multipliers. State/year effects and demographic, lagged-growth and weather controls. Equal-state or prior-year population WLS. Pointwise state-clustered CR1 intervals with $G-1$ t inference; null-imposed Rademacher wild-cluster bootstrap-t, 999 draws, seed 20261020. No multiplicity adjustment.

The controlled equal-state coefficient is 0.000905 log points per funding basis point, with a 95 percent interval of $[-0.000481, 0.002291]$; population weighting gives 0.000284 with interval $[-0.000997, 0.001566]$. Wild-cluster p-values are 0.220 and 0.643. Neither excludes zero. Single-state exclusions retain positive point estimates but no five-percent conventional rejection. Removing pandemic outcomes, restricting to 2013 onward, removing the five largest positive-obligation recipient states (CA, NY, MA, OR, UT), and exploratory lags 1–3 likewise do not yield a precise association.

Missing country codes are concentrated in 2010–2016, while most records retain valid US performance-state fields. In a matched 2018–2024 outcome sample, the lagged funding exposures under explicit-USA-country and broader reported-state rules agree exactly. This is a reporting-regime comparison, not evidence that all partner research occurred at the primary site. Comparing unmatched early and late periods would conflate metadata coverage with economic change.

As a source-quality check, the FERC utility accounts contain cross-border line endpoints, accounting transfers and adjustments, and some inconsistencies between adjacent filing-year book values. These observations do not provide geographically allocated new transmission expenditure. The detailed balance checks are documented in Appendix C and the source manifest.

6 Equal-budget policy diagnostics

We compare solar-only, grid-only, fifty-fifty, and an optimized constant-share allocation under a common *resource* budget. These are hypothetical policy diagnostics. The distinction from a fiscal comparison is material: subsidy dollars and public construction outlays need not create equal installed resources. Fiscal additionality, crowding out, tax distortions, and private cofinancing would be necessary to translate these calculations into equal-fiscal-cost policies.

In discrete annual time, $C_0 = 1$, $G_0 = 0.75$, $A = 1$, and the CES has equal weights with $\sigma = 1/3$ (the code's positive curvature parameter is 2, or $\rho = -2$ in equation 1). Depreciation is 4 percent for clean capacity and 2.5 percent for grid services. The program spends 0.1 normalized resource units in years 0–9, with unit installation costs initially one. Payments occur when investment is initiated; generation begins providing services the following year, with additional exogenous grid delays where specified. Discounting is 3 percent, and the horizon is 40 annual observations, years 0–39. The common present-value expenditure is 0.8786. The no-program benchmark allows existing stocks to depreciate without replacement. This is a finite-horizon incremental exercise, not a stationary forecast; no terminal value is assigned.

Table 7: Hypothetical equal-resource-budget results

Allocation	Solar share	PV budget	PV energy gain
Solar only	1.000	0.879	2.702
Grid only	0.000	0.879	3.410
Coordinated 50:50	0.500	0.879	6.209
Best fixed share	0.490	0.879	6.211

PV energy is a normalized service index summed over time, not dollars, observed MWh, GDP, or welfare. The best fixed share searches increments of 0.01 and is not a solution for unrestricted time-varying controls.

This illustrative near-balanced case is not representative of the full sensitivity grid. Under these assumptions, fifty-fifty allocation produces 6.209 discounted incremental service units, versus 2.702 for solar-only and 3.410 for grid-only. The maximizing fixed solar share is 0.49, giving 6.211. The near-balanced optimum is a consequence of the selected relative stocks, costs, depreciation, and curvature, not an estimated recommendation. The allocation ranking varies with the planning horizon and spending schedule.

The full sensitivity file evaluates five initial grid levels, three relative grid costs, three grid delays, and three curvature values, giving 135 parameter combinations and four allocations per combination. The grid delays are zero, three, and six additional years. They are not calibrated to project queue waiting times. Initial stocks range from 0.5 to 2 for grid capability, and relative unit costs from 0.5 to 2. Because these parameters lack empirical calibration, the sensitivity grid is a set of conditional experiments rather than a confidence distribution. The finite horizon and absence of salvage value penalize delayed benefits; conclusions beyond that horizon require a different terminal treatment. The fifty-fifty allocation exceeds both corners in 79 of the 135 specified combinations, so it fails to do so in 56; these counts describe the chosen parameter grid, not probabilities over true economies. Optimized fixed solar shares range from 0.23 to 1.00. Holding the baseline stocks, curvature, and costs fixed, a three-year extra grid delay reduces grid-only incremental services from 3.410 to 2.705 and a six-year delay reduces them to 2.128. The corresponding fifty-fifty outcomes are 5.609 and 5.005; the best fixed solar shares rise from 0.49 to 0.52 and 0.55. At six years the solar-only corner exceeds grid-only, illustrating a ranking reversal within an unchanged resource envelope.

The exercises implement the resource-allocation mechanism, not the continuous-time planner's

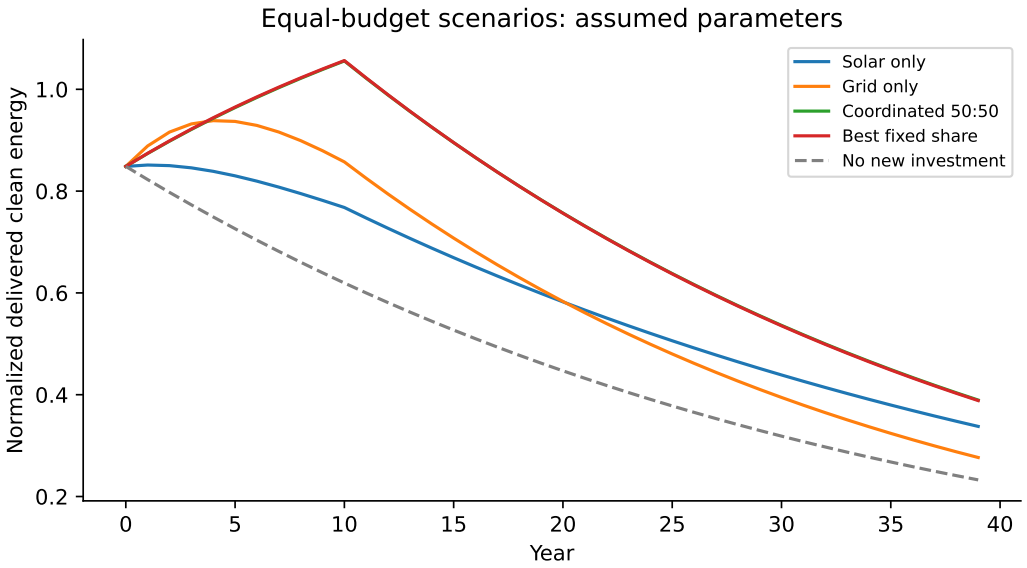


Figure 3: Assumed equal-budget service paths. The figure is generated by the model, not estimated from the observed state or queue panels.

quadratic adjustment costs. They therefore do not numerically solve the full planner equilibrium. Equation 14 and Proposition 2 describe distinct equilibrium and one-time allocation benchmarks. Keeping those objects separate prevents falsely interpreting a finite-grid scenario maximum as a dynamic optimum or a calibrated general-equilibrium result.

7 Identification, policy interpretation, and limitations

The central hypothesis remains causally untested. Grid investment responds to expected demand and renewable projects; projects enter queues where they expect profitability, not randomly. Local opposition, developer quality, financing, network upgrades, and study rules influence both duration and outcomes. Lagging transmission expenditure, adding fixed effects, or using current queue length cannot by itself solve these problems. A credible next empirical design would need a measured deliverability change and a defensible source of policy or construction timing variation, together with explicit spillover exposure. This manuscript supplies neither an instrumental variable nor a claimed natural experiment.

Geography is a second limitation. States do not coincide with electricity-market boundaries, utility territories, or transmission networks. Regional investment can affect generation and prices elsewhere, and GDP includes many sectors with no direct exposure to the local generation mix. State clustering accounts for within-state error correlation but not every cross-state shock. Queue-provider clustering has analogous limits. Nominal significance is consequently less informative than a stable, well-defined estimand supported by a credible design.

Third, public intervention needs a wedge. The baseline model is efficiently decentralized when all benefits are priced. Technical complementarity alone does not explain why government rather than private owners should invest. Network externalities, regulated access, carbon damages, financing constraints, and coordination failures are possible mechanisms, but quantifying any of them requires

additional evidence. A clean-energy subsidy may increase requests faster than delivery capability without proving that a particular grid program has positive net social value.

Fourth, the empirical panels do not calibrate the policy model. We cannot substitute a generation–capacity regression slope for the CES substitution elasticity or a queue completion probability for a unit of G. Nor can we infer a clean-R&D knowledge elasticity from a budget justification. The manuscript therefore avoids reporting invented GDP multipliers, job creation, welfare gains, or dollar-valued emissions benefits. The empirical bounds and conditional allocation analysis provide complementary measurement and theory contributions; neither establishes an identified macroeconomic return.

8 Future research and evidence requirements

The next stage should strengthen measurement and identification before expanding the number of fitted specifications. The present estimates and simulations constitute the completed analysis for this draft; the extensions below are proposed work, not additional findings.

Located investment and grid services. The immediate data priority is a small project-level pilot linking official utility-commission filings, Form 1 accounts and transmission-project records. Each observation should distinguish an approved budget from actual annual cash expenditure, a capitalized plant addition and a book-value stock. Asset routes and substations should establish geography; interstate costs should remain unallocated unless a source supports their division. The Central Maine Power and Green Mountain Power audits provide starting records, not verified state totals. Physical line characteristics should remain separate from transfer capability, congestion and reliability: route miles or nameplate ratings alone cannot identify effective deliverability. Expansion to a state-year panel should follow documented geographic coverage and accounting reconciliation.

Matched clean investment and identification. A complementarity test requires actual clean-generation investment matched to grid investment in geography and time. EIA capacity additions and generation shares can describe deployment but cannot substitute for spending. Renewable-program obligations likewise do not measure realized annual research expense. Project commissioning, construction spending, research location, cofinancing and deobligations require separate treatment. Only adequate overlapping observations would justify an investment interaction, with a stated estimand, network spillover exposure and a credible account of investment targeting. Any causal design must defend its timing variation and exclusion assumptions rather than relying on fixed effects or lags alone.

Bayesian heterogeneity. A parsimonious hierarchical model could examine whether state associations vary while partially pooling short histories. It should specify the outcome, exposure, population of states and pooling structure before fitting, assess sensitivity to weakly informative and skeptical priors, and report posterior predictive checks and time-ordered predictive comparisons with the fixed-effects benchmark. Pooling does not recover missing investment measurements or remove endogenous allocation. Posterior uncertainty would describe a conditional statistical model, not an identified causal return.

State-space dynamics. A later state-space extension could link measured generating capital, grid services and innovation to explicit transition and observation equations. Estimation should wait for informative measurements of grid capability and innovation inputs, followed by scale normalizations, parameter-recovery exercises, and checks that process noise can be distinguished from measurement error. Queue delay cannot simply be equated with a latent grid stock. Short and uneven histories, interstate links and alternative depreciation assumptions should be examined before interpreting latent states economically. Until those conditions are met, state-space estimation is deferred rather than presented as evidence of a hidden transmission effect.

9 Conclusion

The theoretical result is conditional but sharp: grid capability raises the marginal service product of generating capacity, while final-output complementarity depends additionally on substitution and the clean-service share. Optimal investment composition depends on relative user costs and inherited stocks; a fixed balanced package has no general dominance property. A public subsidy rationale requires a market failure beyond complementarity.

The observed-data results show why proposed, installed, and productive capacity cannot be treated interchangeably. Missing queue event dates make fixed-horizon operation a bounded quantity; adjusted technology differences are specification-sensitive; annual capacity is closely associated with generation but cannot reveal physical deliverability by itself. Contemporaneous GDP correlations remain imprecise; positive lagged correlations vary with weighting and sample restriction, including exclusion of the pandemic years. Renewable-program grant-funding models are also imprecise under conventional and wild-cluster inference. The contribution is a conditional allocation theory, transparent deployment measurement, and an assessment of how much the observed macroeconomic associations depend on specification and sample. The estimates do not establish that innovation has no effect. Causal transmission returns, fiscal incidence, and innovation benefits remain questions for the research program in Section 8.

A Descriptive GDP-growth robustness

Table 8: Lag-sum sensitivity: annual GDP growth and renewable-generation-share changes

Sample	Controls	Weight	N	$\widehat{\Theta}$	Cluster p	Wild p
Broad	None	Equal	350	0.004	0.984	0.986
Broad	None	GDP	350	0.694	0.036	0.070
Broad	Full	Equal	350	0.072	0.629	0.647
Broad	Full	GDP	350	0.667	0.036	0.043
No CA	Full	GDP	338	0.395	0.084	0.102
Balanced 13	None	Equal	156	0.727	0.009	0.167
Balanced 13	Full	Equal	156	0.620	0.042	0.135
Balanced 13	Full	GDP	156	1.477	0.060	0.049
No NM	Full	GDP	144	1.593	0.087	0.103

Dependent variable: change in log real state GDP. $\widehat{\Theta}$ sums exposure coefficients at lags 0–3 for annual changes in the wind-plus-solar fraction of within-state generation. All models include state/year fixed effects; each estimate is an observational association. GDP weights are prior-year state real GDP. Full controls: one-year-lagged mining, manufacturing and utilities current-dollar GDP shares, log nominal retail electricity price, state GDP growth and log real GDP. Cluster p: conventional two-sided state-clustered t inference. Wild p: restricted-null state-level Rademacher bootstrap-t (3,999 draws for broad samples, 2026-10-20 seed; exhaustive sign enumeration in the 13-state and 12-state selected panels). For the broad equal-weight baseline the wild p-value is from the 2026-10-19 independently executed robustness check. All p-values are unadjusted for specification search; exact sign enumeration does not solve weak identification or few-cluster concerns.

B Reproducibility and source scope

In the EIA original workbooks, national generation records use source code `US-TOTAL`, whereas national capacity records use `US`; these distinct national aggregate codes were separately validated.

The accompanying package contains `paper.ipynb`, `run.py`, pinned Python dependencies, LaTeX sources, the raw EIA and Berkeley Lab workbooks, the original BEA GDP archive, processed data, and generated figures and tables. The control and nonlinear analyses, together with their derived datasets, coefficient tables, and state-cluster inference scripts, appear under `extensions/controlled_2026_10_20/`. Run `python run.py` for the source-analysis pipeline; the notebook presents the separate robustness estimates and provides code to rerun the control and inference modules. The data are source vintages frozen for replication; optional online refresh can change them. Failed HTTP responses are not converted to observations.

The original-state EIA regressions, selected national totals, queue-bound counts and timing, and conditional allocation values have computational cross-checks. Such checks validate internal arithmetic, not the sources, the assumptions, or causal interpretation. The state-year robustness panel is derived from reported BEA chained-dollar aggregate GDP, nominal industry value added, and EIA electricity data; annual population, personal income, weather and policy sensitivity measures are now included in the separate `empirical_supplement/` extension. Measured physical grid deliverability and verified annual clean-specific research spending remain outside the estimated equations. The demographic supplement can be rebuilt with `python empirical_supplement/run.py`; the notebook exposes the same executed pipeline and its diagnostics. The subsequent funding models and three-year grid audit appear under `funding_grid_extension/`. No proprietary dataset is used.

C Source-accounting and model-check documentation

Official migrated FERC Form 1 archives contain 571 annual utility filings for 2014–2016, of which 465 report transmission-plant additions, with 3,350 pilot line/geographic facts. Among 316 adjacent-filing comparisons, 15 opening balances disagree with the prior reported close (maximum absolute difference \$253.7 million). Central Maine Power's line descriptions span Maine, New Hampshire, and a Massachusetts–New Hampshire boundary despite a Maine service-territory description; reporting-company headquarters therefore cannot justify assigning capitalized additions to one state. Its 2014 beginning/additions/retirements/closing balance difference of −\$1,129,705 is accounted for by reported adjustments of −\$38,067 and transfers of −\$1,091,638. These accounting observations are not a state investment panel, construction expenditure, or a physical deliverability index.

Demographic results were previously cross-checked using an independent NumPy WLS/state-score calculation (430 observations and 40 state clusters; reported standard-error discrepancy below 3×10^{-12}). The funding weighted coefficient and clustered standard error were also cross-checked in the supplied material. Those checks concern local numerical calculations, not independently verified source releases, proof of an identification strategy, or a fresh end-to-end replication. The October 20 author-review revisions did not change any regression or scenario output; full clean-environment recomputation from frozen original raw files remains a pre-submission validation task.

D Data attribution

Electricity data: U.S. Energy Information Administration, historical state electricity workbooks. GDP and industry composition: U.S. Bureau of Economic Analysis, regional annual GDP archive, SAGDP1, SAGDP2 and SAGDP9. Population and income: BEA SAINC1/SAINC30, September 30, 2026 update; Census Vintage 2019 and Vintage 2025 resident estimates. Weather: EIA SEDS 2024F indicators, with NOAA origins. Policy targets: Berkeley Lab June 2026 RPS/CES workbook. Innovation metadata: ARPA-E public project index and USAspending prime assistance transactions, accessed October 20, 2026. The historical grid pilot uses FERC's official migrated 2014–2016 Form 1 XBRL archives, with source migration warnings retained. Exact URLs, units, byte hashes and restrictions appear in the supplement source manifests and documentation. Queue data: Lawrence Berkeley National Laboratory and GridTracker (the workbook also credits Interconnection.fyi); the publisher's applicable CC BY 4.0 attribution and original author information are retained with the source files. The project retains raw methods and codebook sheets. No private information or access credentials are included. The manifest documents unresolved initial-retrieval timestamp precision rather than fabricating timestamps.

References

Daron Acemoglu, Philippe Aghion, Leonardo Bursztyn, and David Hemous. The environment and directed technical change. *American Economic Review*, 102(1):131–166, 2012. doi: 10.1257/aer.102.1.131.

Bebonchu Atems and Chelsea Hotaling. The effect of renewable and nonrenewable electricity generation on economic growth. *Energy Policy*, 112:111–118, 2018. doi: 10.1016/j.enpol.2017.10.015.

Lint Barrage and William D. Nordhaus. Policies, projections, and the social cost of carbon: Results from the DICE-2023 model. *Proceedings of the National Academy of Sciences*, 121(13), 2024. doi: 10.1073/pnas.2312030121.

Audun Botterud, Christopher R. Knittel, John Parsons, Juan Ramon Senga, and Drew Story. Bridging the gaps: The impact of interregional transmission on emissions and reliability. Technical Report 32996, National Bureau of Economic Research, 2024.

A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427, 2008. doi: 10.1162/rest.90.3.414.

Lucas W. Davis, Catherine Hausman, and Nancy L. Rose. Transmission impossible? prospects for decarbonizing the US grid. *Journal of Economic Perspectives*, 37(4):155–180, 2023. doi: 10.1257/jep.37.4.155.

Gaurav Doshi. Chicken and egg: Electricity transmission and investment in renewable energy. *Journal of Public Economics*, 256:105598, 2026. doi: 10.1016/j.jpubeco.2026.105598.

Don Fullerton and Chi L. Ta. What determines effectiveness of renewable energy standards? general equilibrium analytical model and empirical analysis. *Journal of the Association of Environmental and Resource Economists*, 12(1):65–103, 2025. URL https://www.nber.org/papers/w29783.

Todd D. Gerarden, Mar Reguant, and Daniel Yi Xu. The role of industrial policy in the renewable energy sector. *Environmental and Energy Policy and the Economy*, 7, 2026. doi: 10.1086/738542. URL https://www.nber.org/papers/w34079.

Will Gorman, Julie Mulvaney Kemp, Joseph Rand, Joachim Seel, Ryan Wiser, Nick Manderlink, Fredrich Kahrl, Kevin Porter, and Will Cotton. Grid connection barriers to renewable energy deployment in the united states. *Joule*, 9(2):101791, 2025. doi: 10.1016/j.joule.2024.11.008.

Catherine Hausman. Power flows: Transmission lines, allocative efficiency, and corporate profits. *American Economic Review*, 115(8):2574–2615, 2025. doi: 10.1257/aer.20240276. URL https://www.nber.org/papers/w32091.

Hwayoung Jeon. Renewable versus non-renewable: The role of electricity generation to economic growth. *The Electricity Journal*, 35(6):107140, 2022. doi: 10.1016/j.tej.2022.107140.

Sarah Johnston, Yifei Liu, and Chenyu Yang. An empirical analysis of the interconnection queue. Technical Report 31946, National Bureau of Economic Research, 2023. URL https://www.nber.org/papers/w31946.

Sarah Johnston, Yifei Liu, and Chenyu Yang. Grid connection costs as a barrier to new generation: Evidence from PJM and implications for transmission policy. *Energy Economics*, 161:109540, 2026. doi: 10.1016/j.eneco.2026.109540.

Charles F. Manski. *Partial Identification of Probability Distributions*. Springer, New York, 2003. doi: 10.1007/b97478.

Robert Pollin, Heidi Garrett-Peltier, James Heintz, and Bracken Hendricks. Green growth: A u.s. program for controlling climate change and expanding job opportunities. Technical report, Political Economy Research Institute and Center for American Progress, 2014. URL https://peri.umass.edu/publication/green-growth-a-u-s-program-for-controlling-climate-change-and-expanding-job-opportunities/.

Valerie A. Ramey. The macroeconomic consequences of infrastructure investment. Technical Report 27625, National Bureau of Economic Research, 2020.

Joseph Rand, Anna Cheyette, Chris Talley, Steven Zhang, Will Gorman, Ryan H. Wiser, Joachim Seel, Seongeun Jeong, and Fredrich Kahrl. Queued up: 2026 edition, characteristics of power plants seeking transmission interconnection as of the end of 2025. Technical report, Lawrence Berkeley National Laboratory, 2026. URL https://emp.lbl.gov/publications/queued-2026-edition-characteristics. Data

attribution: Lawrence Berkeley National Laboratory and GridTracker; CC BY 4.0. Accessed 2026-10-16.

Paul M. Romer. Endogenous technological change. *Journal of Political Economy*, 98(5, Part 2):S71–S102, 1990. URL https://www.nber.org/papers/w3210.

Mahmoud Salari, Inas Kelly, Nadia Doytch, and Roxana J. Javid. Economic growth and renewable and non-renewable energy consumption: Evidence from the u.s. states. *Renewable Energy*, 178:50–65, 2021. doi: 10.1016/j.renene.2021.06.016.