



**MIT CEEPR**  
Center for Energy and  
Environmental Policy Research

**Newsletter**

**SPRING 2026**



MASSACHUSETTS INSTITUTE OF TECHNOLOGY

# Editorial.

Over the months since publication of our last newsletter, the global energy landscape has been shaped less by a single transition narrative than by a convergence of shocks. Conflict in the Middle East has again exposed the fragility of oil and gas markets, while reminding policy makers that energy security remains inseparable from trade, diplomacy, and military risk. The United States, now a major exporter of crude oil, liquefied natural gas, and refined fuels, has benefited from its position as an energy supplier to allies in Europe and Asia. Yet that same export strength has transmitted global price volatility back to domestic consumers, complicating the promise that abundance alone can insulate households and firms from energy shocks.

Higher fuel prices have sharpened attention to energy affordability, as the electricity sector enters an era of rapid demand growth after years of relative stagnation. Data centers, new manufacturing, and electrification are changing the scale and geography of load growth, with pressures concentrated in regions where grid infrastructure, interconnection queues, and siting constraints have already posed mounting challenges. Electricity rates are rapidly becoming a central political issue. Yet the causes of higher utility bills are multiple: fuel costs, transmission and distribution investment, extreme weather and resilience spending, capacity procurement, and cost allocation for large new loads. As CEEPR research featured in this newsletter shows, careful causal analysis is vital before attributing affordability pressures to any one technology or policy choice.

Federal policy has added deep uncertainty to these dynamics. Following the rollback of major clean energy incentives, administrative delays affecting wind and solar projects, and renewed efforts to loosen environmental regulations, the United States is pursuing a markedly different trajectory from many peers. At the same time, market outcomes have not simply followed the direction of federal policy. Renewable energy and storage deployment continue, even as investment has become more selective, financing risk has increased, and project timelines have lengthened. Conversely, the resurgence of fossil energy has not produced uniform gains: recent profit reporting suggests that some European oil majors have benefited more from price volatility than their U.S. counterparts, owing to differences in trading operations, regional exposure, and hedging strategies.

Internationally, the divergence is widening, but not in a single direction. Europe faces high energy costs, industrial competitiveness concerns, and growing resistance to climate regulation, yet its response remains one of recalibration rather than reversal, including electricity market reform and continued decarbonization measures. China, meanwhile, is combining coal capacity, record clean energy deployment, electrification, and dominance in critical mineral and clean technology supply chains in a strategy that is as much industrial and geopolitical as it is related to the energy transition.

CEEPR work presented in this newsletter underscores how these challenges now cut across the energy system. Climate damages are already imposing significant household burdens, often unevenly distributed across income groups and geographies; oil and critical mineral price shocks affect the economy through different channels; electricity prices reflect grid investment, rate design, and cost allocation as much as generation choices; and the value of storage, nuclear flexibility, and cross-border capacity arrangements depends critically on market rules and institutional design. Meeting today's energy policy challenges requires empirical research, modeling, and institutional analysis to identify where policy can accelerate beneficial market trends, correct market failures, or risk creating new distortions. CEEPR remains committed to this task as the costs of error become ever more consequential.

Michael Mehling

## MIT Center for Energy and Environmental Policy Research

77 Massachusetts Avenue, E19-411  
Cambridge, MA 02139 USA

**MIT CEEPR Newsletter** is published twice yearly by the MIT Center for Energy and Environmental Policy Research.

### Editing/Writing:

Michael Mehling

### Design/Copy-Editing:

Tony Tran

For inquiries and/or for permission to reproduce material in this newsletter, please contact:

Email: [ceepr@mit.edu](mailto:ceepr@mit.edu)  
Phone: (617) 253-3551  
Fax: (617) 253-9845

Copyright © 2026  
MIT

# Contents.

Beyond the Early Era of EVs:  
Evidence from the Staggered  
Rollout of the HOV Lane  
Network in California

**04**

The Differential Impacts of  
Critical Mineral Prices and  
Oil Prices on the Economy

**06**

Who Bears the Burden  
of Climate Inaction?

**08**

How Uncertainty Shapes  
Electricity Storage Decisions:  
Dispatch Policies and  
Capacity Portfolios Under  
Renewable Drought Risk

**10**



From Least-Cost Plans to  
Implementable Electrification:  
A Regulatory–Financial  
Framework to Achieve  
Universal Access

**12**

Designing Second-Best Price  
Zones in Electricity Markets

**14**



Renewables and Electricity  
Affordability: Untangling  
Correlation from Causation


**16**

Planes Overhead:  
How Airplane Noise  
Impacts Home Values

**18**

The Value of Storage in  
Electricity Distribution:  
The Role of Markets

**20**



Cross-Border Solidarity Versus  
National Capacity Markets:  
Risk of Inadequate Capacity  
Procurement

**22**

Abatement or Reallocation?  
Unilateral Methane Standards  
in a Globally Traded Gas  
Market

**24**



Steady Heat, Flexible Power:  
Flexibility by Proxy in  
Nuclear-Thermal Energy  
Storage Systems

**26**

Political Ideology and  
U.S. Electric Vehicle Adoption

**28**

Events

**30**

Publications

**31**





## Research.

# Beyond the Early Era of EVs: Evidence from the Staggered Rollout of the HOV Lane Network in California

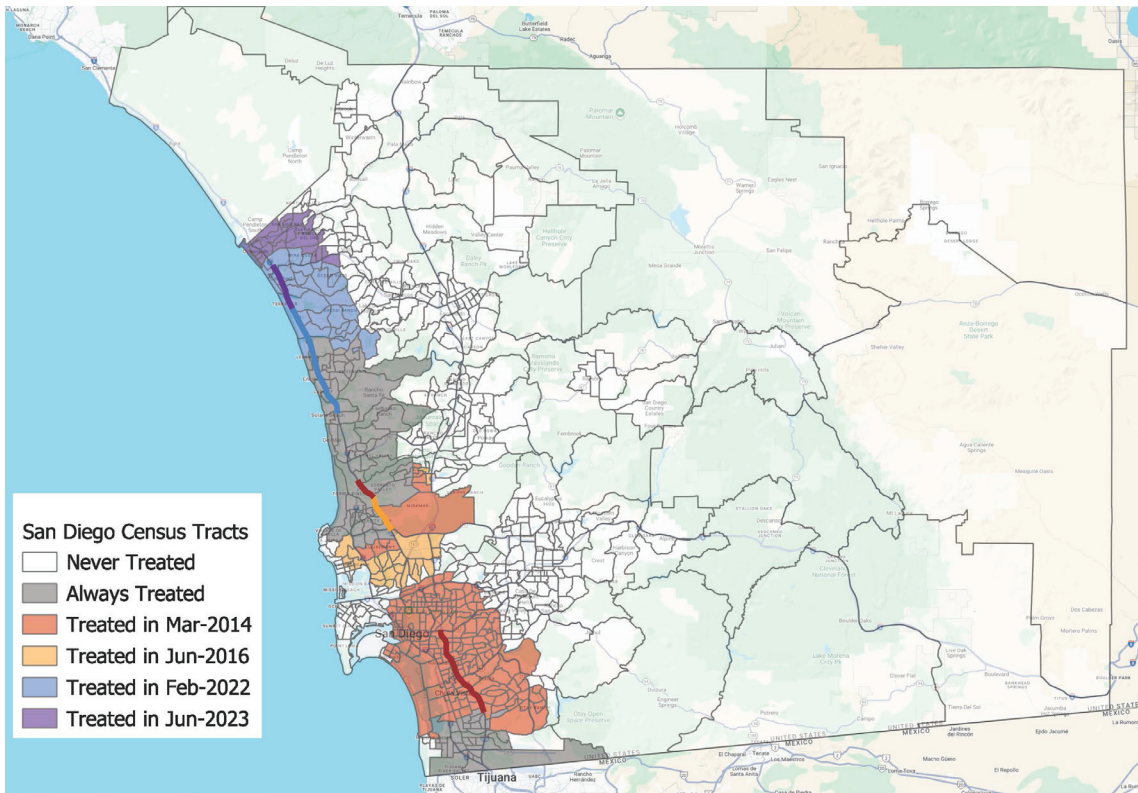
By: Eric Avis, Jing Li, and Katalin Springel

Electric vehicles (EVs) have played a central role in climate and transportation policy efforts, yet many drivers still face practical barriers to adoption. To help make EV ownership more appealing, California's Clean Air Vehicle Decal Program granted eligible EV drivers access to high-occupancy vehicle (HOV) lanes, allowing them to bypass traffic even when traveling alone. The policy offered a direct, daily benefit that made EV ownership more attractive, particularly for commuters in congested areas such as Los Angeles, San Diego, and the San Francisco Bay Area. With the program's expiration in September 2025, an important question arises: how much did these HOV lane privileges influence EV adoption, and which communities benefited the most?

We study this question using a detailed dataset that combines twelve years of monthly vehicle registration records with a spatial history of every HOV lane segment in California between 2012 and 2024. By tracking when and where each HOV segment became operational and matching that information to census tract-level vehicle registrations, we can observe how EV adoption changed in neighborhoods that suddenly gained HOV access compared with those that did not. We

exploit the staggered expansion of the HOV network in California to isolate the policy's causal impact on EV purchases rather than simply reflecting broader market trends. While prior research has documented sizeable effects of HOV incentives in the early years of the U.S. EV market (DeShazo et al., 2017; Jenn et al., 2018), the EV market and access to the HOV network have evolved substantially in recent years. Coupled with a longer time horizon and detailed census tract-level registration data, our study provides new evidence on how the role of HOV exemptions has changed in a more developed EV market and within a rapidly changing policy landscape.

Our findings demonstrate a strong, positive, and statistically significant impact of HOV lane exemptions on EV adoption in California between 2012 and 2024. When a new HOV lane segment opened within five miles of a census tract's population center, the tract-level EV market share rose by about 1.2 percentage points, an increase of roughly 16 percent relative to the average EV share at the time. Our results are robust to how we define HOV lane availability and proximity, specifically whether we use population rather than geographic centroids of census tracts and the choice of the distance threshold. We also find that the effects of HOV policies are not uniform across the state. The effect is strongest in tracts with higher shares of peak-time commuters and in those areas where commutes are more time-intensive, confirming that the value of HOV access stems from the time savings it provides.



**Figure 1. Illustration of HOV Treatment of Census Tracts in San Diego County.**

Notes: This figure illustrates how we define the treatment of census tracts based on the addition of new HOV lane segments within 5 miles of the population centroid of a tract. The map shows all census tracts in San Diego County, color-coded according to their treatment status. Census tracts that are not colored were never treated because no operational HOV lanes were within 5 miles of their population centroid at any time between 2012 and 2024. Grey tracts are always treated, as operational HOV lanes were already within 5 miles of their centroids before the start of the observation period in 2012. Red tracts gained access to an HOV lane within 5 miles of their population centroid when new HOV segments were added to the HOV lane network in March 2014. Similarly, yellow tracts were treated following HOV lane additions in June 2016, blue tracts in February 2022, and purple tracts in June 2023. The map also includes the HOV lanes added to the road network during the sample period, color-coded to match the treatment of tracts. For example, red lane segments represent those that became operational in March 2014.

Consumers may also be influenced not only by whether they have access to at least one HOV lane within their vicinity, but by the extent of HOV lane availability in their area. To capture this, our study looks beyond simple access and examines whether EV adoption increases as the number of nearby HOV segments increases. The results show that communities with more HOV lanes close by see higher EV adoption, suggesting that greater HOV lane availability also plays an important role in driving EV adoption.

However, the distribution of benefits from HOV exemptions is uneven and reveals an important equity concern. We find that the estimated impact of HOV lane access is concentrated entirely in higher-income tracts, suggesting that the policy disproportionately encouraged EV adoption among wealthier households. This distributional pattern highlights that while HOV exemptions can be effective, they may also reinforce existing disparities in access to clean technology. ■

A natural question, extending beyond prior research that focused on the early U.S. EV market, is whether and how the impact of HOV lane exemptions has changed over time as the market matured. Unlike in the early stages of the market, our study spans recent years with a more mature EV market, in which hundreds of EV models, including larger vehicle classes, are commercially available, battery technology has advanced considerably, charging networks have expanded substantially, and EVs are better understood by consumers. Our findings indicate that the policy had a larger impact in the later years of the study period, suggesting that as battery ranges improved, more EV models became available, and consumer awareness of EVs grew, drivers became increasingly responsive to non-financial incentives such as HOV access.



Avis, E., Li, J., and Springel, K. (2025), "Beyond the Early Era of EVs: Evidence from the Staggered Rollout of the HOV Lane Network in California," MIT CEEPR Working Paper 2025-21, November 2025.

Concordel, A., Ho, P., and Knittel, C. R. (2026), "The Differential Impacts of Critical Mineral Prices and Oil Prices on the Economy," MIT CEEPR Working Paper 2026-01, January 2026.



## Research.

# The Differential Impacts of Critical Mineral Prices and Oil Prices on the Economy

By: Adrien Concordel, Phuong Ho, and Christopher R. Knittel

### The New Commodity Frontier

The global economy is undergoing a fundamental structural pivot. As we transition from internal combustion to electrification, we are swapping a century-long dependence on oil for a new reliance on critical minerals like lithium, cobalt, and copper. This shift is not merely a change in resources but a change in strategic risk.

Critical mineral supply chains are significantly more concentrated than oil. China currently controls up to 80% of specific rare earth elements, exposing the energy transition to unique geopolitical bottlenecks. This vulnerability is compounded by shifting demand: by 2040, the IEA projects copper demand to surge by 50% while oil consumption is expected to fall by 25%. Consequently, policymakers must address a vital question: Do these new mineral price shocks pose the same recessionary threat as the historic oil shocks that triggered post-war downturns?

### The Fundamental Distinction: Flow Inputs vs. Capital Formation

To answer this, we provide a unified neoclassical growth model that contrasts the impacts of oil and mineral price shocks. This framework reveals that the economic impact of a shock is determined by how a commodity enters the production cycle. We distinguish between these

roles using a simple analogy: oil is the "gas in the tank," whereas minerals are the "steel and battery in the frame."

- **Oil as a Flow Input:** Oil acts as a variable operating cost. It is required to utilize existing capital (e.g., fuel to run a car already on the road). Because it is consumed contemporaneously, oil-price spikes act like an adverse cost/productivity shock, immediately reducing output and welfare.
- **Minerals as Investment Components:** Critical minerals are essential to capital formation. They affect the cost of creating new capital (e.g., the price of a new EV battery) without altering the productivity or operating cost of the existing capital stock.

Because minerals are embedded in investment goods—machinery, vehicles, and electrical equipment—their shocks propagate in a "slow but persistent" manner. While oil shocks hit the economy with immediate, fast-acting operational costs, mineral shocks act as an intertemporal drag on future capacity by raising the cost of investment.

### Why Oil Remains the Greater Threat

Our findings demonstrate that both shocks lower output in the long run, but oil price increases are systematically more contractionary and damaging to aggregate welfare.

Key Comparative Impacts:

- **Output and Welfare:** Our results show that oil shocks are more damaging. As shown in Figure 1 (Benchmark Adjustment), a doubling of oil prices (solid blue line) produces a deeper, more immediate collapse in welfare and output compared to a mineral doubling (dashed orange line).
- **Short-Term "Blips" vs. Long-Run Decline:** While oil shocks cause an immediate drop, mineral price doublings can actually produce a "small positive blip" in output on impact (approximately 0.43% in certain scenarios) before gradually drifting into a milder long-run decline.
- **Capital Contraction:** Mineral shocks generate a sharper contraction in capital stocks than oil shocks. Because minerals directly hit investment costs, firms cut new capital formation more aggressively.
- **Labor Dynamics and the Wealth Effect:** Surprisingly, mineral shocks can increase long-run employment. We find this is driven by a two-fold mechanism: firms substitute labor for expensive capital, and households supply more labor to reconstruct their balance sheets and offset the erosion of net foreign wealth.

### Policy Implications: Shifting the Stabilization Toolkit

Our model assumes prices adjust quickly. Mineral shocks mainly hit

investment and foreign borrowing. Demand-side stimulus plays a smaller role in this setting. So, policy should shift toward investment stability and financial resilience. For example:

**Macprudential Tools:**

Prioritize countercyclical capital buffers and stabilization funds to smooth the balance-sheet volatility and investment fluctuations caused by mineral price cycles.

**Technological Diversification:**

Success depends on increasing the "elasticity of substitution." Investment must be funneled into R&D for material efficiency and "swap-ready" chemistries such as transitioning to sodium-ion batteries or rare-earth-lean motor technologies. This allows the economy to swap chemistries as prices fluctuate, effectively muting the impact of any single mineral spike.

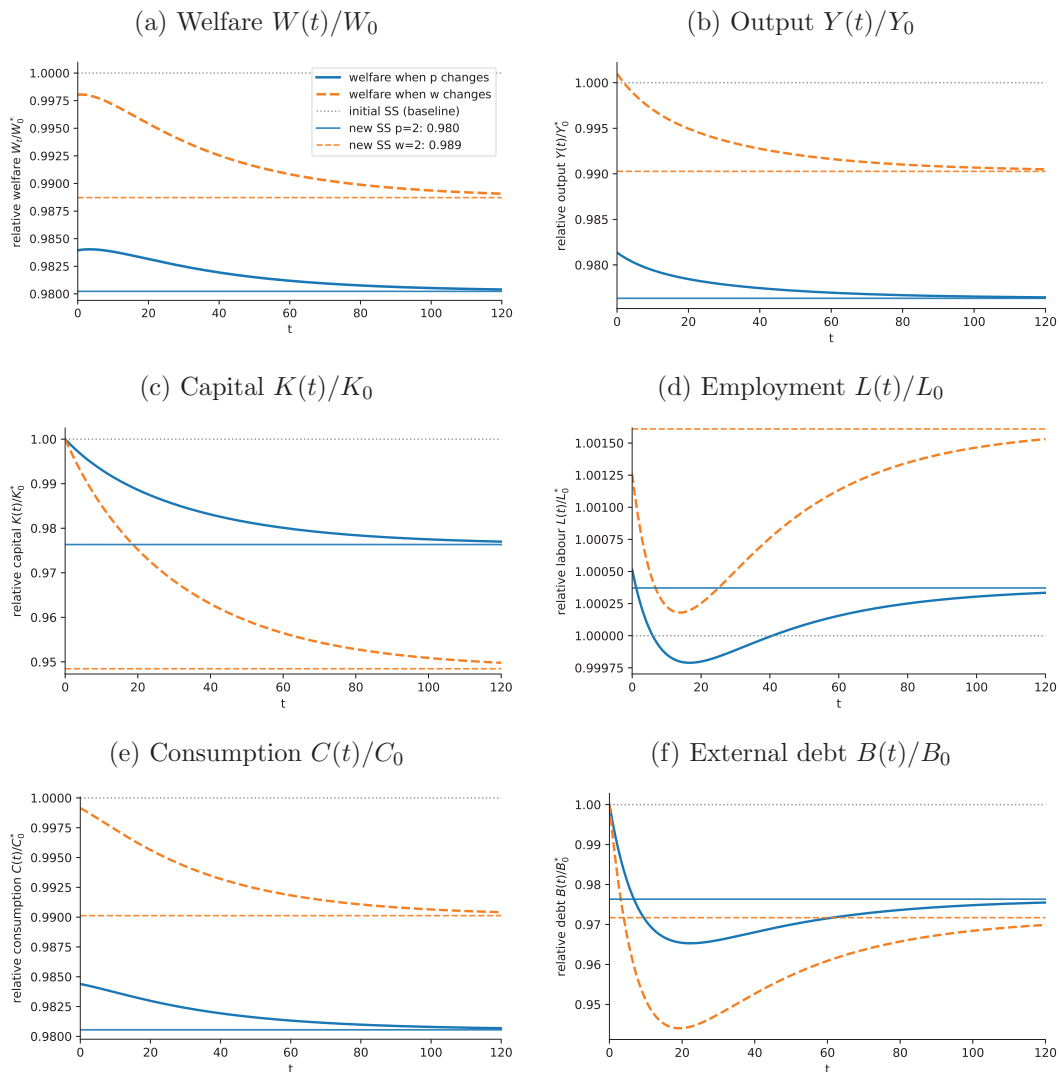
**Financial Resilience:**

Maintain precautionary wealth buffers (e.g., foreign-exchange reserves) to manage the debt-cycle impacts of mineral shocks.

**Navigating the Energy Transition**

The transition to minerals introduces new complexities but does not replicate the existential vulnerability of the "oil era." The economic nature of minerals as an investment component provides a natural buffer for aggregate activity.

For advanced economies, mineral price spikes are less likely than oil spikes to cause an economy-wide slump. They mainly affect investment and foreign debt. By focusing on financial stability and technological flexibility, we can manage the "wealth/user-cost shocks" of the new commodity frontier without derailing long-term growth.



**Figure 1. How the Economy Adjusts After Oil Prices Double Versus Critical Mineral Prices Double (OECD Calibration).**

Notes: The solid line reports the response to an oil price doubling. The dashed line reports the response to a mineral price doubling. Values are shown relative to the initial steady-state values (pre-shock level 1.0). The horizontal axis is time after the shock. Calibration is based on OECD economy parameters.

Clausing, K. A., Knittel, C. R., and Wolfram, C. (2026), "Who Bears the Cost of Climate Inaction?," MIT CEEPR Working Paper 2026-02, February 2026.



## Research.

# Who Bears the Burden of Climate Inaction?

By: **Kimberly A. Clausing, Christopher R. Knittel, and Catherine Wolfram**

### Extreme Weather as the Dominant Cost Driver

A central finding of the paper is that extreme weather events—not incremental warming—account for most household burdens today. Over the past three decades, the United States has experienced rising frequencies of billion-dollar disasters, from hurricanes and floods to wildfires and severe storms. These events create financial pressures through higher insurance premiums, infrastructure disruptions, and utility price increases, while also imposing significant threats to life. Although heat exposure has increased across most U.S. counties, its net effect on mortality remains small. In many regions, declines in cold-related mortality have roughly offset increases in heat-related mortality. By contrast, wildfire smoke and extreme weather events have created sharp increases in mortality and financial damages, and these trends have accelerated over the past decade.

### Insurance Markets: A Growing Financial Strain

Household insurance premiums have risen rapidly, particularly in disaster-prone regions. Drawing on mortgage escrow data from Keys and Mulder (2024), the authors estimate that climate-driven increases in hazard risk have raised homeowners' insurance costs by an average of \$73 per year under more conservative assumptions and up to \$356 per year under less conservative attributions to climate change. These increases are notably higher in the South, Gulf Coast, and wildfire-prone West.

These estimates exclude major perils—such as flooding and storm surge—that are not covered by standard homeowners' policies. Supplementing with First Street Foundation data, the authors find additional flood- and wind-related losses that households implicitly bear through either insurance gaps or taxpayer-funded programs.

When indirect commercial pass-through is included, insurance-related costs rise by another \$32–\$145 per household.

The distributional pattern is stark: losses are relatively flat in dollar terms across income groups but sharply regressive when viewed at its share of households' income.

### Energy Expenditures: Modest Physical Use Impacts, Large Price Effects

Climate change affects household energy spending in two ways: Changes in the quantity of energy used (driven by shifts in heating and cooling needs) and changes in the price of energy (driven by disaster-related infrastructure costs).

Using machine-learning models applied to RECS microdata, the authors estimate that warming since 1990 has increased cooling expenditures by \$25–\$33 per household, while reducing winter heating bills from natural gas, propane, and kerosene, resulting in a net but modest increase in energy consumption of about \$11 per household. The more significant effect comes from utility rate increases.

Disaster recovery, wildfire mitigation, and storm hardening have placed new financial burdens on electric utilities, which are increasingly passed on to ratepayers. Based on electricity price models using NOAA disaster data, households now pay an additional \$30 per year on average for electricity due to climate-related infrastructure costs—with some regions, such as the South and West, experiencing increases above \$150–\$200 per year.

These energy cost increases are also regressive: low-income households devote a larger share of their income to energy and have fewer adaptation options.

### Public Expenditures and Taxpayer Costs

Governments also bear significant costs from climate-related disasters which the authors measure through the disaster relief outlays dedicated towards infrastructure repair, social insurance programs, and ad hoc emergency assistance to affected households and firms. By aggregating FEMA, HUD, state, and local expenditures from 2017-2021 and delineating allocations for storm- and wildfire-related disasters, the authors find that the between \$41 and \$77 of the taxes paid by the average American household go towards these relief efforts. Since these disbursements are ultimately funded by taxpayers, the greater frequency of catastrophic events may induce greater tax burdens which themselves vary greatly across states.

### Mortality Impacts: Wildfire Smoke Dominates

The mortality effects of climate change differ sharply across geography. While heat-related mortality has changed only modestly (though higher among Black Americans and low-income individuals), wildfire smoke has emerged as a major and rapidly growing health threat.

Using updated particulate-matter exposure data from Childs et al. (2022) and quasi-experimental mortality estimates by Deryugina et al. (2019), the authors find that wildfire smoke contributed to over 35,000 deaths in 2024, with an average annual household mortality cost of \$100 attributed to climate change. These mortality outcomes are concentrated in Western states and disproportionately harm lower-

income counties highlighting regressivity in particulate matter impacts.


These estimates include only mortality. Because emergency department visits, hospitalizations, medication use, and lost workdays are excluded, the true health cost is significantly higher.

**Conclusion: Climate Change Is Already Costly—and Unequally So**


Across insurance markets, energy bills, public infrastructure, and health outcomes, the authors find that the costs of recent climate change are substantial, growing, and regressive. The most consequential impacts

today come not from incremental warming but from extreme weather events and wildfire smoke, which disproportionately affect vulnerable households and regions.

Even under conservative assumptions, the analysis suggests that climate inaction already costs U.S. households tens of billions of dollars each year. As extreme weather intensifies in coming decades, these burdens are likely to increase dramatically unless mitigation, adaptation, and resilience measures keep pace. ■



(a) More Conservative Estimate (winsorized at 95%)



(b) Less Conservative Estimate (winsorized at 95%)

Figure 1. Aggregate Costs per Household by County.

## Research.

# How Uncertainty Shapes Electricity Storage Decisions: Dispatch Policies and Capacity Portfolios Under Renewable Drought Risk

By: Cem Keske, Bjarne Steffen, and John E. Parsons

Electricity storage can serve the electric grid in a variety of ways. Historically, it has enabled utilities to exploit economies of scale in generation capacity in the face of predictable diurnal variations in hourly load. Different storage technologies have been matched with different generation technologies. The large-scale expansion of pumped hydro storage in many countries between 1960 and 1990 complemented investments in baseload nuclear and coal plants. Thermal storage has been integrated into concentrated solar power plants to economize on the size of the power island in the face of dramatic, predictable diurnal variations in insolation. Grid-scale battery storage is proving a valuable complement to the large-scale penetration of solar PV, especially with regard to managing the regular evening ramp. Storage is also useful to manage the stochastic variation in load and renewable generation at short time scales. In many electric grids, battery storage is becoming an important source of frequency regulation and fast response operating reserves.

Storage assets are also now beginning to be assessed for how they can help manage stochastic variation at longer time scales. Hydro reservoir dominated grids have long managed the inventory of water in the face of seasonal and inter-annual variation in water inflows. Now, other operators managing the increasing penetration of wind and solar PV generation are also learning how to assess the danger of renewable resource droughts and attend to the stock of stored energy in their systems. Many grid operators have been working to update their metrics for resource adequacy. Capacity markets in many systems are being adapted to incorporate more sophisticated measures of how each resource contribute to system capacity, including storage assets (a.k.a. energy duration limited resources).

Our model is tailored to help illuminate how storage serves resource

adequacy, especially in light of the danger of renewable resource droughts. We eliminate any deterministic calendar variations (e.g., daily cycles) and focus exclusively on optimization against a stationary uncertain inflow of renewable energy. In spirit then, this is a model more inspired by the volatility characterizing wind dominated systems than solar PV dominated systems. Our model is constructed with hourly granularity, focusing on the provision of energy, so we do not consider the provision of services at a shorter time scale such as frequency regulation, operating reserves, or other ancillary services. Subject to that proviso, the storage asset is optimized in the face of variability at all other time scales, whether hourly variations or the long-run danger of droughts.

Figure 1 illustrates the effect of a key parameter of our model of the uncertain availability factor of our renewable energy, the rate of decay. The figure shows two simulations of the renewable availability factor across three weeks. The two simulations have the same long-term distributions. However, they differ in the rate of decay: in the blue simulation, when the availability factor is above or below its mean value, it can be expected to return to the mean faster, while in the red simulation it is expected to return to the mean more slowly. The blue simulation experiences more frequent, but shorter scarcity events—i.e., fewer prolonged droughts. The red simulation experiences fewer, but longer lasting scarcity events—i.e., more prolonged drought episodes.

We solve for the optimal operation of a battery as a function of the rate of decay, among other drivers. Figure 2 illustrates the optimal dispatch strategy for an 8 MW, 8-hour duration battery with a 90% roundtrip efficiency on a hypothetical grid with an hourly load of 150 MW, 200 MW of renewable capacity, 60 MW of baseload thermal generation, and 40 MW of peaking thermal generation. The two dimensions of the figure define the state-space on which the dispatch can be conditioned.




Figure 1. Two Simulations of Renewable Availability Factors Reflecting Different Parameters for Rate of Decay.

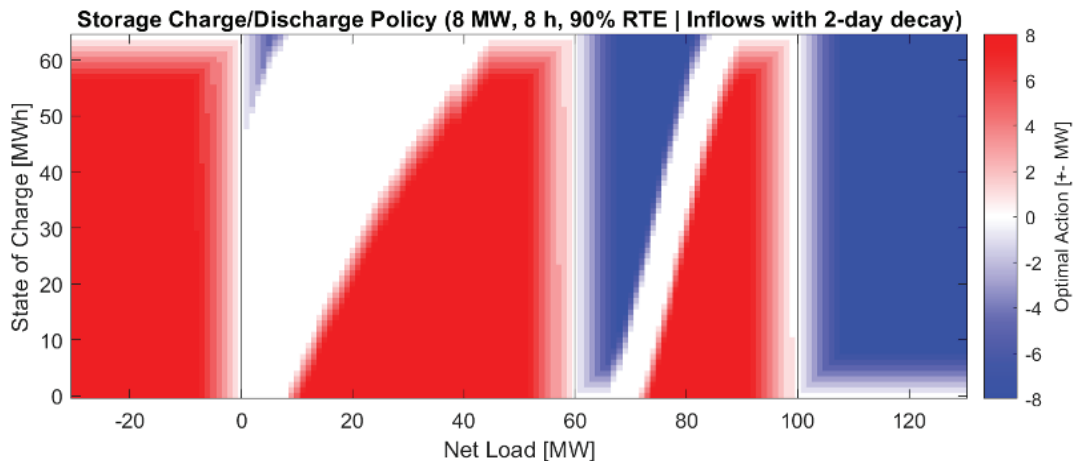


Figure 2. The Optimal Dispatch Policy.



The horizontal axis is this hour's net load before operation of the battery: when the renewable capacity factor is high, the net load is low. On the far left, below 0, renewable generation would be curtailed unless the battery is charging. On the far right, beyond 100, non-battery generation is insufficient to supply load and without the battery discharging there would be a loss-of-load event. Between 0 and 60 on the horizontal axis, the baseload generator is the marginal unit, while between 60 and 100, the peaking generator is marginal. The decay parameter determines how long net load can be expected to remain far from the center after an unexpectedly high or low hour.

The vertical axis is the battery's own state-of-charge. At the top, at 64 MWh, the battery is full. At the bottom, the battery is empty.

The color scheme describes the dispatch strategy: red means it is charging, blue means it is discharging and white means it is holding its charge. Unsurprisingly, when renewables would otherwise be curtailed—on the far left—the battery is charging, and when there would otherwise be a loss-of-load event—on the far right—the battery is discharging. What the optimization informs us about is how the battery is managed in less clear cut cases: for example when net load is between 60 and 100 and the expensive peaker plant is operating. Net load is expected to move back below 60, but there is always the possibility that it might move above 100. If the battery discharges above 60 and net load does move below 60, then the battery can be recharged at the lower price determined by the baseload plant and the system has avoided utilizing the expensive peaker when net load was briefly above 60. However, if the net load stays above 60, the battery will eventually be exhausted. Then, if net load increases further, towards 100 and above, if the battery is to provide any resource adequacy, it must recharge itself using the expensive peaking generation. The optimization is the result of balancing these possibilities.

As we decrease the rate of decay, making a persistent excursion towards a loss-of-load event more likely, the space for arbitraging the peaking and baseload generator is reduced and the battery holds its charge more often.

The paper also shows how the dispatch strategy varies with the storage asset's duration and roundtrip efficiency. We also illustrate how the dispatch of one storage asset depends on the features and state of charge of other storage assets in the system, and demonstrate conditions under which it is optimal for one asset to charge while another is discharging. Finally, we show how these drivers affect the optimal portfolio of capacities, system cost, and loss-of-load events. ■■



Keske, C., Steffen, B., and Parsons, J. E. (2026), "How Uncertainty Shapes Electricity Storage Decisions: Dispatch Policies and Capacity Portfolios Under Renewable Drought Risk," MIT CEEPR Working Paper 2026-03, February 2026.

## Research.

# From Least-Cost Plans to Implementable Electrification: A Regulatory–Financial Framework to Achieve Universal Access

By: Santos J. Díaz-Pastor, Carlos de Abajo, Robert Stoner, and Ignacio J. Pérez-Arriaga

Universal electricity access remains a development priority and a climate enabler. Yet the global trajectory still falls short. IEA projections suggest that about 645 million people could remain without electricity by 2030, with more than 80 percent in Africa (IEA et al., 2025). Africa also needs a sharp increase in capital inflows, on the order of \$23 billion per year between 2025 and 2035, to meet access targets (IEA, 2025). The main barrier is no longer the lack of least-cost planning tools. It is the difficulty of turning technically coherent plans into financeable investment programs.

Modern geospatial electrification algorithms have achieved unprecedented analytical precision. Today, these computational tools can seamlessly delineate the economically optimal deployment of grid extensions, localized mini-grids, and stand-alone systems down to the household level, offering highly granular spatial resolutions that were previously unattainable. Governments and development partners now use these tools to define national roadmaps. Many of these roadmaps then stall at implementation because planners treat finance as an exogenous input, not as an outcome of regulation, utility performance, and payment security. Recent work has shown how ignoring financing assumptions can bias electrification outcomes and understate implementation constraints (Agutu et al., 2022).

This working paper proposes an integrated regulatory and financial framework that addresses this issue. It links three components that often sit in separate conversations. It starts from a least-cost plan that produces a detailed time profile of capital and operating expenditures across on-grid and off-grid options. It then assigns clear implementation responsibilities through licenses or concessions, with a universal service obligation, and places these responsibilities under cost-of-service regulation. It finally designs a financing architecture that coordinates public resources, concessional instruments, and commercial capital in a way that matches cash needs over time.

The context matters for bankability. In many low- and middle-income power sectors, distribution companies face chronic financial distress.

Governments cap tariffs for affordability. Utilities also lose revenue through technical losses, weak billing, and low collection. These frictions create quasi-fiscal deficits and deter investment (Foster and Rana, 2020). They also raise perceived risk and the cost of capital. On African power projects, risk premiums and interest rates can be two to three times those in advanced economies. Private finance then concentrates on generation projects with contracted revenues, while networks and last-mile delivery lag.

Cost-of-service regulation provides a disciplined way to define what it should cost to deliver service and what an efficient operator should earn. The regulator calculates an allowed annual revenue requirement. It includes efficient operating costs, depreciation of invested assets, and an allowed return on the regulated asset base. This revenue requirement is not a planning aspiration. It is the benchmark that investors and lenders use to assess revenue adequacy and credit risk.

Implementation depends on a simple cash logic that many plans leave implicit. Each year, an operator must receive enough cash to cover the allowed revenue requirement. Customer tariffs can provide part of that cash flow, but affordability constraints often bind. When customer payments fall short, a supplementary payment must be made to cover the difference in a timely and predictable way. Annual revenue reconciliation is imperative to bridge the gap between allowed revenues and what the operator can actually collect.

This cash-flow view separates two distinct constraints. The first is the viability gap. It arises when affordable tariffs and weak retail performance imply that cash revenues remain below the regulated cost of service over the relevant horizon. In that case, the business is structurally unable to recover efficient costs. No financial engineering can fix that problem. Policymakers must either raise revenues, reduce efficient costs, or provide explicit compensation that is rule-based and predictable.

The second constraint is financing needs. It comes from timing, not from long-run profitability. Electrification requires front-loaded capital expenditures and working-capital outlays. Cost-of-service regulation recognizes these costs over the useful life of the assets, through a regulated revenue requirement that spreads recovery over time. In practice, the regulator does not “pay back” CAPEX in the year it is incurred. Instead, it allows (i) depreciation, which returns principal gradually, and (ii) a return on the undepreciated regulated asset base (RAB), which compensates the investor for tying up capital. Revenues therefore follow the path of regulated cost recognition and tariff pass-through, while the operator pays real cash costs upfront when it builds lines, installs transformers, or deploys off-grid assets. This intertemporal mismatch between cash outlays and regulated revenue accrual creates early-year cash deficits and liquidity stress even when the program is viable in present value terms. The key point is that “profitability” is assessed over the asset life, but lenders and suppliers require cash in the construction and ramp-up years.

These two constraints require different instruments. Viability gaps require an explicit and credible compensation mechanism. This mechanism can combine targeted budget transfers from the general government budget, earmarked resources through an access fund that can be capitalized by the state, donors, or levies, or regulated cross-subsidies within the tariff system. The instrument matters less than its credibility as a predictable cash-flow commitment. If payments arrive late, remain

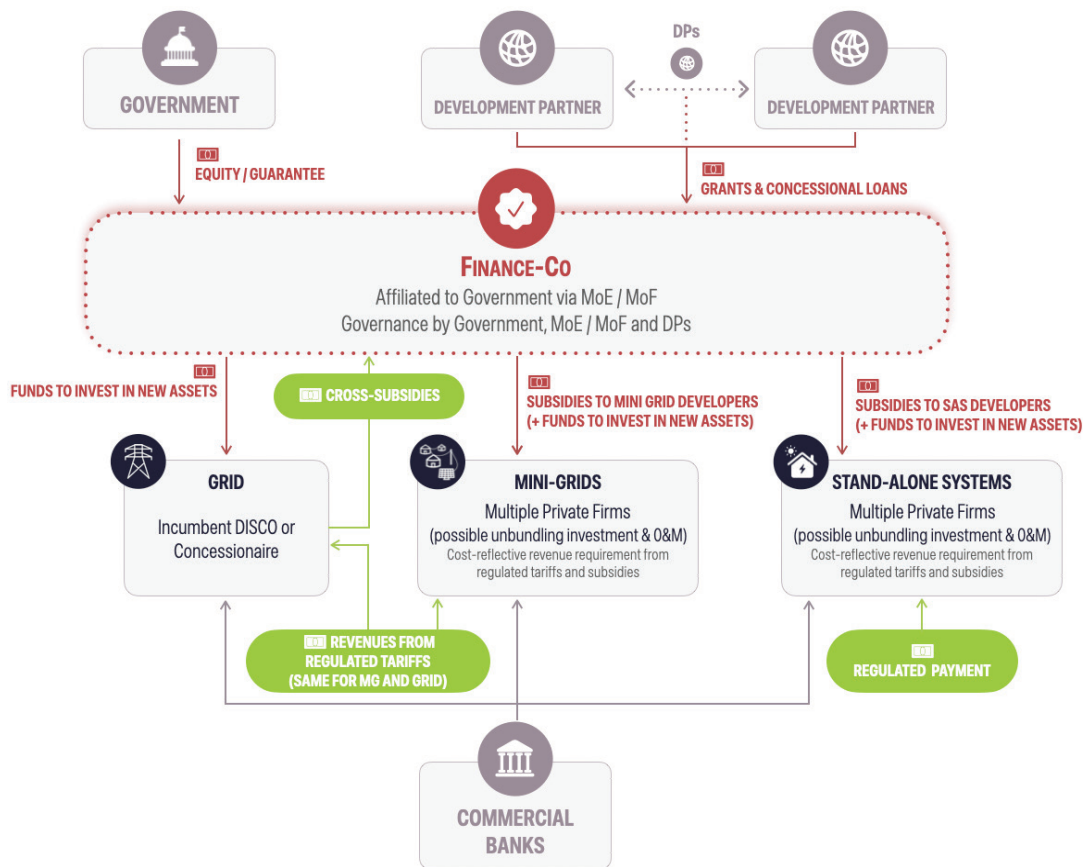


Figure 1. Stylized Cash Flow Structure of an Integrated Financing Vehicle for Universal Electrification, Showing Funding Sources, Subsidy Channels, and Implementing Entities Across Grid Extension, Mini-Grids, and Stand-Alone Systems.

discretionary, or are exposed to annual appropriation uncertainty, investors will price that risk, and the cost of capital will increase. Financing needs instead require bridge finance and capital structures aligned with the regulatory recovery schedule. Longer tenors, grace periods, sculpted debt service, and liquidity facilities can keep early-year cash deficits from becoming implementation failures.

The framework treats delivery models as technology-neutral. Grid extension, mini-grids, and stand-alone systems all face the same economics when governments expect them to provide permanent service. Each implementing agent needs a stable revenue model anchored in efficient costs and enforceable service standards. For off-grid provision, this implies a “utility-like” approach rather than a pilot logic. Competitive procurement and concessions can still play a central role. They can allocate territories and reveal the minimum subsidy required to meet a defined service standard, while regulation anchors long-run revenue adequacy.

A central contribution of the paper is a financing architecture that can operationalize these principles at scale. As defined in Figure 1, we propose a pooled platform (“Finance-Co”) that consolidates public resources, development finance, and sector-based cross-subsidies. Finance-Co ringfences inflows and standardizes disbursement rules. It channels funds to grid, and off-grid implementers based on verified delivery and the annual revenue reconciliation described above. This design reduces fragmentation across donor programs, improves transparency, and strengthens payment security.

Finance-Co does not replace strong regulation. It strengthens it by making cash transfers auditable and predictable. It also creates a natural home for blended finance. Concessional loans, guarantees, and grants can reduce the sector’s effective cost of capital and lower the long-run burden on tariffs and budgets. The platform can allocate these instruments to the specific constraint that binds. Grants and recurrent transfers address the viability gap when affordability is structurally binding. Liquidity facilities and long-tenor debt address financing needs during rollout.

The policy message is direct. Least-cost planning remains necessary, because it reduces the investment requirement and the subsidy envelope. It is not sufficient. A plan becomes implementable only when it specifies enforceable obligations, a transparent annual revenue reconciliation, and a financing structure that covers early-year cash deficits. This paper provides a replicable methodology to do that in a way that is consistent with affordability and fiscal constraints. The full working paper formalizes the cash-flow metrics and the institutional design choices. It is intended as a practical guide for regulators, ministries of finance, development partners, and investors who need universal access plans that can actually be financed and delivered. ■

Díaz-Pastor, S. J., de Abajo, C., Stoner, R., and Pérez-Arriaga, I. J. (2026), “From Least-Cost Plans to Implementable Electrification: A Regulatory–Financial Framework to Achieve Universal Access,” MIT CEEPR Working Paper 2026-04, March 2026.



## Research.

# Designing Second-Best Price Zones in Electricity Markets

By: Jonas Boeschemeier and Sebastian Schwenen

### Electricity Markets in North America

More than two-thirds of electricity in North America is traded in deregulated wholesale markets. These markets use a system called locational marginal pricing (LMP), where generators are paid prices that reflect local supply and demand conditions at each point in the grid. These prices can vary enormously across locations, driven mainly by transmission congestion.

Consumers, however, do not pay these fine-grained prices. Instead, thousands of grid locations are grouped into a small number of broad price zones, where everyone pays the same average price. When real local scarcity is hidden behind a zone-wide average, some consumers overpay and others underpay, a mismatch that reduces the efficiency of the market.

The problem is that most zone boundaries were drawn when these markets were first set up, in some cases more than 25 years ago. In California, just three zones cover the entire state. These configurations have not kept pace with changing generation patterns, load growth, or grid constraints, as our paper shows.

### A New Tool for Measuring Zone Efficiency

We develop a straightforward framework to measure how well price zones reflect actual price variation across the grid. Building on Jacobsen et al. (2020), our central metric is a new efficiency measure running from 0% to 100%, which captures how close a given zone configuration comes to perfectly efficient pricing. A score of 100% means zone prices fully reflect local conditions, with no efficiency loss from aggregation. A score of 0% means zone prices are no better than charging everyone in the market the same flat rate.

### How Do Current Zones Perform?

We apply this framework to three major U.S. electricity markets using five years of hourly price data from 2020 to 2024.

- *New York (NYISO)* has 547 pricing nodes in our sample, grouped into 11 zones, and scores 79%. Its zones do a good job of tracking local price conditions.
- *New England (ISO-NE)* has 1,005 nodes in our sample and 8 zones, scoring 37%. Zones explain less of the local variation, though spatial price differences in New England are relatively small to begin with, so the efficiency loss in dollar terms remains limited.
- *California (CAISO)* has 1,398 nodes in our sample but only 3 zones, scoring just 29%. Performance has also been getting worse over recent years, with growing price differences within zones pointing to a configuration increasingly out of step with how the grid actually operates.

Taken together, these results suggest that existing zone boundaries, particularly in California, no longer reflect grid conditions well and that redesign is worth serious consideration.

### What Would Better Zones Look Like?

To identify efficiency-maximizing alternatives, we apply machine learning clustering techniques that group nodes with similar price patterns together, following the approach of Astier (2021). Our algorithm can incorporate price data alone, or augment it with geographic coordinates and institutional boundaries to produce configurations that are both economically efficient and practically implementable.

A first takeaway is that a handful of well-designed zones already captures most of the spatial variation in all three markets, with gains diminishing quickly as more zones are added. More importantly, a substantial share of efficiency losses in each market stems not from

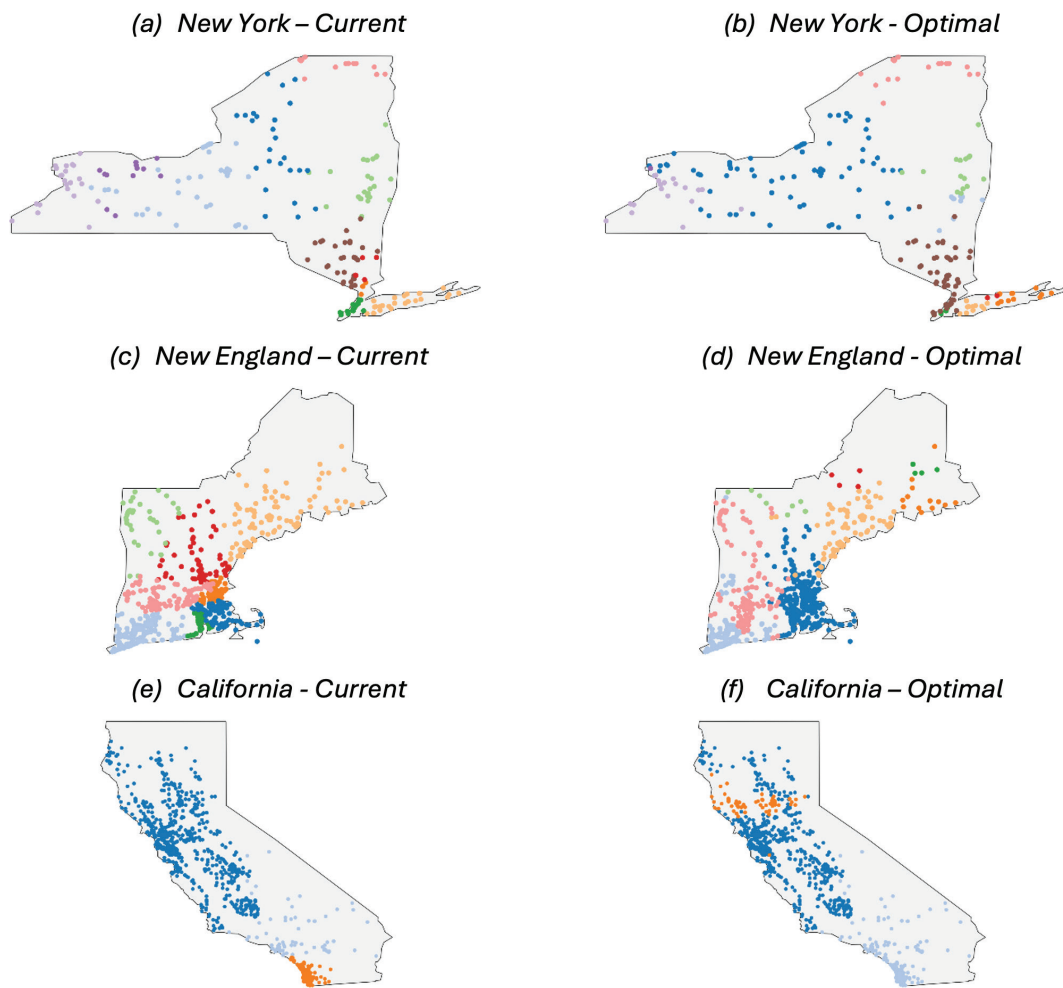


Figure 1. Current Price Zones Versus Optimal Counterfactual Price Zones.

having too few zones, but simply from having boundaries drawn in the wrong places.

That misalignment is modest in New York, where the current 11-zone configuration is close to optimal. In California, restructuring the same three zones based on price patterns would raise efficiency from 29% to 37%. The gap is largest in New England, where the current 8-zone structure scores just 37%, while an optimally assigned 8-zone configuration reaches 73%. In all three markets, an alternative fix also delivers meaningful gains: just splitting the single least efficient zone raises efficiency substantially without requiring a full redesign.

When we map the counterfactual zones, as shown in Figure 1, a striking pattern emerges: price-based clustering produces geographically contiguous regions even without using location as an input, because prices already encode the spatial structure of the grid.

In New York, the main change is that Long Island, currently one zone, is split into three subzones.

In New England, Rhode Island and eastern Massachusetts are merged into one zone, while Maine is split into four.

In California, the two southern zones are consolidated into one, while the large northern zone of Pacific Gas & Electric is subdivided.

## Conclusion and Policy Implications

Our results suggest that zone boundaries deserve more attention than they typically receive. For regulators and market operators, meaningful efficiency improvements appear achievable without radical change, often by splitting a single poorly designed zone or modestly redrawing a few boundaries. Given that grid conditions continue to evolve with the energy transition, periodic reassessment of zone configurations seems a reasonable and relatively low-cost policy step. ■



Boeschemeier, J. and Schwenen, S. (2026), "Designing Second-Best Price Zones in Electricity Markets," MIT CEEPR Working Paper 2026-05, March 2026.

## Research.

# Renewables and Electricity Affordability: Untangling Correlation from Causation

By: Fischer J. Espiritu Argosino and  
Christopher R. Knittel

Energy prices nationwide have been rising at roughly twice the rate of inflation. Many observers link this trend to the rapid growth of renewable generators, particularly utility-scale wind and solar. This has contributed to the rescissions of the incentives and policies that previously helped renewables account for over 80% of new U.S. generation capacity in 2024 (FERC 2025). Prior studies argue that RPSs can raise electricity prices by imposing high compliance costs on utilities in certain states (Wiser 2025). Building on this literature, this paper uses a fixed-effects econometric approach to estimate plausibly causal effects of RPSs and renewable generation on residential electricity prices and utility expenditures over time and across states.

We draw upon data from the Energy Information Administration, the Federal Energy Regulatory Commission, and the Lawrence Berkeley National Laboratory to construct a state-utility year panel dataset that links residential electricity prices, RPSs, generation mixes, and utility spending. We focus our analysis on the most recent 25 years of complete data from 1998 through 2023. This dataset enables us to address two crucial questions: to what extent can we attribute increases in residential electricity prices to a decarbonizing electricity sector, and how does greater renewable generation impact utility spending?

Using a two way fixed effects regression model that controls for unobserved differences across years and between states, we find that RPSs and utility scale renewables are not associated with higher residential electricity prices. Instead, RPSs have virtually zero impact on prices and utility scale renewables are modestly associated with lower prices; in particular, utility-scale solar is strongly linked with lower prices. One explanation is that large solar plants may reduce network congestion and place downward pressure on overall generation costs because they do not require fuel inputs. We test these cost reduction channels using a three way fixed effects model that also accounts for differences between utilities. Our results show that utility scale solar is strongly associated with lower transmission, distribution, and total operation and maintenance (O&M) costs. These cost reductions may be passed through to customers in the form of lower retail electricity prices.

We find rooftop solar, however, to be significantly correlated with higher prices. One possible reason for this is that current rate structures

in some states shift fixed costs associated with rooftop solar from the owners onto other ratepayers. Many utility bills include per kilowatt hour charges for the utility to compensate rooftop solar owners for the excess electricity they export to the grid. This rate structure can result in some rooftop solar owners receiving negative electricity bills, meaning they end the month with a credit rather than a charge. In effect, this means that those without rooftop solar often fund those who do through an additional volumetric rate. In addition to this equity concern, increased rooftop solar generation is often used as a reason to delay investment in the distribution network because it can reduce peak demand hours by allowing more homes to self consume electricity. While certainly true in some scenarios, our analysis finds that rooftop solar is strongly correlated with greater distribution O&M costs. One reason for this is that the distribution network was not designed for bidirectional power flows. By deferring investment in distribution networks while offering incentives for rooftop solar exports, utilities may need to spend more to manage the physical and administrative complexity of harmonizing small, decentralized generators with increasingly antiquated infrastructure.

Even though our results present strong evidence for technological and policy drivers of both residential electricity prices and utility spending, additional confounders may remain. Our ongoing research applies more rigorous econometric techniques to further strengthen the causal interpretation of these estimates. Nevertheless, our work offers robust evidence that asserts neither RPSs nor utility scale wind and solar are responsible for increased residential electricity prices. We find instead that rooftop solar generation is linked with higher prices, likely because of current cost-shifting rate structures that may not fully account for the


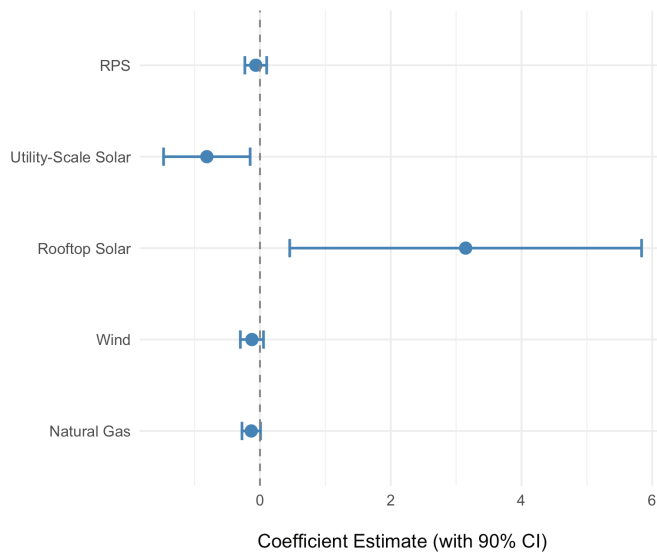


Figure 1. US Residential Electricity Price Trends by State.

This chart plots the nominal percentage changes of residential electricity prices relative to 1998 values. Over the past 25 years, household residential electricity prices have risen substantially, with many that have exceeded the rate of inflation as indicated by the solid black CPI line.

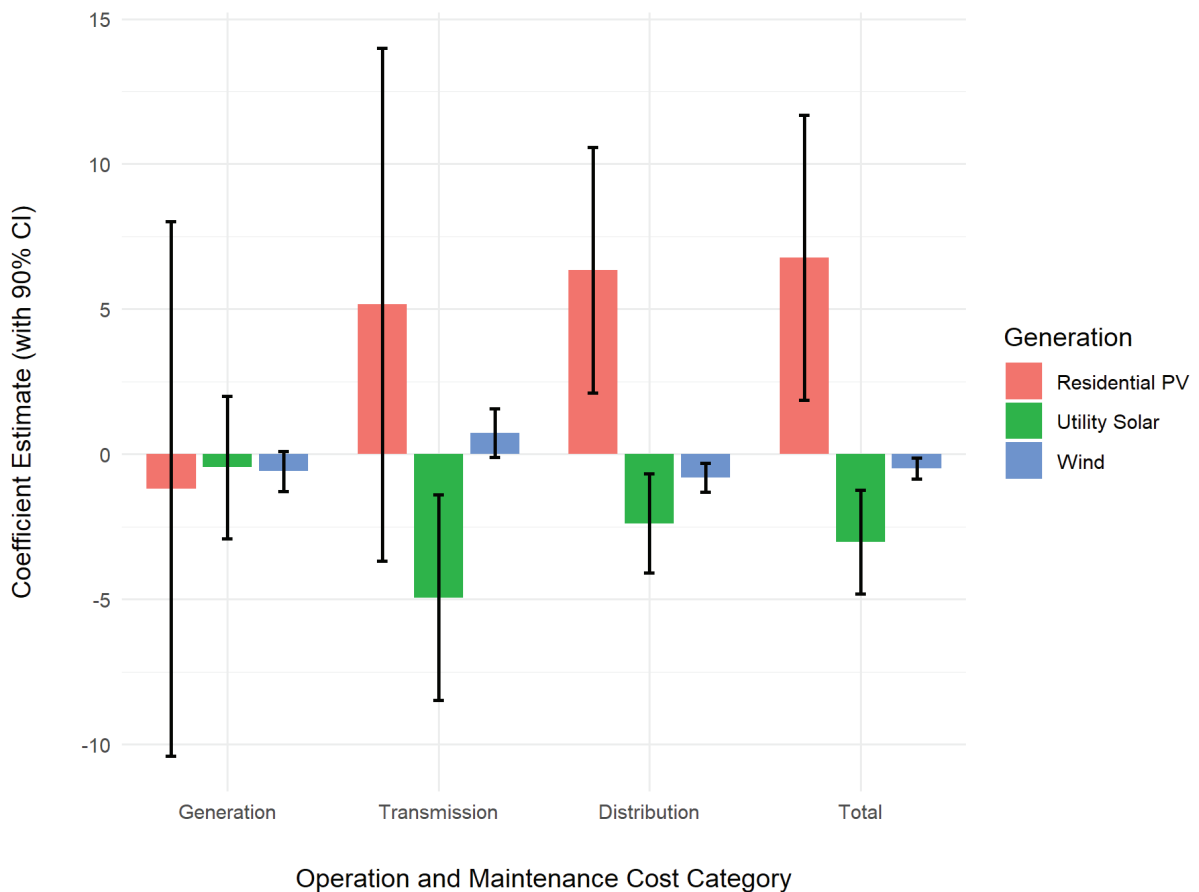


strain that exports place on the distribution network. These results affirm that the proper integration of renewables is an important aspect of a multifaceted solution that directly addresses household energy costs in the face of unprecedented demands on our infrastructure. Utilities are already investing heavily in grid upgrades to accommodate the impacts of climate-enhanced extreme weather events (such as heatwaves and wildfires) along with increased data center demand. Policymakers can secure an affordable energy future by improving the integration of low-cost renewables while implementing more equitable cost recovery strategies. ■■



Argosino, F. J. E. and Knittel, C. R. (2025), "Renewables and Electricity Affordability: Untangling Correlation from Causation," MIT CEEPR Working Paper 2025-22, December 2025.


**Figure 2. RPS and Generation Impacts on Household Electricity Prices.** This is a coefficient plot of a fixed-effects regression that estimates the impact of RPSs and generation shares on residential electricity prices. These results indicate that higher generation shares of utility-scale solar correlate with lower residential electricity prices, whereas higher shares of rooftop solar correlate with higher prices.



**Figure 3. Generation Technology Impacts on Utility O&M Costs.**

This coefficient plot displays three separate fixed-effects regression models that compare the impacts of different generation technologies on self-reported utility spending on generation, transmission, and distribution O&M costs. Residential PV (or solar) generation is consistently correlated with higher distribution and total costs for utilities whereas utility-scale solar and wind are associated with delivery and total costs.

Allroggen, F., Hansman, R. J., Knittel, C. R., Li, J., Wan, X., and Wang, J. (2026), "Planes Overhead: How Airplane Noise Impacts Home Values," MIT CEEPR Working Paper 2026-06, March 2026.



## Research.

# Planes Overhead: How Airplane Noise Impacts Home Values

By: Florian Allroggen, R. John Hansman,  
Christopher R. Knittel, Jing Li, Xibo Wan,  
and Juju Wang

Utilizing a robust quasi-experimental design, this study analyzes property value fluctuations across three major U.S. transportation hubs: Boston Logan International (BOS), Chicago O'Hare International (ORD), and Seattle-Tacoma International (SEA). By exploiting exogenous shifts in flight path architecture, the research provides a precise econometric foundation for understanding key components of the social costs of aviation noise. A rigorous examination of these impacts requires an analytical framework capable of isolating noise from other market variables through high-resolution data.

### Methodology: Exploiting Changes in Flight Path Architecture

Determining the causal impact of noise on property values often suffers from "self-selection bias," where noise-tolerant residents move into affected areas. To overcome this, the study utilizes the implementation of Performance-Based Navigation (PBN) and major runway reconfigurations as quasi-random events. These shifts altered noise exposure in ways that residents could not have anticipated, providing an ideal environment for a hedonic price model to measure market reactions.

The research achieves a technical superiority over traditional studies by utilizing a spatial resolution of 0.25 nautical miles. This high-resolution approach, powered by the Aviation Environmental Design Tool (AEDT) and ASDE-X radar-based flight trajectory data, ensures that noise is measured as it is actually "flown," rather than relying on generalized contours. The study identifies two primary sources of noise variation:



- **Performance-Based Navigation (PBN):** Satellite-based GNSS technology transitioned flight paths from broad, dispersed corridors into highly concentrated, precise tracks. This effectively "funneled" noise over narrow, high-intensity corridors.
- **Runway Reconfigurations:** Infrastructure projects, most notably the O'Hare Modernization Program (OMP), reconfigured intersecting runways into parallel alignments, fundamentally shifting traffic flows over neighborhoods that previously experienced minimal exposure.

These precise methodological controls allow for a definitive evaluation of how specific decibel increases translate into financial losses for homeowners.

### Quantitative Results: The Cost of a Decibel

To bridge the gap between acoustics and economics, the study employs the **Day-Night Average Sound Level (DNL)**. As the primary regulatory metric, DNL calculates 24-hour average exposure but applies a 10-decibel penalty to nighttime flights (10:00 p.m. to 7:00 a.m.) to reflect the heightened social and biological costs of sleep disruption.

The central finding across all three markets is a significant and consistent depreciation in housing prices relative to noise intensity. This is expressed through the **Noise Depreciation Index (NDI)**, which measures the percentage reduction in a home's value for every one-decibel (dBA) increase in noise.

The data indicates that a 1-decibel increase in noise exposure is associated with a reduction in housing prices between 0.6% and 1.0%. This capitalization effect suggests that airplane noise is a stable economic deterrent across diverse urban geographies. However, while the market response is consistent, the underlying "willingness to pay" for quiet reveals significant human and demographic heterogeneity.

### Heterogeneity and Resident Preferences for Quietness

Economic "Revealed Preferences" suggest that the amount a resident is willing to sacrifice for a quieter environment is the truest measure of that amenity's value. To derive annualized costs from property transactions, the researchers applied a 5% User Cost of Housing metric. This provides a clear view of the **Marginal Willingness to Pay (MWTP)**—the annual price a typical household would pay for a 1 dBA decrease in noise:

1. **Seattle (SEA):** \$221.36 (Highest average MWTP, reflecting high regional property values and a strong preference for quiet).
2. **Boston (BOS):** \$152.15.
3. **Chicago (ORD):** \$104.45.

Critically, the research reveals that this "taste for quietness" is deeply linked to sociodemographic profiles. Analysis of census block groups shows that the preference parameter for quietness is **positively correlated with higher household income but negatively correlated with the percentage of the non-white population**. These findings indicate that noise exposure is not only an economic issue but one of environmental justice, as more affluent and predominantly white communities demonstrate a higher financial threshold for avoiding aviation noise. While these results are descriptive rather than causal, they highlight disparities that may be relevant for federal and local policy design.

### Policy Implications and Conclusion

The implementation of PBN procedures highlights a stark policy conflict. For the FAA, PBN offers systemic efficiencies: reduced fuel burn, lower carbon emissions, and increased airspace capacity. However, these regional benefits come at the expense of localized social costs. For example, the study shows that people in Boston and Seattle would be willing to pay to go back to the 2011 noise levels, while in Chicago, the average household would actually need about \$122 in compensation to return to 2011, because the changes there made things better for many residents.

These findings necessitate a shift in how aviation infrastructure is planned and mitigated. Policymakers must move beyond "generalized noise contours" and adopt as-flown trajectory data at high spatial resolutions to accurately assess the impact of flight path concentration. Furthermore, because the "taste for quietness" correlates with income and race, future mitigation strategies must account for these demographic disparities to ensure equitable outcomes.

Integrating housing market data into aviation planning is no longer a secondary consideration; it is a fundamental requirement for sustainable urban infrastructure. Future FAA environmental impact assessments must prioritize the quantification of localized capitalization effects to develop targeted, data-driven mitigation measures that reconcile the demand for global connectivity with the right to local quiet. ■

Region	Noise Depreciation Index (NDI)
<i>Boston (BOS)</i>	0.61%
<i>Chicago (ORD)</i>	0.60%
<i>Seattle (SEA)</i>	0.97%

Table 1. Regional Noise Depreciation Impacts.

## Research.

# The Value of Storage in Electricity Distribution: The Role of Markets

By: Dirk Lauinger, Deepjyoti Deka, and Sungho Shin

In the 1980s and 1990s, deregulation was en vogue. By breaking up established monopolies, the hope was to increase market efficiency by stimulating competition across various industries, including electricity. In the US, for example, the state of Massachusetts separated electricity generation from distribution and transmission, which were to be kept as regulated monopolies. The rationale for the separation was to increase competition in generation while avoiding any “cross-subsidization of competitive businesses from regulated businesses and discriminatory policies affecting access to distribution and transmission networks upon which all competitive suppliers depend” (Joskow, 2008, p. 12).

In the context of the separation between generation and distribution, storage investment and operation within electricity markets warrant a closer look. Storage has generally been considered generation, but can be considered distribution if it serves grid reliability or defers grid investments such as substation or line upgrades. In the latter case, utilities have been allowed to own and operate storage for distribution needs. Such needs are generally infrequent, especially if they arise from reliability during extreme events. Storage built to address grid reliability thus experiences low utilization, on the order of one discharge cycle per month in some cases (Orange and Rockland Utilities, 2024), and foregoes potential revenue from participating in electricity markets. Allowing for market participation would improve storage economics but risk incentivizing investments that go beyond serving local distribution needs, which would run contrary to the separation of distribution from generation.

Here, we study policy designs that allow for utility-owned storage to participate in wholesale electricity markets while constraining the level of storage investment to local distribution needs. To identify permissible storage investments, we design a model that calculates optimal degrees of investment in grid assets, such as substation and line upgrades, and non-grid resources, such as storage and backup generation. Crucially, this investment model incorporates legal market participation constraints, which limit the generation from non-grid resources to serving local distribution needs. Next, we quantify the economic gains from participation in arbitrage and capacity markets, and analyze if the gains lead to storage investments that go beyond meeting local distribution needs.

Our study is timely as US utility spending on distribution grids has increased by 45% from the year 2018 through 2023 and now stands at over \$50 billion annually (Aniti, 2024). In addition, some deregulated states, i.e., Maryland (2019, p. 2) and New York (2021, pp. 12–13), have recently allowed market participation, hoping to reduce costs and recognizing that utility-owned distribution-grid storage is unlikely to exert substantial market power as it constitutes less than 1% of total US electricity generation capacity. On the other side of the Atlantic, the European Commission (2023, § 5) recommends exploring the full flexibility potential of energy storage in distribution grids and entertains proposals to create a market for local services that would help alleviate distribution grid constraints (ACER, 2025, Art. 29, 34, 40, 41, 44).

Specifically, we address three research questions about the value of storage in distribution grids:

1. How to model market participation constraints in storage operation and investment planning;
2. How do profits from market participation compare to savings from reducing peak demand;
3. Would market participation generate storage investments that go beyond distribution needs, and if so, how can this be detected?

We model market participation constraints by limiting the supply from non-grid resources, such as storage, to the shortfall of grid capacity relative to electricity demand. We integrate these constraints into an optimization problem that determines investment decisions and mimics integrated resource planning by distribution grid companies. As the constraints are nonconvex, they increase the computational complexity of the problem, raising questions about tractability. In a Massachusetts case study, we find that problems with a 26 year horizon and hourly resolution can be solved within several hours on current compute servers.

The case study further reveals that arbitrage and capacity market participation each generate about 50% of the capital cost savings from reducing or deferring grid investments by increasing storage utilization, see Figures 1 and 2. We determine storage investment levels appropriate for serving local distribution needs by solving the planning problems with the market participation constraints. We find that under current technology costs, market participation does not generate any storage investment that go beyond distribution needs.

Battery storage can thus be used more efficiently by providing distribution grid services and participating in electricity markets. These efficiency gains could be realized by (i) allowing and incentivizing distribution companies to participate in electricity markets, (ii) incentivizing distribution companies to contract third-party storage that also participates in electricity markets, (iii) creating a market for distribution grid services, potentially relying on nodal pricing in distribution networks (Sotkiewicz and Vignolo, 2006), or a combination of these approaches. We now discuss how our findings inform each of these options.

Regarding market participation of storage owned by distribution companies, our model can be used to estimate the value of such a policy change. In the past, the prospect of reduced or deferred grid investments was large enough to cast aside fears about market distortion, and regulators allowed storage investment for distribution

needs. Market participation promising similar savings may prompt regulators to reconsider existing policies that restrict market participation for storage assets owned by distribution companies. There is a risk that allowing for market participation would incentivize distribution companies to invest in storage solely for serving the market. Under current market conditions, we find that such investments are not profitable. Even if they were profitable, our model could be used to audit proposed investments and limit them to levels appropriate for addressing distribution needs. Jurisdictions that have already allowed market participation, e.g., New York and Maryland, may find our model useful to audit proposed storage investments.

Regarding distribution companies contracting third-party storage, our

model can be used to inform contract terms, e.g., storage quantities, prices, and commitment durations. The substantial savings from deferring grid investments shown in our case study suggest that such contracts could be beneficial to both distribution companies and storage developers, which may make previously unprofitable storage projects viable.

A market for distribution grid services could in theory be organized similarly to existing capacity markets on the transmission level. However, each distribution grid would likely have far fewer sellers than a capacity market, which may create concerns about liquidity and market power. Our model could be used to inform the terms of the capacity auctions in such a market, e.g., price caps and product durations. ■




Figure 1. Savings from Storage Investment Relative to Total Costs Without Storage Investment.

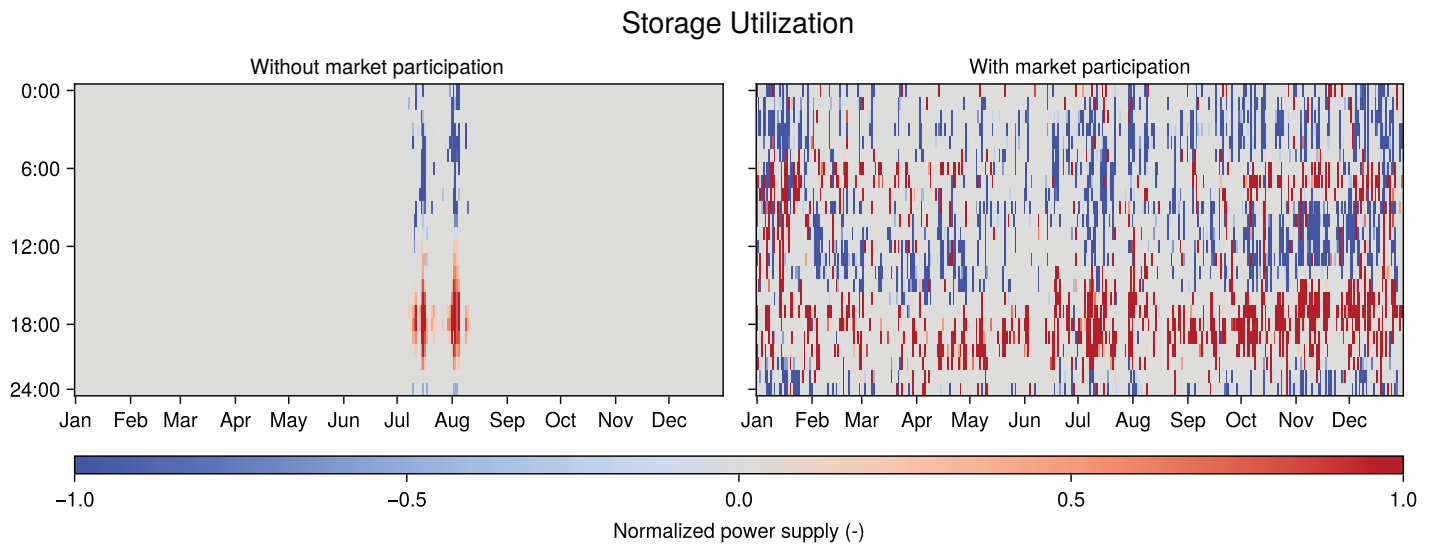


Figure 2. Storage Utilization With and Without Market Participation.



Lauinger, D., Deka, D., Shin, S. (2026), "The Value of Storage in Electricity Distribution: The Role of Markets," MIT CEEPR Working Paper 2026-07, April 2026. For references cited in this story, full bibliographical information can be found in the Working Paper.



Research.

# Cross-Border Solidarity Versus National Capacity Markets: Risk of Inadequate Capacity Procurement

By: Emma Menegatti and Leonardo Meeus

## Non-cooperative risk with national capacity markets

When capacity markets are operated nationally, each country might adopt a procurement strategy which maximizes its national welfare, rather than the regional welfare. Arises therefore a risk of non-cooperative behavior.

To evaluate this risk, we modeled a stylized electricity system of two interconnected countries, inspired by the approach of Lambin & Léautier (2019). Both countries are symmetrical, operate a common regional wholesale energy market, and a national capacity market. We verified whether the optimal regional capacity level is a stable Nash equilibrium. In other words, we checked that no country can increase its national welfare by unilaterally increasing or decreasing its capacity demand.

## Non-cooperative risk depends on how capacity market costs and benefits are allocated across the region

First, the national capacity market benefits—reduced likelihood of consumers curtailment—depend on how curtailment is allocated across

the two countries during common scarcity periods. We identified three possible allocation options:

- Curtailment-local, where domestic generation is allocated in priority to domestic consumers,
- Curtailment-reservation, where generation capacity (be it domestic or foreign) is allocated to the country which procured it
- And Curtailment-sharing, where curtailment ratios are equalized across the two countries. This is currently the prevailing rule, as implemented by default in the EU single day-ahead market coupling algorithm EUPHEMIA.

Second, the national capacity market costs depend on how capacity costs recovery is allocated across consumers of the two countries. We identify three possible allocation options:

Costs-local, where domestic consumers pay only for domestic capacity, Costs-reservation, where consumers pay for the capacity reserved (be it domestic or foreign). This is currently the prevailing rule, as the EU Electricity Regulation mandates Member States to allow for explicit cross-border capacity participation into national capacity markets. Costs-sharing, where domestic consumers pay for half of the regional capacity.

Combining these different curtailment and cost allocation options, we evaluated their compatibility, which is summarized in Table 1.

## The current EU framework is prone to free-riding

We found a risk of non-cooperative behavior with the current EU regulatory set-up. Sharing is the prevailing curtailment allocation rule in

the day-ahead market clearing algorithm EUPHEMIA. In case of common scarcity periods, curtailment ratios are equalized, as much as physically possible, across the two countries. Moreover, capacity markets costs are allocated at national level with either explicit cross-border participation (Cost-reservation), or implicit cross-border participation (Cost-local).

In such a configuration, the optimal capacity level is not a Nash equilibrium. Because the benefits (reduced consumers' curtailment) of capacity procurement are shared at regional level, both countries have an incentive to reduce their capacity demand to reduce their capacity costs. Injecting numerical values in the equilibrium condition, we find that we can expect capacity under procurement reaching around -1.4%. Though this relative change appears quite low, it increases the annual loss of load expectation by +54% and unserved energy by +60%.

Additionally, the capacity each country can access during scarcity does not match the capacity it procured. Therefore, EU countries don't have an incentive to procure capacity abroad, i.e. to implement explicit cross-border participation as mandated the EU Electricity Regulation.

**Not a solution: purely domestic approach**

An alternative approach would consist in letting countries a full responsibility over their national security of supply. Domestic consumers would only pay for domestic capacity (Costs-local). Additionally, domestic consumers would get the available domestic generation in priority during scarcity periods, which corresponds to minimizing cross-border trade (Curtailment-local).

The problem with such an approach is that it distorts capacity procurement towards domestic generation. If foreign generation is cheaper to build, then it would be more cost-efficient from the regional point of view to procure it abroad and import it during scarcity. The domestic approach would result in procuring generators in the wrong location, and higher capacity costs from the regional point of view.

**First solution: changing cross-border energy allocation during scarcity**

What if we maintain national capacity markets with explicit cross-border participation (Costs-reservation)? In such a case, the allocation of curtailment during scarcity periods should be based on the capacity procured. This ensures that each country gets exactly what it paid for, and eliminates the risk of non-cooperative behavior.

This solution might be difficult to implement in practice, considering time-dependent interconnectivity constraints, non-perfectly correlated scarcity events, and flow-based market coupling (see Elia, 2022). Additionally, "Curtailment sharing" in the day-ahead market clearing only takes place once supply bids are exhausted and market prices reach the common cap. Before that, cross-border flows would be directed towards the country in which consumers submitted most high price bids, which might not be the one which actually procured most capacity.

**Second solution: sharing fairly capacity costs ex-ante or ex-post**

What if we maintain the current day-ahead market clearing algorithm with "curtailment sharing"? Then, we should ensure that each country pays for the capacity that it gets. In other words, capacity costs should be fairly shared across the region. This can be done ex-ante, or ex-post.

First is ex-post. Capacity procurement and cost allocation could remain at national level. However, cross-border compensations could be allocated based on the observed cross-border flows during scarcity periods. If a country benefitted from more capacity than it paid for, then it would compensate its neighbors for the corresponding imports (i.e. the savings from reduced curtailment) it benefitted from.

Second is ex-ante. The capacity demand(s) could be defined cooperatively by the concerned countries through, for example, a regional capacity market. Costs could be recovered across the region proportionally to the expected benefits of the measure within each country. We discuss in Menegatti & Meeus (2025) why and how to implement regional capacity markets in the EU. As adequacy assessments and capacity auctions are conducted years ahead with a significant uncertainty, ex-post corrections could also be applied. ■



Menegatti, E. and Meeus, L. (2026), "Cross-Border Solidarity Versus National Capacity Markets: Risk of Inadequate Capacity Procurement," MIT CEEPR Working Paper 2026-08, April 2026

		Curtailment		
		Local	Reservation	Sharing
Costs	Local	Incompatible	N/A	Incompatible (Underprocurement)
	Reservation	N/A	Compatible	
	Sharing	Incompatible (overprocurement)		Compatible

Table 1. Compatibility of Capacity Costs and Curtailment Allocation Rules for the Achievement of the Optimal Capacity Level in a Non-Cooperative Scenario.

Shahabuddin, M. and Chyong, C. K. (2026), "Abatement or Reallocation? Unilateral Methane Standards in a Globally Traded Gas Market," MIT CEEPR Working Paper 2026-09, May 2026



## Research.

# Abatement or Reallocation? Unilateral Methane Standards in a Globally Traded Gas Market

By: Muntasir Shahabuddin and  
Chi Kong Chyong

Conventional fuel rankings typically place pipeline gas below LNG, below light fuel oil, with coal highest on a full lifecycle basis. Combustion emissions remain broadly stable across gas sources and continue to dominate total lifecycle emissions under a 100-year global warming

potential. However, once upstream heterogeneity is incorporated, the margin between fuels narrows substantially for high-intensity supply chains. Under median assumptions, export-oriented gas supply chains (e.g., North America) do not exceed coal in total from production to combustion. Under upper-bound methane intensities, however, some supply chains narrow the gap with coal and in the most extreme cases (e.g., Southeast Asia) close it entirely. The fuel hierarchy is therefore conditional on origin and asset performance rather than intrinsic to the fuel type. Treating gas as homogeneous obscures differences that are relevant to both climate accounting and regulatory design.

Upstream emissions are not evenly distributed across producers. In Russia, the top quintile of fields by methane intensity accounts for roughly one-quarter of regional upstream emissions, from around 12% of production. In Southeast Asia, the top quintile accounts for approximately 38% of emissions from 27% of output. In North America, dispersion is broader but still material: the top quintile represents around 10% of emissions from roughly 6% of production. This concentration has practical implications. Differentiated accountability can exert pressure on a relatively small subset of high-intensity assets without requiring disengagement from entire regions.

These findings are directly relevant to the EU Methane Regulation (Regulation (EU) 2024/1787), which introduces traceability of origin and producer identity, monitoring, reporting and verification equivalence, methane intensity disclosure from 2028 and the possibility of maximum intensity thresholds from 2030. The Regulation creates transparency and a screening mechanism at the point of import. Whether it generates global abatement or primarily reshapes trade flows depends on the scale of upstream heterogeneity, abatement costs relative to the discount required to place high-intensity gas in non-EU markets, the price elasticity of demand in those markets, and the EU's share of basin-level offtake, which together determine whether reshuffling is cheaper than mitigation.

The principal risk is trade reshuffling: higher-emissions gas displaced from EU trade is redirected to less-regulated destinations. In coal-dominated markets such as China and India, LNG competes mainly with thermal coal, and in the base carbon price scenario shown in the



box below, LNG remains materially above coal parity, meaning large-scale diversion would require deep discounts and sharply reduce producer revenues. Under prevailing 2024 cost and price assumptions, seamless redirection at scale therefore appears constrained. This conclusion weakens if climate policies in China and India become much more stringent: under the high carbon price scenario, rising coal costs close the gap with LNG, making diversion more plausible. The climate effect then depends on which supply chains are diverted. Under GWP100, gas emissions generally remain lower than coal emissions, so switching still reduces emissions on average. Under GWP20, methane carries greater weight, and the gap narrows for methane-intensive supply chains, strengthening the case for upstream-inclusive carbon pricing or for the EU-type methane regulation.

A separate, supply-side adjustment channel lies within exporting jurisdictions (e.g., North America). Where methane intensities vary significantly across assets, producers can route lower-intensity supply to regulated markets while absorbing higher-intensity production domestically. This portfolio sorting improves the average methane intensity of gas entering Europe without necessarily reducing aggregate production or basin-level methane emissions. Import intensity can fall even if global emissions remain broadly unchanged, provided high-intensity assets continue operating at capacity for non-EU markets. The distinction between asset-level abatement and portfolio reallocation is therefore central to evaluating the Regulation’s climate impact.

Shipment-level disclosure alone cannot resolve this distinction. Import statistics reveal what enters the EU market, but not how emissions evolve across producing regions as a whole. If EU import intensity declines while methane intensity across exporting basins remains unchanged, the evidence points to supply rerouting rather than genuine emissions reduction. To address this issue, we propose a Commission-led monitoring function that publishes production-weighted methane intensity and absolute emissions for major exporting regions supplying the EU, drawing on producer disclosures, production data, and satellite

verification. This would serve an evaluative rather than a compliance purpose. It will provide policymakers with the evidence base to assess whether the Regulation is driving real mitigation or portfolio reallocation, and to progressively tighten intensity thresholds where rerouting is observed.

For policymakers, the results suggest that unilateral methane regulation does not mechanically induce leakage, but neither does it automatically deliver global abatement. The effectiveness of demand-side standards depends on upstream heterogeneity, abatement costs, the price elasticity of alternative destination markets, and the EU’s share of exporting basin offtake. Carbon pricing amplifies lifecycle cost differentials between high- and low-intensity supply chains, thereby strengthening the commercial case for abatement over relocation. Contract structures and domestic demand conditions in exporting jurisdictions will further shape adjustment pathways.

For producers and traders, methane intensity is increasingly a competitive parameter within fuel types. Asset-level performance and portfolio composition affect market access under emerging screening regimes. The concentration of emissions in a limited subset of high-intensity fields implies that targeted mitigation investments can materially shift relative positioning. When emissions are concentrated in a few high-intensity fields, relatively small operational changes at those fields can significantly lower a producer’s overall emissions.

The design of methane accountability frameworks needs to anticipate portfolio responses and build in mechanisms to verify whether reductions in import intensity reflect genuine mitigation or commercial reallocation. Differences in upstream methane intensity across supply chains are large enough to warrant differentiated accountability rather than average-based treatment. The central question is not simply whether methane standards alter trade flows, but whether they change the economics of high-intensity assets. That answer determines whether unilateral methane regulation results in abatement or reallocation. ■■

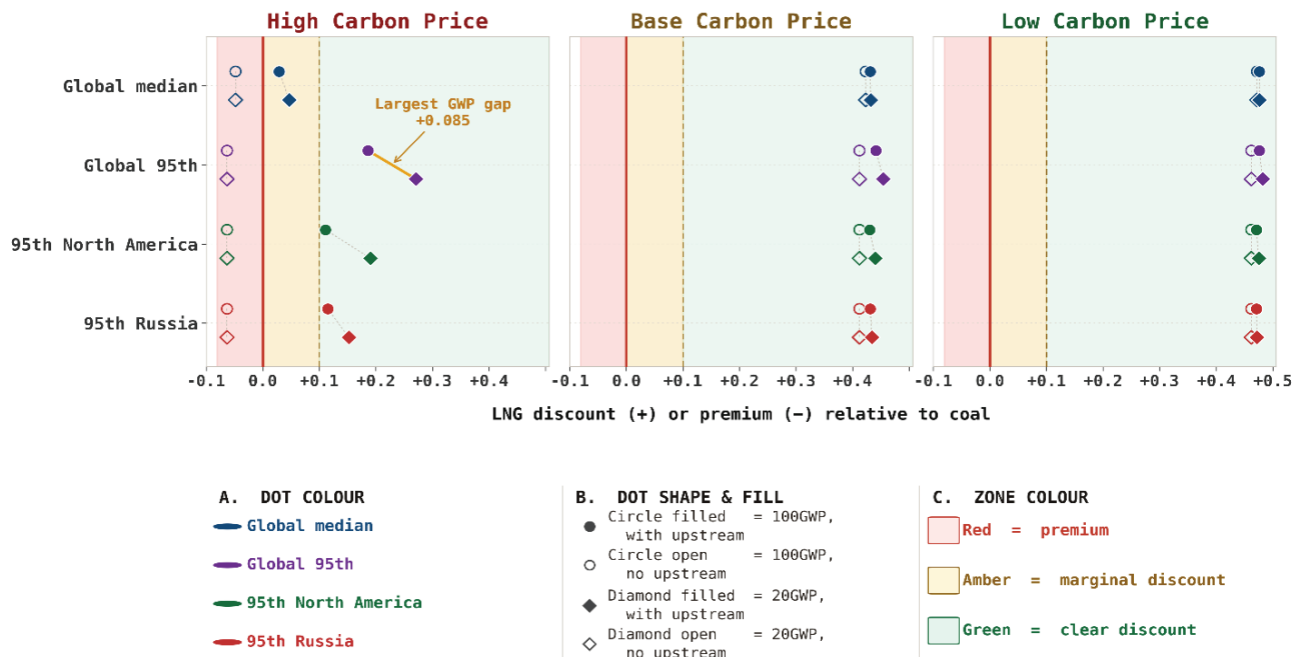


Figure 1. LNG-Coal Breakeven Under Carbon Price and Methane Accounting Scenarios for China and India.

## Research.

# Steady Heat, Flexible Power: Flexibility by Proxy in Nuclear-Thermal Energy Storage Systems

By: Nicolas Trosino, John E. Parsons, Ruaridh Macdonald, Charles Forsberg, and Yannick Perez

Thermal Energy Storage (TES) is a promising technology to reconcile the economic constraints of capital-intensive low-carbon thermal generators such as nuclear power plants with the operational flexibility required in deeply decarbonized systems. The increasing penetration of variable renewable energy drives a growing spread in the marginal value of generation across hours through a day and throughout the year. On the one hand, when renewables are plentiful and net load is small or even zero, the marginal value of generation is zero. On the other hand, when there is a drought of renewable energy and net load threatens to exceed firm capacity, the marginal value of generation spikes. Integrating Thermal Energy Storage enables baseload production of heat, even during hours when the marginal value of electricity is low, which can then be converted to electricity when the marginal value of generation is high. This creates a new operational paradigm in which nuclear assets can behave simultaneously as firm baseload providers to maximize rate-of-return and as flexible resources to match system needs.

We employ the GenX Capacity Expansion and Dispatch Model to assess the system value of integrating Thermal Energy Storage into a fusion power plant. We parameterize our GenX model to match a 2050, deeply decarbonized New England grid, and we use GenX to co-optimize investment and hourly operation across a portfolio of competing low-carbon technologies. Our results show that Thermal Energy Storage fundamentally reshapes the economic frontier for fusion, dramatically increasing the breakeven cost at which investment is cost-effective for the grid, from 8,000 \$/kW<sub>e</sub> up to 19,000 \$/kW<sub>e</sub>. Incorporating TES and therefore deploying fusion compresses the spread in the marginal value of generation across hours of the day and through seasons of the year. Our model optimizes the size of the TES energy reservoir as well as the size of the turbine generating electricity. The optimized TES includes a very large energy reservoir, giving the TES the ability to deliver power for a long duration of 10 hours. We analyze the operation of the TES and contrast it with the operation of other storage technologies.

We use the GenX model described in Armstrong et al. 2024. Generating technologies available in the model include utility-scale PV, commercial rooftop PV, residential PV, onshore wind, fixed offshore wind, floating offshore wind, legacy run-of-river hydro, natural gas combined-cycle, and natural gas combined-cycle outfitted with carbon capture and sequestration. Storage technologies include legacy pumped hydro-facilities as well as Li-ion batteries. Notably, since New England imports a significant amount of power from Quebec's large hydro system, the model also includes an hourly load profile in Quebec, together with an hourly water inflow profile, and the model optimizes the exports of power to New England while respecting specified reservoir system constraints such as minflow and seasonal maximum capacities. A distinctive feature of that model is its parameterization with 20 scenario years for hourly load and wind and solar resource availability, so that the chosen portfolio of investments must be robust across the range of scenario years. In this paper, the carbon limit is 12 gCO<sub>2</sub>/kWh<sub>e</sub>, a 95% reduction relative to 1990.

Alongside the other generation technologies, Armstrong et al. 2024 introduce a tokamak-type fusion power plant. The representation of a




Figure 1. Median Weekly Heat and Electricity Production Through an Average Year.

thermal power plant in GenX involves only a few essential parameters, such as thermal efficiency and operational constraints, and does not require elaboration of the many details particular to this or that fusion plant. Nevertheless, to fix ideas, we briefly describe a few selected features of a generic, imagined deuterium-tritium magnetic confinement fusion power plant. The base case plant's total thermal power capacity to 1,095 MW<sub>th</sub>. Turbine efficiency is 40 %, in which case the plant's gross electric power capacity is 438 MW<sub>e</sub>. Some 111 MW<sub>e</sub> of power is needed for station load within the plant to cool the magnets, drive the fusion reactor, pump the molten blanket, and supply electric power for the tritium processing system and other plant operations. Thus, the net electric capacity is 327 MW<sub>e</sub>. For our modeling, it will be important to understand how the station load varies as the thermal power is ramped down. Besides a permanent station load of 10 MWe, we assume FixedSL = 10 MW<sub>e</sub> is the fixed load and VarSL = 0.083 MW<sub>e</sub>/MW<sub>th</sub> is the variable load. The Fusion Power Plant can balance its self-consumption by collecting part of its electricity generation or by consuming electricity from the grid. Overall, the plant's heat-to-electricity efficiency is  $\eta_{\text{FPP}} = 29.4\%$ .

We model the TES in a direct architecture in which the fusion plant provides thermal energy to the salt loop which independently supplies

the turbine loop. The TES charges when the plant provides more energy to the salt loop than the salt loop discharges to the turbine, and excess salt is transferred from the cold tank (light blue) to the hot tank (light red). Conversely, the TES discharges when the salt loop discharges more energy to the turbine than it receives from the plant, transferring salt from the hot tank to the cold tank and boosting generation at the turbine.

Table 1 shows the optimized portfolio and dispatch, first assuming the fusion plant is only available without TES, and second assuming TES is integrated into the fusion power plant. The assumed cost of the base fusion plant is \$8,500/kW<sub>e</sub>. At that cost, and without TES, the optimized portfolio does not include any investment in fusion capacity. However, when TES is integrated into the fusion plant, the optimized portfolio does include investment in fusion capacity—7.7 GW<sub>e</sub> of turbine capacity generating 17.46 TWh of electricity.

Our model optimally sizes both the power and the energy capacity of the TES, and we find that it chooses very long durations. The TES has a lower turnover than grid batteries, but a higher utilization. The TES captures a higher spread on its turnovers than grid batteries, and its operating profitability is more concentrated in a fewer number of hours. ■

	Without TES			With TES		
	Capacity (GW <sub>e</sub> )	Generation (TWh <sub>e</sub> )	Capacity Factor	Capacity (GW <sub>e</sub> )	Generation (TWh <sub>e</sub> )	Capacity Factor
Natural gas	8.9	2.21	3%	9.8	2.83	3%
Natural gas w/ CCS	23.1	50.83	25%	19.4	46.11	27%
Solar PV	22.0	27.59	14%	22.0	27.7	14%
Onshore wind	9.0	30.27	38%	9.0	30.47	39%
Offshore wind (fixed)	27.3	114.79	48%	23.3	98.37	48%
Run-of-river hydro	2.0	7.78	45%	2.0	7.78	44%
Fusion	0.0	0.00		7.7	17.46	26%
Pumped hydro	1.8	2.61		1.8	1.97	
Battery	18.6	25.48		12.5	13.55	
Total	112.7	261.56		107.5	246.24	
Net imports	2.3	13.15		2.3	13.59	
Total generation		233.48			230.72	
Storage charge		-33.32			-18.44	
Storage discharge		28.09			15.52	
Net Energy supplied		241.39			241.39	

Table 1. Optimized Capacity Portfolios and Dispatch, Without and With TES Integrated with the Fusion Power Plant.



Trosino, N., Parsons, J. E., Macdonald, R., Forsberg, C., & Perez, Y. (2026), "Steady Heat, Flexible Power: Flexibility by Proxy in Nuclear-Thermal Energy Storage Systems," MIT CEEPR Working Paper 2026-10, May 2026



## Research.

# Political Ideology and U.S. Electric Vehicle Adoption

By: Lucas W. Davis, Jing Li, and  
Katalin Springel



Davis, L., Li, J., Springel, K. (2026), "Political Ideology and U.S. Electric Vehicle Adoption," MIT CEEPR Working Paper 2026-11, May 2026

This paper examines the relationship between political ideology and electric vehicle (EV) adoption in the United States, using detailed county-level data on new vehicle registrations in the U.S. from 2012 to 2023. Our data reveal a strong, persistent correlation between Democratic voting patterns and higher rates of EV adoption, raising important implications for climate policy and technology diffusion.

We begin by examining the new vehicle registration data at the state level from 2021 to 2023, when electric trucks and vans were brought to the market in meaningful numbers. Figure 1 shows the correlation between political ideology and EV adoption, separately by vehicle type. The x-axis is the Democrat vote share, and the y-axis is the share of that vehicle type that is electric. For example, in California, 23% of all new SUVs were EVs. Trucks are more popular in Republican-leaning areas, suggesting that introducing electric versions of trucks could help broaden EV adoption across the political spectrum. However, the positive association between political ideology and EV adoption is pronounced for all four vehicle types.

Turning to the data at the county level, we find that about half of all new EVs in the U.S. were registered in the 10% most Democratic counties, and about one-third went to the top 5%. These patterns remain stable over time, even as the EV market expanded dramatically in the range of offerings from manufacturers and in total units sold. Top U.S. counties for EVs are urban, high-income, and in Democratic states, such as California's Bay Area, King County, WA, Multnomah County, OR, and Middlesex County, MA. By the end of our sample period, in 2023, 45% of EVs still went to the 10% most Democratic counties, and 32% still

went to the top 5%. Overall, we find limited diffusion of EV adoption across the political spectrum.

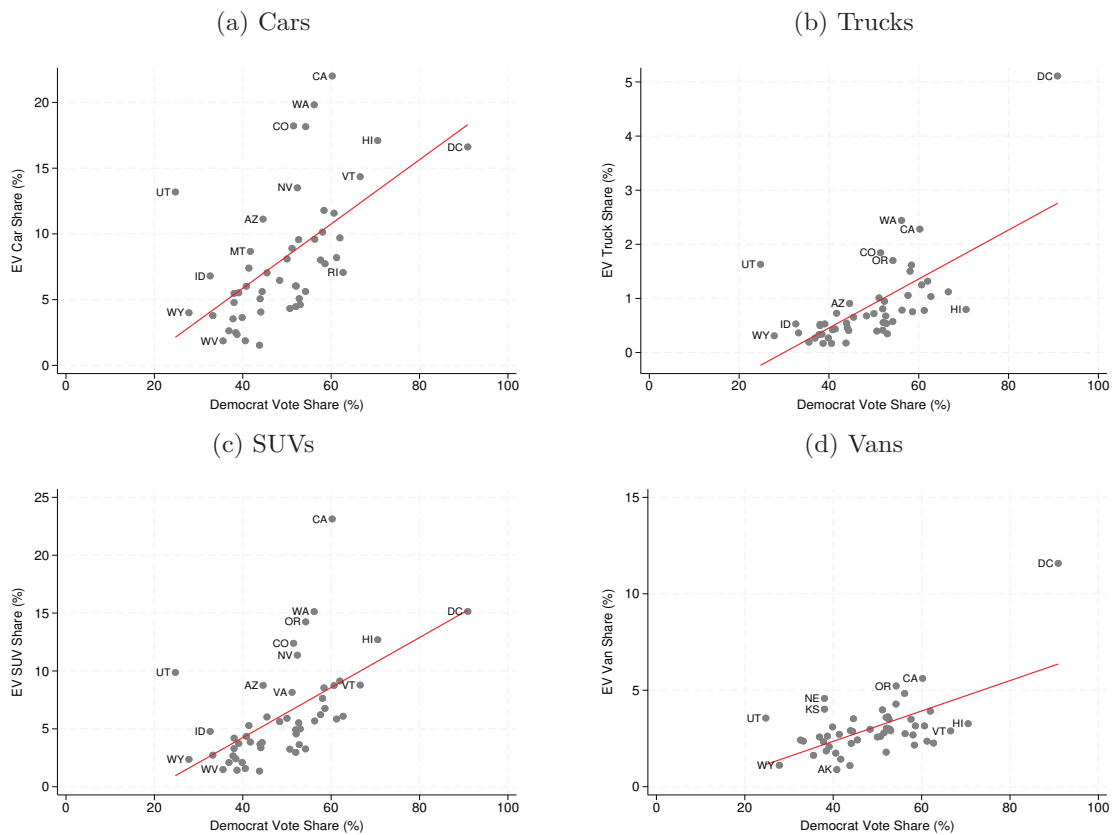
The association between EV adoption and political affiliation is descriptive rather than causal. We cannot randomly assign political ideology to examine its impact on EV adoption. However, we can use regression analysis to control for other key economic and demographic factors, including household income, population density, and gasoline prices. We find that political affiliation is an economically and statistically meaningful determinant of EV adoption even after controlling for these other potential explanatory factors: ...a one percentage point increase in Democrat vote share (e.g., going from 45% to 46%) is associated with a 0.026 percentage point increase in EV adoption (e.g., going from 1.0 percent to 1.026 percent). Mean EV share is less than 1 percent, so this is a large effect.

We then conduct a survey to examine whether a gap in knowledge about EVs can explain the differences in EV adoption between Democrats and Republicans. Our survey consisted of three parts. In the first part, respondents were shown images of vehicles and asked to classify them as EVs or non-EVs. The second part of the survey contained questions about EV knowledge, including charging and batteries. The third and last part of the survey asked about general automotive knowledge. Our survey revealed minimal differences between

Democrats and Republicans in their ability to recognize EVs. While Democrats demonstrate slightly higher understanding of EV charging, they scored slightly worse on general automotive knowledge.

Our paper contributes to a broader literature on “green” consumption, which has identified two main mechanisms for how political ideology may factor into households’ consumption decisions. First, intrinsic motivations, such as environmental preferences or “warm glow” from reducing emissions, may differ across political groups. Second, extrinsic motivations, such as social signaling, may play a role. EV ownership could signal environmental identity in Democratic-leaning communities but carry less social value in Republican-leaning areas.

In this paper, we find that political affiliation is an independent and meaningful factor in U.S. households’ EV adoption decisions, even after controlling for other demographic and socio-economic variables. Given the persistence of the association between political affiliation and EV adoption from 2012 to 2023, political polarization may be a critical challenge for scaling green technology adoption. We emphasize that our findings are purely correlational and descriptive. To understand how to design policies that can bridge the partisan divide in green technology adoption, future work can delve further into the underlying motivations for EV adoption and how those motivations are determined by political affiliation. ■



**Figure 1. Political Ideology and EV Adoption by Vehicle Type.**

Notes: Each scatterplot has 51 observations, one for each state and one for Washington, DC. The x-axis is the share of voters in the 2012 U.S. presidential election who voted for Barack Obama. The y-axis is EV cars, trucks, SUVs, and vans as a share of all new vehicles registered in that vehicle type from 2021 to 2023, ranging from 0% to 100%. We also plot least squares linear regression lines.



On April 22 in Chicago, Joshua Hodge, CEEPR's Executive Director, presented to the Board of Directors of Midcontinent Independent System Operator (MISO) as part of a closed strategic planning session. The discussion touched on the integration of large new electricity loads, including data centers, as well as future transmission planning and cost allocation considerations.



## Events.

### Recent and Upcoming Conferences:



Information on these events is available on our website, where Associates can also access presentation slides and recordings: [ceepr.mit.edu/events](https://ceepr.mit.edu/events)

#### **2026 Spring MIT CEEPR Research Workshop**

**May 20-21, 2026**  
MIT Samberg Conference Center  
Cambridge, Massachusetts

#### **2026 CEEPR-FSR Conference on the Future of Energy and Climate Policy**

**October 7-8, 2026**  
Florence, Italy

*in partnership with the  
Florence School of Regulation.*

#### **2026 Fall MIT CEEPR Research Workshop**

**December 10-11, 2026**  
Hotel Washington  
Washington, D.C.

# Publications.

## Recent Working Papers:

### WP-2026-11

#### **Political Ideology and U.S. Electric Vehicle Adoption**

Lucas W. Davis, Jing Li, and Katalin Springel, May 2026

### WP-2026-10

#### **Steady Heat, Flexible Power: Flexibility by Proxy in Nuclear-Thermal Energy Storage Systems**

Nicolas Trosino, John E. Parsons, Ruairidh Macdonald, Charles Forsberg, Yannick Perez, May 2026

### WP-2026-09

#### **Abatement or Reallocation? Unilateral Methane Standards in a Globally Traded Gas Market**

Muntasir Shahabuddin and Chi Kong Chyong, May 2026

### WP-2026-08

#### **Cross-Border Solidarity Versus National Capacity Markets: Risk of Inadequate Capacity Procurement**

Emma Menegatti and Leonardo Meeus, April 2026

### RC-2026-03

#### **Commentary: The Economics of Continued Operation of the Diablo Canyon Nuclear Power Plant, 2030-2045**

John E. Parsons, April 2026

### WP-2025-07

#### **The Value of Storage in Electricity Distribution: The Role of Markets**

Dirk Lauinger, Deepjyoti Deka, and Sungho Shin, April 2026

### RC-2026-02

#### **Commentary: Concentrating Risk Without Promoting Resilience: What the MP Materials Deal Means for U.S. Critical Mineral Strategy**

Brian Deese, Zehra Khan, and Robert Reese, April 2026

### WP-2026-06

#### **Planes Overhead: How Airplane Noise Impacts Home Values**

Florian Allroggen, R. John Hansman, Christopher R. Knittel, Jing Li, Xibo Wan, and Juju Wang, April 2026

### WP-2026-05

#### **Designing Second-Best Price Zones in Electricity Markets**

Jonas Boeschemeier and Sebastian Schwenen, March 2026

### WP-2026-04

#### **From Least-Cost Plans to Implementable Electrification: A Regulatory-Financial Framework to Achieve Universal Access**

Santos J. Díaz-Pastor, Carlos de Abajo, Robert Stoner, and Ignacio J. Pérez-Arriaga, March 2026

### WP-2026-03

#### **How Uncertainty Shapes Electricity Storage Decisions: Dispatch Policies and Capacity Portfolios Under Renewable Drought Risk**

Cem Keske, Bjarne Steffen, and John E. Parsons, February 2026

### WP-2026-02

#### **Who Bears the Burden of Climate Inaction?**

Kimberly A. Clausing, Christopher R. Knittel, and Catherine Wolfram, February 2026

### RC-2026-01

#### **Commentary: Lessons for Scope 2 Greenhouse Gas Emissions Accounting from the Debate Over Non-GAAP Earnings and the SEC's Regulation G**

John E. Parsons, February 2026

### WP-2026-01

#### **The Differential Impacts of Critical Mineral Prices and Oil Prices on the Economy**

Adrien Concordel, Phuong Ho, and Christopher R. Knittel, January 2026

### WP-2025-22

#### **Renewables and Electricity Affordability: Untangling Correlation from Causation**

Fischer J. Espiritu Argosino and Christopher R. Knittel, December 2025

### WP-2025-21

#### **Beyond the Early Era of EVs: Evidence from the Staggered Rollout of the HOV Lane Network in California**

Eric Avis, Jing Li, and Katalin Springel, November 2025



All listed working papers in this newsletter are available on our website at: [cepr.link/workingpapers](https://cepr.link/workingpapers)



**MIT CEEPR**  
Center for Energy and  
Environmental Policy Research

**MIT Center for Energy and  
Environmental Policy Research**  
Massachusetts Institute of Technology  
77 Massachusetts Avenue, E19-411  
Cambridge, MA 02139-4307  
USA

[ceepr.mit.edu](http://ceepr.mit.edu)



MIT CEEPR Graduate Research Assistants and friends from MIT CS3 and MIT TPP at the Charles River on an early morning walk before Spring 2026 finals week.

*Pictured from left to right starting from the top row:*

*Jaclyn Rambarran, Fischer Argosino, Kevin Lin Yang, Rohil Agarwal, Lance Pangilinan, Eli Duggan, Grant Lee,  
Khyati Garg, Daniel Xue, Julia Lukens, Reese Dobson, and Shreeya Parekh.*



MASSACHUSETTS INSTITUTE OF TECHNOLOGY