MIT CEEPR
Center for Energy and Environmental Policy Research

Research Commentary

Smart Meters, Dumb Prices:
Why Electricity Tariffs Are Becoming (More) Inefficient, Regressive, and Outdated

Christopher R. Knittel

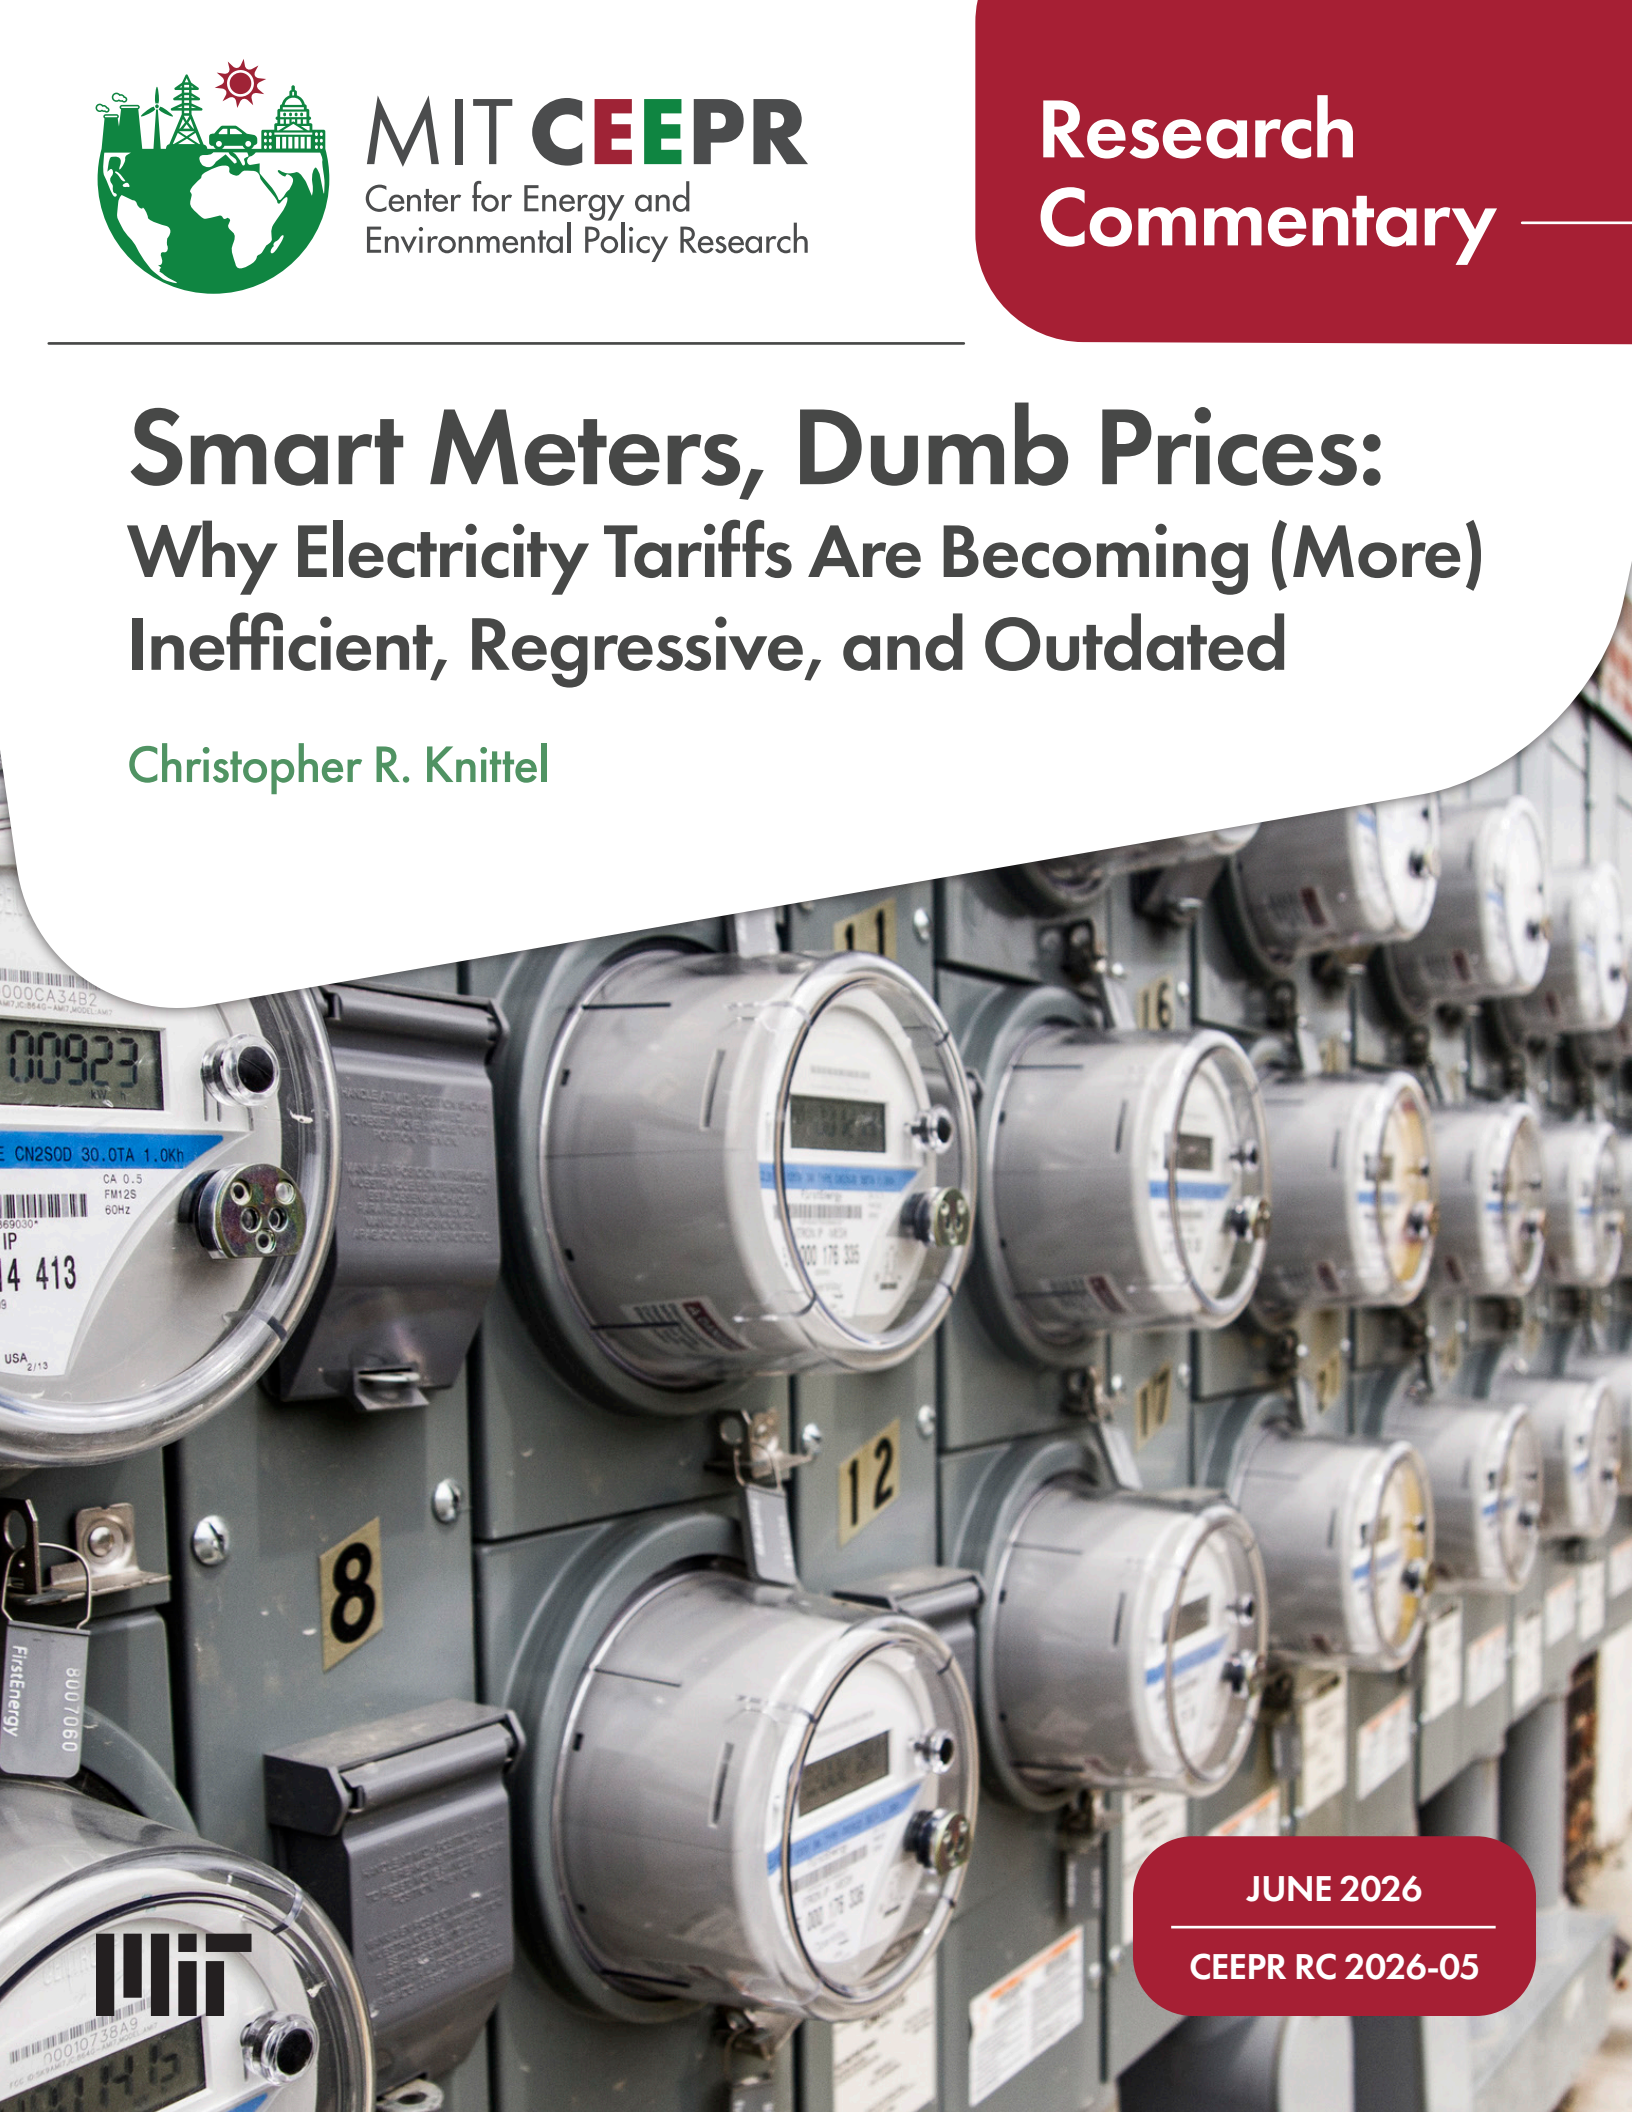

JUNE 2026

CEEPR RC 2026-05

Research Commentary Series.

Since 1977, the Center for Energy and Environmental Policy Research (CEEPR) has been a focal point for research on energy and environmental policy at MIT. CEEPR promotes rigorous, objective research for improved decision making in government and the private sector, and secures the relevance of its work through close cooperation with industry partners from around the globe. Drawing on the unparalleled resources available at MIT, affiliated faculty and research staff as well as international research associates contribute to the empirical study of a wide range of policy issues related to energy supply, energy demand, and the environment.

As a Research Commentary, views expressed within are solely those of the authors. Interested parties may contact the authors directly for further feedback on their Commentaries.

Smart Meters, Dumb Prices:

Why Electricity Tariffs Are Becoming (More) Inefficient, Regressive, and Outdated

Christopher R. Knittel*

June 10, 2026

Abstract

Electricity tariffs were designed for an earlier power system: one with electromechanical meters, one-way power flows, limited customer-side generation, and relatively weak substitutes for grid electricity. That world is changing. Smart meters make time-varying prices feasible. Rooftop solar, batteries, and behind-the-meter generation allow customers to reduce measured net consumption while remaining connected to the grid. Electrification makes electricity prices central to decarbonizing transportation, buildings, and industry. Data centers are creating large, concentrated new loads and are increasingly considering on-site generation and special grid-service arrangements. These changes expose two long-standing weaknesses in retail electricity pricing. First, many costs of the electricity system are fixed, residual, or driven by peak capacity needs, yet they are commonly recovered through volumetric charges on kilowatt-hours. Second, most residential and small commercial customers face one retail energy price even though the cost of serving electricity varies across hours, locations, and system conditions. In short, we have smart meters, but too often still have dumb prices.

This primer focuses primarily on retail tariff design in the United States, though it draws on international experience where useful. It explains why that tariff structure is becoming increasingly inefficient, regressive, and outdated. It begins by unpacking the electricity bill and distinguishing energy costs, peak-driven capacity and network costs, and residual costs. It then describes the economically efficient tariff: time-varying prices for energy, coincident peak charges for costs driven by system peaks, and transparent fixed or connection charges for residual costs. The primer explains why current tariffs looked reasonable historically, why rooftop solar and data centers now put pressure on volumetric cost recovery, why flat energy prices are increasingly poorly suited to renewable-heavy power systems, and how other jurisdictions are experimenting with dynamic prices, net billing, two-way tariffs, and cost-reflective network charges. It concludes with principles for reform: charge customers for the costs they impose, pay them for the value they provide, and address affordability explicitly rather than hiding redistribution in distorted kilowatt-hour prices.

* *Christopher R. Knittel*: George P. Shultz Professor of Energy Economics and Associate Dean for Climate and Sustainability, MIT Sloan School of Management; Director, MIT Center for Energy and Environmental Policy Research; Director, MIT Climate Policy Center; and NBER; knittel@mit.edu. I thank Severin Borenstein, Meredith Fowlie, Thad Moody, and Graham Turk for helpful comments.

Executive Summary

Electricity tariffs were designed for an earlier power system: electromechanical meters, one-way power flows, limited customer-side generation, and customers with little ability to respond to prices in real time. In that world, recovering most costs through cents-per-kilowatt-hour charges was simple, administratively feasible, and often distributionally attractive. That world is changing. Smart meters can measure when electricity is used. Rooftop solar and behind-the-meter generation let customers reduce net purchases while still relying on the grid. Electric vehicles, heat pumps, batteries, and data centers are reshaping when and where load appears. Renewable generation makes the cost of serving electricity vary sharply across hours. We increasingly have smart meters, dumb prices.

The diagnosis of this primer is that a cluster of problems usually treated as separate share a single structural cause. Transmission, distribution, policy, and residual costs are recovered through a volumetric kilowatt-hour charge that is also asked to signal the marginal cost of energy. That single number cannot do both jobs. It is too high to encourage efficient electrification in many hours, too low to reflect scarcity in others, too easy to avoid through behind-the-meter generation, and too poorly targeted to do equity work efficiently. The cost shift from rooftop solar, the implicit tax on electrification, the proliferation of appliance-specific rates, the upward pressure on rates from new load, and a substantial share of the regressivity visible in residential bills are downstream consequences of overloading one price. Many of the second-best instruments now layered onto retail tariffs exist to patch those consequences.

The efficient tariff has three components. Energy should be priced based on what it costs to serve electricity at that time and place. Costs driven by system peaks or local network constraints should be recovered through coincident peak or capacity charges. Residual costs that do not vary with short-run consumption should be recovered transparently—through fixed charges for residential customers, capacity-based instruments for large customers, or funding outside the electricity bill where the underlying purpose is broader than electricity service.

Two cases illustrate why the current structure is breaking down. Rooftop solar under traditional net metering avoids paying the full bundled retail rate, including grid costs that do not disappear when the customer installs solar; because adoption is positively correlated with income, this shifts costs onto customers who cannot install solar. Data centers face a different version of the same problem: standard demand charges were calibrated to a prior generation of large customers, do not capture the transmission and generation costs hyperscale loads trigger, and do not price the standby value of grid connection for facilities that self-generate during peaks. In both cases, net consumption is not the same as grid reliance.

Equity concerns about moving residual costs out of volumetric rates are real, but the answer is not to preserve inefficient prices indefinitely. It is to address distribution explicitly—through income-graduated fixed charges, low-income discounts, or tax-funded support—rather than hide redistribution in distorted energy prices. The grid is not becoming less important. It is becoming more important and harder to pay for correctly. A modern electricity tariff should charge customers for what they use, pay them for what they provide, and fund shared infrastructure transparently. Today's tariffs are not crazy. They are obsolete.

1 Introduction: Why Electricity Prices Are Sending the Wrong Signals

The US electricity tariff system was built for a different era. It was built for a world of slow meters, limited behind-the-meter generation, and customers who had little ability to respond to prices in real time. That world is fading. Smart meters can now record when electricity is used. Rooftop solar, batteries, electric vehicles, and controllable appliances can change how and when households draw power. Data centers are becoming large new loads. The electricity system is no longer just a one-way pipe that delivers anonymous kilowatt-hours to passive customers.

That shift creates a tariff problem. The price customers see on their bill often does not match the cost their electricity use creates for the system. In many places, the bill folds together costs that vary hour by hour with costs that do not. It also gives customers the same energy price whether electricity is scarce or abundant, even though the system value of electricity can differ sharply across the day.

What this primer adds to a now-substantial literature on electricity tariff design is a single diagnostic claim: a cluster of problems usually treated as separate—the cost shift from rooftop solar, the implicit tax on electrification, the proliferation of appliance-specific rates, the upward pressure on rates from new load, and a substantial share of the regressivity visible in residential bills—are not independent. They are downstream consequences of one structural choice: recovering residual transmission, distribution, and policy costs through a volumetric kilowatt-hour charge that is also asked to signal the marginal cost of energy. Once that single price is asked to do several incompatible jobs, it does each of them badly, and a long list of second-best instruments accumulates to patch the consequences. Recognizing the common cause changes which reforms matter most: fix the underlying prices, and a substantial share of the patches become smaller, narrower in scope, or unnecessary.

This primer explains why. It argues that electricity prices should be built around three distinct ideas: charge for the energy when and where it is used, charge for the capacity and network service a customer depends on, and recover residual costs transparently instead of hiding them in the kilowatt-hour price. That sounds technical, but the practical issue is simple: the old tariff structure is increasingly making electricity look more expensive than it really is when served, while at the same time letting some customers avoid paying for the grid they still rely on.

2 The Bill Is a Bundle

A retail electricity bill is not just a bill for electricity (Burger et al., 2019; Borenstein and Bushnell, 2022*a*). It is a bundle of several different charges. Some cover the cost of energy. Some cover transmission and distribution. Some cover reliability, public-purpose programs, taxes, and other residual costs. Some reflect the value of being connected to the grid at all. A good tariff should not treat all of those costs as if they were the same.

The first distinction is between costs that change when a customer uses one more kilowatt-hour and costs that do not. Energy costs are the clearest example of the first type. Other costs are driven by peak demand or local network constraints. Still others are customer-specific or residual.

If a customer uses one more kilowatt-hour at 2 am on a mild spring night, the system cost may be low. If the same customer uses one more kilowatt-hour at 7 pm on a hot summer evening, the system cost may be much higher.

A useful way to name this first distinction is through two concepts the electricity industry treats as fundamentally different but that retail bills routinely blur: *energy* and *capacity*. Energy is the cost of actually generating the kilowatt-hours a customer consumes—primarily fuel and variable operating costs. It scales with how much electricity is used and is measured in kilowatt-hours. Capacity is the cost of having enough generation, transmission, and distribution capability standing ready to meet demand at the moment it is highest, regardless of how much energy is consumed over the year. It is built to the peak and is measured in kilowatts. Two customers who consume identical annual energy can impose very different capacity costs: the one whose usage spikes during the system peak requires the grid to carry more standing capability on its behalf than the one whose usage is spread evenly across the day.

This distinction does more organizing work in what follows than its simplicity suggests. Capacity costs are increasingly the fastest-growing component of the bill, driven by peak demand rather than by total consumption, and a flat per-kilowatt-hour charge cannot see the difference between a kilowatt-hour consumed at the peak and one consumed at 3 am. It therefore recovers a peak-driven cost through a price that has nothing to do with peak contribution. This is where flat rates do their quietest damage: because the charge tracks energy, customers face no price consequence for *when* they place load on the system, even though the capacity half of the cost is determined almost entirely by timing. Recent wholesale capacity auctions in several US regions have cleared at levels several times their historical averages, translating directly into the retail rate increases that have made capacity costs newly visible to customers and policymakers. Much of the rest of this primer is an argument for pricing the two through separate instruments—time-varying energy prices for energy, coincident-peak charges for capacity—rather than collapsing both into a single volumetric rate.

The second distinction is between the cost of using electricity and the cost of being connected to the grid. A household with rooftop solar or a data center with on-site generation may buy very few kilowatt-hours from the utility. But that does not mean they no longer depend on the grid. They still want backup service, balancing, exports, and reliable interconnection. That value should be recognized explicitly (MIT Energy Initiative, 2016; Pérez-Arriaga, Jenkins, and Batlle, 2017).

Figure 1 makes the bundle concrete on a real residential bill. Energy supply, which has costs that vary by hour, accounts for 43 percent of the total. Transmission and distribution charges, driven by contributions to peak demand, account for another 41 percent. Policy programs that do not vary with usage account for 14 percent, embedded in the volumetric rate as cents-per-kilowatt-hour riders. A small per-customer fixed charge accounts for the remaining 2 percent. The bill collapses four economically distinct cost categories into a single dollar amount, with most of the charges measured in cents per kilowatt-hour even though only one category—energy supply—is actually driven by kilowatt-hours consumed in a given hour.

Table 1 summarizes the basic logic.

A theme that runs through what follows is that many tariff problems that appear distinct share a common origin. When residual transmission, distribution, and policy costs are bundled

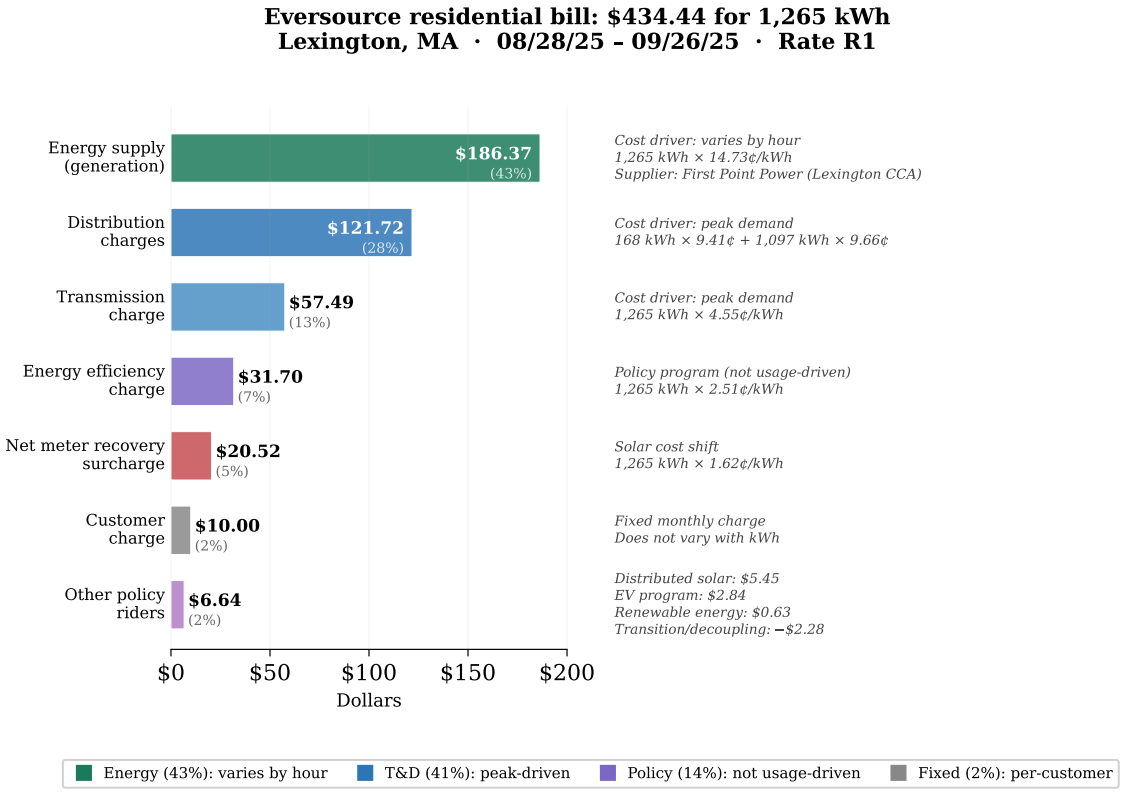

Figure 1: Composition of a typical residential electricity bill, decomposed by cost driver. The bill bundles charges with very different economic structures: energy supply that varies by hour, transmission and distribution charges that are peak-driven, policy programs that do not vary with usage, and a small per-customer fixed charge. Most categories are recovered through cents-per-kilowatt-hour charges regardless of their actual cost driver. Source: Eversource residential bill, Lexington, MA, Rate R1, August–September 2025.

Table 1: Cost categories and economically relevant tariff instruments

Cost category	**Primary cost driver**	**Better tariff instrument**
Short-run energy	Use in a specific hour and location	Time-varying volumetric energy price
Losses and congestion	Use or injection at constrained locations	Locational or zonal price adjustment
Generation capacity	Contribution to scarcity or system peak	Capacity charge or coincident peak charge
Transmission capacity	Contribution to regional peak flows or reserved service	Coincident peak charge, contract demand, reservation charge
Distribution capacity	Local peak demand, customer density, local hosting needs	Local demand charge, customer charge, hosting/export charge where appropriate
Customer-specific costs	Metering, billing, interconnection, service drop, customer support	Customer or connection charge
Residual and policy costs	Costs not attributable to hourly use	Fixed charge, income-adjusted charge, tax funding, explicit rider

into the volumetric energy price, and that price is held flat across hours and locations, a single number on the bill has to perform several incompatible jobs: signal the short-run marginal cost of electricity, recover capacity and network investment, fund policy programs, and carry whatever redistribution regulators have built into the rate base. The result is a price that is too high to encourage efficient electrification in many hours, too low to reflect scarcity in others, too easy to avoid through behind-the-meter generation, and too poorly targeted to do equity work efficiently. Much of the rest of this primer can be read as tracing specific consequences of this overloading. **The corollary is that many second-best instruments in current tariff design—peak-time rebates, EV-specific rates, and the contested compensation structures around rooftop solar, among others—exist to patch those consequences and would be smaller, more narrowly scoped, or unnecessary if the underlying prices were correct in the first place**.

3 Price Electricity When It Matters

The efficient tariff is easiest to understand by asking a simple question: what does it actually cost the system when a customer uses one more kilowatt-hour at a particular time and place (Bonbright, 1961; Schweppe et al., 1988; Borenstein, 2005). Sometimes the answer is: not much. If solar generation is abundant at noon and the grid is unconstrained, one more kilowatt-hour may be cheap to serve. Sometimes the answer is: a lot. If the system is stressed at 7 pm on a hot summer evening, one more kilowatt-hour may require expensive generation, worsen congestion, or contribute to the need for future grid investment. A well-designed energy price should reflect those differences. Electricity should be cheap when it is cheap to serve and expensive when it is expensive

to serve.

But charging the right price for energy is not enough. Some system costs do not change when a customer uses one more kilowatt-hour. The grid contains large fixed costs. Capacity is built in large increments, not one kilowatt-hour at a time. Policy programs are often layered onto rates. Some costs simply cannot be traced to a customer's decision to use electricity in a particular hour. Those residual costs must be recovered somehow. A practical efficient tariff therefore has three elements.

3.1 Time-varying energy prices

The volumetric price for energy should reflect what it costs to serve electricity at that time and place. In its most complete form, this means real-time pricing: the retail price changes hour by hour, or more frequently, with wholesale conditions, losses, congestion, scarcity, and environmental costs. In less complete but more administratively tractable forms, it can mean time-of-use pricing, critical peak pricing, or peak-time rebates.

The purpose is not to expose households to unnecessary volatility for its own sake (Allcott, 2011; Wolak, 2011; Jessoe and Rapson, 2014). The purpose is to send a simple signal: use more electricity when the system is cheap to run, and use less when the system is expensive or constrained. A heat pump, electric vehicle, water heater, battery, industrial process, or data center can often shift consumption in time. If retail prices are flat, the customer has no reason to move consumption from high-cost hours to low-cost hours. If prices vary with system conditions, customers and automated devices can reduce system costs. Experimental evidence consistently shows that residential customers do respond to dynamic prices by reducing peak demand (**??**).

3.2 Peak-coincident demand or capacity charges

Some network and capacity costs are driven not by annual kilowatt-hours but by usage during stressed periods. The relevant peak is often not the customer's own individual maximum demand. A customer who reaches maximum demand at 3 am may impose little capacity cost if the system peak occurs at 6 pm The economically relevant measure is the customer's contribution to system peaks, local feeder peaks, or other constrained periods.

What the peak charge should do, and how high it should be set, is inseparable from the question of what the rest of the tariff already conveys. The cleanest way to see this is through a thought experiment. Suppose retail energy prices passed through distribution-level locational marginal prices (DLMPs)—hour-by-hour, location-by-location signals reflecting the marginal cost of energy plus local distribution capacity constraints. In that world, a separate peak-coincident demand charge would be redundant. The energy rate would already reward customers for shifting consumption away from constrained hours and locations, and adding a peak charge on top would double-count costs already in the price. But DLMPs are mostly theoretical at the residential and small-commercial level today (Pérez-Arriaga, Jenkins, and Batlle, 2017). Peak-coincident demand charges exist as the administratively feasible quantity-based approximation to a DLMP regime that does not yet exist. The right level for the charge is whatever it takes to recover the cost-causation work the energy rate is not doing.

At a minimum, the peak charge should capture distribution-level capacity costs, since these are precisely the costs that wholesale price signals do not reach. Real-time prices and locational marginal prices in wholesale markets reflect the marginal cost of energy plus congestion and losses *at the transmission level*. They do not generally reflect distribution-level capacity constraints, which are local to a feeder or substation and driven by the local coincident peak, not the system or wholesale peak. A residential feeder's binding constraint may be a 7 pm summer evening when air conditioners are running, even if wholesale prices are higher in the afternoon. How much more the peak charge needs to capture depends on the energy rate. When retail rates pass through wholesale real-time prices—such as real-time pricing—transmission-level peak and generation capacity scarcity costs are already in the energy charge, so the peak demand charge should be calibrated to the marginal cost of distribution capacity alone. Under time-varying retail rates that only partially track wholesale conditions—typical multi-block time-of-use designs, for example—the peak charge would also need to capture whatever portion of transmission-level peak cost the energy rate fails to convey, particularly the scarcity tail that TOU blocks tend to flatten. Under flat volumetric rates, neither transmission nor distribution peak costs are signaled through the energy price, and the peak charge would have to absorb both.

In practice, such a charge is based on the customer's usage during the hours that drive system or network investment. A simplified version might charge customers based on their average usage during the top few system peak hours of the month or year. A more granular version could vary by location and by the type of constraint (Hledik and Greenstein, 2016). That is, you might have a distribution-peak coincident demand charge and a transmission-peak coincident demand charge.

The level of the charge follows from the same cost-causation logic. From an economic efficiency perspective, the right price is the long-run marginal cost of expanding capacity at the relevant level—distribution alone where the energy rate already captures transmission and generation peak costs, and more where it does not. In practice, regulators may also want to allocate some residual or historical network costs through the same instrument, both because long-run marginal cost estimates are uncertain and because historical investments reflect prior cost-causation decisions that should continue to be recovered. The design choice is how to balance forward-looking cost reflectiveness against full historical cost recovery.

A natural concern about this framework is that transmission and distribution capacity exhibits substantial economies of scale, so most network cost is largely fixed in any short-run sense rather than incrementally driven by peak demand. The implication sometimes drawn is that a peak charge can recover only a fraction of network revenue and is therefore the wrong instrument. The right response is to recognize what a peak charge is for. Utility engineering studies size networks to coincident peak demand, so a peak-coincident charge priced at the long-run marginal cost of expanding capacity is the correct cost-causation signal whether or not it recovers full embedded cost. Economies of scale mean that LRMC-priced charges will, in most mature systems, recover only a minority of total network revenue. The remainder—the residual—should be recovered through a separate instrument, typically a fixed or customer-count charge (MIT Energy Initiative, 2016; Brown, Faruqui, and Lessem, 2018).

Recent estimates for US distribution networks place LRMC at roughly \$50 to \$150 per kW-year (Turk et al., 2025), drawn from California Public Utilities Commission avoided-cost studies.

A typical residential coincident peak of three to five kW per household therefore implies an LRMC-priced peak charge of roughly $150 to $750 per customer per year, applied alongside an energy rate that passes through wholesale prices. Current embedded distribution revenue, at flat volumetric rates of around ten to twelve cents per kWh applied to typical US household consumption, runs around $700 to $900 per customer per year. Plausible central estimates suggest the LRMC peak charge would recover on the order of one-quarter to one-half of embedded residential distribution revenue, with substantial variation by utility, region, and rate of system growth. The point of these calculations is not that peak charges are unnecessary, but that they are inadequate on their own to recover transmission and distribution costs: a complete network tariff combines a cost-reflective peak charge to size the system with a fixed or customer-count residual charge to close the revenue gap that economies of scale create.

At the opposite extreme, some argue that the entire transmission and distribution revenue requirement should be recovered through peak-coincident demand charges, on the grounds that peak demand is what drives infrastructure investment in the first place. The intuition is appealing, but the economics is wrong. A peak charge set well above the long-run marginal cost of capacity expansion distorts the very margin the charge is meant to price correctly. Customers facing inflated peak charges would over-invest in peak avoidance—oversized behind-the-meter batteries, aggressive load-shifting controls, demand-side equipment installed past its social break-even point—reducing peak consumption past the point where the social value of the reduction equals the social cost of achieving it. The error is the mirror image of overpricing the kilowatt-hour: volumetric rates above marginal energy cost distort the kilowatt-hour margin, and peak charges above LRMC distort the peak-kilowatt margin. The right tariff prices each margin at its marginal cost and recovers any residual through a separate instrument that does not distort either decision.

Coincident peak charges are especially relevant for large customers. A data center that requires firm service during peak hours may impose substantial capacity and transmission costs even if it self-supplies much of its energy over the year (Knittel, Senga, and Wang, 2025). Conversely, a data center that can curtail load or shift computation during constrained hours imposes lower system costs and, under a well-designed coincident peak charge, pays a correspondingly lower bill. The compensation for flexibility is built into the charge itself: a separate payment for shifting away from peak hours would double-count the saving the customer already captures by avoiding the peak.

3.3 Transparent residual cost recovery

After energy costs and peak-driven network costs are recovered, some costs remain. These include customer-specific costs, residual network costs, policy costs, and other revenue requirements that cannot be sensibly assigned to one more kilowatt-hour consumed in a particular hour. These costs should be recovered through transparent instruments: connection charges, fixed charges, income-adjusted fixed charges, contract-demand charges, subscription charges, or funding outside the electricity bill (Ramsey, 1927; Brown and Sibley, 1986).

The appropriate instrument differs sharply by customer class. For residential customers, where the population is large and individual customer sizes are relatively similar, a per-customer fixed or income-graduated charge is administratively natural. For commercial and industrial customers, where customer sizes vary by orders of magnitude, a flat per-connection charge is meaningless: a

hyperscale data center cannot sensibly be charged the same connection fee as a small office building. The commercial and industrial analogs of residual cost recovery are capacity-based: contract demand charges, reservation charges, minimum bill provisions, standby charges, and contributions in aid of construction. These instruments scale with the customer's claim on the system rather than with metered consumption, and they perform the same economic function as a fixed charge does in the residential context: recovering costs that are not driven by the customer's marginal kilowatt-hour decisions.

The central point is that residual costs should not be hidden in the cents-per-kilowatt-hour price of energy. Doing so makes electricity look more expensive than it is to serve in many hours. That discourages efficient consumption, distorts electrification decisions, overcompensates behind-the-meter generation, and shifts costs to customers who cannot easily reduce net consumption.

3.4 What belongs on the bill in the first place

The framework so far has answered an allocation question: given that a set of costs has to be recovered through electricity rates, how should each cost be matched to an instrument that reflects its driver? That framing takes the revenue requirement as given and asks how to slice it across customers. But there is a logically prior question that cost causation alone cannot answer: which costs should be in the revenue requirement to begin with?

The question matters because what is currently bundled into electricity rates has grown well beyond the cost of producing and delivering electrons. Figure 1 makes this visible on a single residential bill: 14 percent of the total is policy programs that do not vary with usage, recovered through cents-per-kilowatt-hour riders. Those programs do not have to be on the bill—they are there because a series of past policy decisions put them there, and each could in principle be funded another way. Wildfire risk mitigation in California, accelerated distribution upgrades to keep up with electrification, public-purpose programs of various kinds, low-income assistance, and an array of energy-policy initiatives are all increasingly recovered through utility rates. Each can be defended on a narrow cost-causation argument—wildfire spending is caused by the existence of distribution infrastructure in fire-prone areas; electrification upgrades are caused by the electrification policies driving them. But the same logic could justify funding almost any climate or social policy through electricity rates, and at some point the question is not how to allocate the cost but whether the electricity bill is the right collection mechanism.

The alternatives are familiar: general taxation, dedicated funds (wildfire funds, climate adjustment funds, infrastructure banks), securitization with state backing, federal cost-sharing where the underlying policy is national in scope. Each has its own efficiency and equity properties, and each removes the cost from the per-kWh price seen by customers making electrification, behind-the-meter generation, and consumption decisions. The case for keeping a cost on the bill is strongest when the cost is genuinely driven by electricity use by the customers paying it; it weakens as the link between the customer's marginal decisions and the cost grows more attenuated.

A practical concern reinforces the point. As rates climb to absorb costs of this kind, the customers most able to leave—through rooftop solar, behind-the-meter storage, on-site generation, or full grid defection—have stronger incentives to do so, taking their contribution to fixed-cost recovery with them. The cost-shift dynamic that drives the rooftop-solar discussion in Section 6 is

not specific to solar; it is a general property of recovering large residual costs through any instrument that customers can avoid. Grid defection at the top of the rate compounds the problem: each departure raises rates for those who remain, accelerating the next round of departures. Containing the rate-recovery burden is not a separate consideration from cost allocation; it is part of the same problem.

The principle this suggests is to be conservative on inclusion and aggressive on transparency. Costs that customers genuinely cause through their electricity use belong on the electricity bill, allocated by the instruments described above. Costs that are driven by broader public objectives, or that arise from circumstances outside any individual customer's control, are candidates for funding outside the bill, particularly where the alternative funding source has a broader and more progressive base. The choice should be made explicitly, in light of the alternatives, rather than by default through whatever historical accident left the cost in rates.

4 How Dumb Prices Became Normal

Before criticizing existing tariffs, it is important to understand why they emerged. The current structure is the product of technological, institutional, and social constraints.

4.1 Metering technology

For most of the history of the electric industry, residential meters measured cumulative monthly consumption. They did not record when electricity was used. A utility could bill a household for total kilowatt-hours consumed during the month, but not for kilowatt-hours consumed during the system peak or during a particular wholesale price interval. Under those conditions, a flat volumetric rate was administratively natural (Houthakker, 1951).

Commercial and industrial customers were sometimes metered differently. Electromechanical demand meters could record a customer's maximum demand during the billing period, and traditional commercial and industrial demand charges relied on this capability. But a customer's individual non-coincident maximum demand is not necessarily aligned with the system or local peak that drives capacity costs. Old metering was therefore a workable proxy for cost causation only in a limited sense, and modern interval metering allows a much better match between charges and the periods that actually drive system investment.

Smart meters change this constraint. They allow utilities to measure interval consumption. That makes it feasible to implement time-of-use rates, critical peak pricing, real-time pricing, demand response, and more sophisticated export compensation. The technological justification for flat rates is therefore much weaker than it once was.

4.2 Limited customer response

Historically, most residential customers had limited ability to respond to time-varying prices. They could turn lights and appliances on or off, but they lacked automated devices capable of responding to system conditions. Today, that is changing. Electric vehicles can charge overnight or during solar-rich hours. Water heaters can preheat. Batteries can charge and discharge. Thermostats can

pre-cool. Data centers can, in some cases, shift computational workloads across time or geography (Knittel, Senga, and Wang, 2025). The more flexible load becomes, the more valuable efficient price signals become.

4.3 Distributional side effects

Recovering fixed costs through volumetric charges also had a distributional appeal (Borenstein, 2012; Burger et al., 2019). Higher-income households tend to live in larger homes and consume more electricity. If fixed grid costs are recovered through kilowatt-hour charges, higher-income households pay more. That made volumetric cost recovery a crude but politically attractive redistribution mechanism.

The problem is that this mechanism depends on the correlation between income and net electricity purchases (Borenstein, 2012; Burger et al., 2019). Rooftop solar weakens that correlation. A high-income household with rooftop solar may consume substantial electricity services while purchasing few net kilowatt-hours from the grid. If fixed costs are recovered volumetrically, that household may avoid paying for grid services it still benefits from. The burden shifts to customers without rooftop solar, who are disproportionately renters, lower-income households, or households with unsuitable roofs (Borenstein, Fowlie, and Sallee, 2021, 2022).

5 Volumetric Rates Are a Hidden Tax on Electrification

When fixed and residual costs are recovered through per-kilowatt-hour charges, the retail price can be far above the cost of serving one more kilowatt-hour in many hours (Burger et al., 2019; Borenstein and Bushnell, 2022*a*). That creates a wedge between what customers pay and what their consumption actually costs the system.

Suppose the system cost of serving one more kilowatt-hour at a particular hour is 8 cents, but the retail rate is 25 cents because it includes energy, transmission, distribution, policy costs, and residual cost recovery. A customer deciding whether to charge an electric vehicle, operate a heat pump, or run an industrial process sees 25 cents. The system cost may be only 8 cents. The customer may therefore use too little electricity from society's perspective, even when electricity is the cleaner or cheaper option.

This matters for decarbonization. Many climate strategies require electrification: vehicles, space heating, water heating, cooking, and some industrial processes must shift from direct fossil fuel use to electricity. If electricity prices are inflated by fixed-cost recovery, electrification becomes less attractive even when the system can serve additional electricity cheaply and cleanly (Borenstein and Bushnell, 2022*a*,*b*). The distortion is not unique to electricity: natural gas and gasoline are also mispriced, but the net wedge between retail price and social marginal cost is typically largest for electricity (Borenstein and Bushnell, 2022*b*).

Figure 2 illustrates the effect for electric vehicles using calculations from Clinton, Knittel, and Metaxoglou (2020), who compute, for every battery price, the oil price at which an EV is cheaper over the life of the vehicle than a comparable internal-combustion-engine vehicle. This break-even oil price depends on how electricity is priced. The figure plots break-even curves under a typical bundled retail rate (around 25 cents per kWh) and under marginal-cost pricing (around 8 cents per

kWh). At today's battery prices of roughly $120 per kWh, EVs are cost-competitive with gasoline vehicles whenever oil exceeds about $64 per barrel under marginal-cost pricing—below current oil prices—but they require oil above $126 per barrel before they become cost-competitive under the bundled rate. The shaded gap is the hidden tax on electrification.

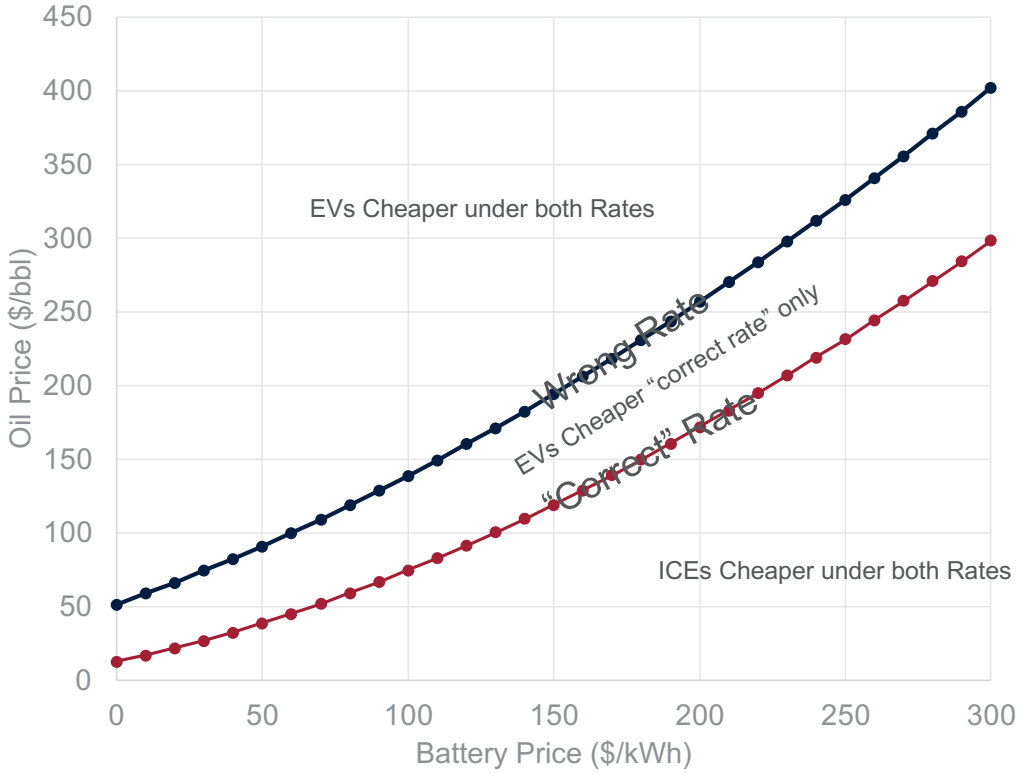

Figure 2: Volumetric T&D charges raise the break-even oil price for electric vehicles. Each curve shows the oil–battery price pairs at which an EV and a gasoline vehicle have equal total cost of ownership. At today's battery prices ($120/kWh), the break-even oil price is roughly $126/bbl under the bundled retail electricity rate but only $64/bbl under marginal-cost pricing—below current oil prices. The shaded gap is the hidden tax on electrification. Adapted from Clinton, Knittel, and Metaxoglou (2020).

The same distortion affects behind-the-meter generation. If a customer produces one kilowatt-hour behind the meter, the private savings equal the avoided retail rate. But if the retail rate includes fixed grid and policy costs, the private savings exceed the costs the system actually avoids. This over-incentivizes self-generation relative to grid-scale alternatives and encourages customers to reduce net purchases even when they continue to rely on the grid. Recent evidence suggests that rooftop solar is associated with higher retail electricity prices, likely because rate structures shift fixed costs onto non-solar customers (Argosino and Knittel, 2025).

A common response is that customers should pay less if they use less. That is partly right. Customers should pay less when they reduce costs. But reducing net kilowatt-hours does not necessarily reduce transmission costs, distribution costs, policy costs, or reliability costs. Tariff design should distinguish between cost-reducing behavior and revenue-avoiding behavior.

High volumetric rates also create political constituencies (Borenstein, 2012; Burger et al., 2019). They increase the private return to rooftop solar, because every kilowatt-hour generated behind the meter avoids the full bundled retail rate. They also increase the apparent payoff from energy efficiency, because every avoided kilowatt-hour saves not only energy costs but also a share of grid, policy, and residual costs embedded in the volumetric rate. This means that the solar industry and parts of the energy-efficiency industry often benefit from maintaining high volumetric rates.

Solar deployment and energy efficiency are worthy societal goals (Borenstein, 2012; Burger et al., 2019). The problem is not the goals; it is the instrument. Inflating the cents-per-kilowatt-hour price to encourage solar or efficiency is a poorly targeted subsidy. It encourages some socially valuable behavior, but it also discourages other socially valuable behavior, including electrification of vehicles and buildings. It shifts costs toward customers who cannot install solar or make large efficiency investments. And it makes electricity consumption look expensive even in hours when the system can serve additional load cheaply. If policymakers want more solar or more efficiency, they should support those goals directly and transparently, not by distorting the retail price of electricity.

It is worth being explicit that the wedge this section describes is itself a consequence of recovering residual costs through the volumetric rate. Under the residual-cost recovery instruments described in Section 3—fixed or income-graduated charges for residential customers, contract-demand and standby charges for large customers—the volumetric rate would track marginal cost much more closely, and the wedge would shrink substantially. This argument has a mirror image that has to be addressed directly, because it appears to cut the other way. If volumetric rates are a tax on electrification, and electrification is something we want more of, why are utilities across the country pointing to electrification as the fastest-growing source of upward pressure on rates? The answer is that the two claims describe different hours. The hidden tax operates in the many hours when the system is unconstrained and a marginal kilowatt-hour is cheap to serve: in those hours the inflated volumetric rate discourages exactly the beneficial load—overnight EV charging, midday water heating—that the grid could absorb at low cost. The rate pressure operates in the few hours when new electric load lands on the system peak: a winter cold snap that turns on a region's heat pumps at once, or a fleet of EVs that all begin charging at 6 pm, adds coincident peak demand, and coincident peak is what drives capacity investment. Both are real, and a flat rate gets both wrong at the same time—it overprices the off-peak kilowatt-hour the system wants and sends no signal at all about the on-peak kilowatt-hour the system struggles to serve (Borenstein and Bushnell, 2022*b*).

The reconciliation is the central argument of this primer. The problem is not electrification; it is *unmanaged* electrification that lands on the peak because nothing in the price tells it to do otherwise. Time-varying prices are precisely the instrument that moves new electric load into the hours when the system can serve it cheaply: they lower the price in the valleys, where electrification belongs, and raise it at the peak, where unmanaged electrification would otherwise drive capacity cost. This makes electrification and time-varying pricing *linked* reforms rather than separate ones. Electrify on a flat rate and the new load inflates the peak and raises rates for everyone; electrify with prices that reflect system conditions and most of that load can be steered into the valleys, deferring the capacity investment that would otherwise be required. Pursued apart, electrification without time-varying pricing produces higher peaks and higher rates—which is exactly the trend

utilities are now reporting.

The electrification penalty, the over-subsidy to behind-the-meter generation, and the political constituencies that depend on a high volumetric price are not features of electricity pricing in general. They are features of one particular choice about where to put residual costs.

6 Under Current Rate Design the Grid Is Insurance and a Free "Battery": Rooftop Solar, Data Centers, and the Cost Shift

Rooftop solar is the clearest residential example of a broader tariff-design problem: customers can increasingly reduce measured net consumption without proportionally reducing their reliance on the grid (Schittekatte, Momber, and Meeus, 2018). The issue is not self-consumption itself. A customer who produces a kilowatt-hour behind the meter and consumes it on the same premises has used no grid service for that kilowatt-hour, and a well-designed marginal price would correctly charge them nothing for it. The economic problem with traditional net metering arises specifically with exports. A kilowatt-hour exported to the grid offsets a future kilowatt-hour of imports one-for-one at the full retail rate. That would be economically sensible if the retail rate reflected only the avoided marginal cost of energy at the relevant location and time. But in most jurisdictions, the retail rate also includes transmission, distribution, policy, and other residual costs that the exported kilowatt-hour does not save the system from incurring: the neighbor who consumes the exported kilowatt-hour still relies on the distribution network to receive it, and the policy programs funded through the volumetric rate continue to operate. Net metering at the full retail rate therefore overpays the customer for the value the export actually delivers.

There is a second, related distortion. Because residual transmission, distribution, and policy costs are folded into the volumetric rate, any reduction in measured net imports—through efficiency, self-consumption, or net-metered exports—reduces a customer's contribution to recovering those costs even though the utility's revenue requirement is unchanged. The unrecovered amount is shifted to other customers through higher rates, riders, or future rate cases. The mechanism is not self-consumption itself, which is benign, but volumetric residual recovery, which makes residual cost contributions erodable through any reduction in imports.

6.1 A simple rooftop solar example

Consider a utility with 1,000 residential customers. Suppose each customer initially pays $100 per month toward transmission, distribution, policy, and other residual costs, so the utility collects $100,000 per month in non-energy revenue.

Now suppose 200 of the highest-income customers install rooftop solar systems sized so that they become net-zero customers for these grid-related charges. Their energy purchases fall, which is real and valuable. But their need for the grid does not disappear: they still rely on the system at night, during cloudy periods, during outages, and as a platform for exporting surplus power. If those 200 customers stop contributing their $100 monthly share of grid costs, the utility must still recover the same $100,000 from the remaining 800 customers. Each remaining customer's bill for non-energy costs rises from $100 to $125 per month.

What began as a technology subsidy becomes a cost shift: the 200 rooftop solar customers shift $20,000 per month in grid costs onto everyone else, raising the non-energy bill of the remaining customers by 25 percent. Empirical work in high-solar jurisdictions confirms this pattern: in Queensland, Australia, the shift from solar to non-solar households amounted to hundreds of dollars per year (Simshauser, 2016).

6.2 Why "solar lowers rates" is not an argument for solar subsidies

A common defense of net metering responds to the cost-shift literature by arguing that rooftop solar generates ratepayer benefits that should be netted against it: lower wholesale prices, deferred utility investments, suppressed retail rates. The empirical adjustments embedded in this argument—accurate capacity factors, proper accounting for self-consumption, the share of solar customers receiving low-income rate discounts—are legitimate refinements to cost-shift calculations and can move the headline number around. But the central conceptual claim, that rooftop solar's downward pressure on rates is itself a benefit warranting compensation through above-marginal-cost export rates, is a conceptual error.

The cleanest way to see this is by analogy. Suppose new oil and natural gas fields come online. The additional supply lowers the world oil price and, especially for natural gas, the regional wholesale price. Consumers of gasoline and natural gas pay less; existing producers earn less per unit; oil- and gas-exporting governments collect less royalty revenue. No serious economist would argue that the consumer savings constitute a social benefit warranting a subsidy to the new fields, equal to the consumer savings, beyond whatever price the fields receive in the market. The price effect is a transfer from producers to consumers, not a real-resource saving. From a social-welfare perspective, the consumer gain and the producer loss net to approximately zero in equilibrium. What economists call this is a *pecuniary externality*—an externality that operates through prices—and the standard treatment is to ignore it in benefit-cost analysis. We contrast that with *technological* externalities—emissions, learning-by-doing, network spillovers, real reliability effects—that do change real resource costs and are legitimate grounds for policy intervention.

Let's be clear, rooftop solar generates real technological externalities—avoided emissions, learning effects, possibly grid resilience contributions during specific contingencies—and the avoided-cost framework that California uses for net billing is designed precisely to capture these along with the real-resource savings the export delivers (avoided fuel and variable O&M from displaced generation, avoided generation capacity, avoided ancillary services procurement, deferred transmission and distribution capacity in present-value terms, and avoided line losses). What the framework deliberately excludes is the pecuniary effect of rooftop solar on retail rates paid by other customers, because that effect is a transfer rather than a resource saving. Counting it as a benefit would be inconsistent with how every other technology policy is evaluated.

Applied consistently, the "solar lowers rates" argument would justify subsidies for any technology or behavior that reduces utility revenue requirements, including utility-scale solar, demand response, energy efficiency, and even customer departures from the system. The final example is useful: a household that moves away and disconnects from the grid also reduces the revenue the utility must collect from remaining customers; no one proposes paying departing households for the favor. The reason the argument seems plausible specifically in the rooftop solar context is that

the technology is politically favored and the price effect happens to flow in a politically attractive direction. Neither consideration changes the underlying economics.

The deeper reason the "solar lowers rates" argument retains some intuitive appeal is that rooftop solar does deliver a real technological externality through avoided emissions, and current retail rates fail to price that externality in any direct way. Defenders of generous net metering can therefore frame above-marginal-cost export compensation as a rough vehicle for the environmental benefit, even when their stated justification appeals to rate suppression. If retail energy prices reflected social marginal cost—wholesale energy cost plus environmental damages at the relevant time and location—the avoided cost on which export rates should be based would already include the emissions value, and the residual case for paying exports above that level would mostly disappear. The pattern is the same one this primer has traced repeatedly: a second-best instrument that looks reasonable in isolation looks much less reasonable once the underlying prices are corrected. Fix the prices and a substantial share of the political constituency for retail-rate net metering loses its strongest non-pecuniary argument.

6.3 Data centers as a different challenge

Data centers raise related issues, but the mechanics differ from residential solar in important ways (Knittel, Senga, and Wang, 2025). Large customers typically face demand-dominated tariffs, so a data center cannot zero out its contribution to fixed-cost recovery simply by reducing volumetric purchases. The cost-shift story that applies cleanly to residential rooftop solar applies less cleanly to large industrial customers whose bills are anchored by demand charges, minimum bills, and contract terms.

The deeper problem is that the tariffs themselves were designed for a different scale of customer. Standard demand charges are calibrated against historical embedded cost-of-service studies that divided existing rate base across customer classes. They were built around customers whose interconnection triggered shallow incremental upgrades at the distribution level—a transformer replacement, a feeder upgrade, perhaps a local substation modification. Hyperscale data centers are different in kind, and they connect to the grid differently as well: they typically interconnect directly with the transmission system at high voltage, bypassing the distribution network entirely. The infrastructure their growth drives is therefore concentrated in transmission and generation rather than distribution. New transmission substations, transmission upgrades, and sometimes new generation capacity can be required, and the marginal cost of this buildout can substantially exceed the average embedded cost reflected in legacy demand charges. Even a customer paying the full standard rate can impose net costs on other ratepayers if it triggers infrastructure that the demand charge was never sized to recover. The transmission-level connection has a further implication that is sometimes obscured in policy debates: hyperscale data center load growth does not help recover distribution system costs from a broader base, because these customers bypass the distribution network. Whatever cost-spreading benefits new load might provide accrue at the transmission and generation level, not at the level where distribution residual recovery is most acute.

Behind-the-meter generation introduces a further complication. Demand charges typically price observed peak demand at the meter. A data center using on-site generation to ride through high-

price intervals or system peaks can reduce its measured peak demand, and therefore its demand-charge bill, without reducing the underlying capacity it requires from the grid. The grid still must stand ready as backup if the on-site generator trips, if fuel delivery is interrupted, if maintenance is needed, or if self-supply becomes uneconomic. This is the insurance value of grid connection, and it is rarely priced explicitly in standard tariffs.

It is sometimes argued that the insurance value of grid connection is priced indirectly: a data center that does draw on the grid pays the prevailing energy and demand charges, so as long as those charges reflect short-run system conditions during the moments of use, the data center's expected payment is the correct insurance premium. In principle, this is right. Efficient scarcity prices, integrated over the year, would in equilibrium deliver expected revenues sufficient to fund the capacity the system commits to in advance. The trouble is that actual wholesale prices fall short of that benchmark. Bid caps, operating reserve demand curves that flatten the scarcity tail, out-of-market interventions during system stress, and political resistance to extreme prices all suppress observed prices below efficient scarcity levels. The gap between what spot prices generate and what efficient pricing would generate is the missing-money problem familiar from wholesale capacity markets, which exist precisely to provide a separate revenue stream funding capacity adequacy. Retail real-time prices inherit this suppression because they pass through wholesale prices. A data center charged only for energy and demand at suppressed real-time prices does not pay an expected amount sufficient to cover the capacity the system has reserved on its behalf. The natural instrument is therefore a standby or reservation charge paid on contracted backup capacity, regardless of actual draw. **A customer that rarely uses the grid is not receiving little value from it; the option to draw is itself valuable, and providing that option has a cost.**

The same on-site generation that complicates cost recovery can, pointed the other way, be a system asset—and it is economically equivalent to a margin the data-center literature has begun to take seriously: shifting computational workloads in time. A data center that defers a training job from a scarcity hour to a low-price hour reduces its draw on the grid in that hour by exactly the same amount as one that meets the same load from an on-site generator. From the grid's perspective the two responses are indistinguishable: both reduce the facility's contribution to the system peak, both relieve capacity that would otherwise need to be built or held in reserve, and both should be rewarded through the same instrument. Recent work shows that the temporal flexibility margin is substantial (**?**), which matters because workload shifting is cleaner, cheaper, and faces fewer permitting and air-quality constraints than running diesel backup generation. The broader point is that a facility that shifts, curtails, or self-generates during scarcity is supplying reliability through whichever margin is cheapest for it, and this is the beneficial mirror image of the insurance-avoidance problem above: the concern there is a facility that uses on-site generation to shrink its measured peak while still leaning on the grid for reliability; a facility that makes its flexibility available when the *system* is stressed is doing the opposite. Used this way, data-center flexibility can let utilities interconnect large new loads sooner than firm grid capacity alone would allow, bridging the gap until dispatchable balancing resources can be built. It is also fully consistent with the flexibility principle developed below: a well-designed tariff rewards these responses through the peak charge they avoid rather than through a separate payment. The emissions caveat applies specifically to the self-generation margin: emergency backup generation is typically diesel or gas,

and although its annual emissions footprint is modest when run only for genuine system emergencies, local air-quality limits and permitting constraints are real and can bind well before the annual-hours arithmetic suggests.

Stranded-cost risk is a third issue. Concentrated large loads can leave—through technology changes, business failures, or relocations—and the cost of purpose-built infrastructure can then fall on remaining ratepayers. Standard demand charges do not price this risk. That is why emerging large-load tariffs include long contract terms, minimum bill provisions, and contributions in aid of construction that ask the new customer to commit in advance to a portion of the cost of the facilities built to serve it.

The key distinction is that net consumption is not the same as grid reliance. A household with rooftop solar and a data center with on-site generation may both rely on the grid for backup, balancing, exports, reliability, and reserved capacity, even when they purchase few net kilowatt-hours. The instruments that recover those costs differ across customer classes—volumetric for residential, demand-based for large customers—but the underlying problem is similar: tariffs that were calibrated to a prior generation of customer behavior do not fully capture what customers cost the system today.

7 Appliance-Specific Rates Are a Bandaid

A recent response to high volumetric electricity rates is to create special rates for particular appliances or uses, especially electric vehicles and heat pumps. These rates are understandable. Policymakers want to encourage electrification, and they recognize that high bundled electricity prices make it harder to switch from gasoline cars or fossil heating to electric alternatives. But appliance-specific rates are a bandaid on a deeper tariff problem. The efficient price of electricity does not depend on which appliance uses the kilowatt-hour. It depends on when and where the kilowatt-hour is used, whether the grid is constrained, and what costs the customer imposes on the system (Borenstein and Bushnell, 2022*b*).

An electric vehicle charging at 2 am and a water heater running at 2 am impose similar energy costs if they are served at the same location and time. A heat pump running during a winter peak and a resistance heater running during that same winter peak both contribute to the same capacity problem. The system does not care whether the electron goes into a vehicle battery, a compressor, a server rack, or a dishwasher. What matters is timing, location, peak contribution, flexibility, and grid reliance. A tariff that varies by appliance rather than by these cost drivers is trying to solve the right problem with the wrong instrument.

A more theoretically grounded defense of appliance-specific rates appeals to Ramsey-Boiteux pricing logic (Ramsey, 1927; Boiteux, 1956; Borenstein, 2025). Some electrification uses have close substitutes outside the electricity bill—gasoline for transportation, natural gas for space and water heating—so demand for electricity for those uses may be more elastic than demand for captive uses like lighting and refrigeration. Ramsey logic says that when residual costs must be recovered through markups over marginal cost, markups should be inversely proportional to demand elasticity: less elastic uses should bear higher markups, more elastic uses lower markups. A discount rate for electric vehicles and heat pumps is, in this view, the application of Ramsey pricing to electrification:

charge less on the margins that would otherwise be lost to a substitute fuel.

The Ramsey defense is internally consistent, but it rests on a premise that the primer has been arguing against. Ramsey-Boiteux pricing is a second-best result. It only becomes relevant once a regulator has decided to recover residual costs through markups over marginal cost in the first place. Under the residual-cost recovery instruments described in Section 3—fixed or income-graduated charges for residential customers, capacity-based instruments for commercial and industrial customers—the markup over marginal cost shrinks toward zero, and the Ramsey case for differentiating rates by appliance shrinks with it. The appliance-specific rate is in this sense a second-order patch on a first-order distortion: it tries to undo, for selected uses, a markup that should not have been there in the first place. Removing the markup itself is the cleaner fix, and it has the advantage of restoring efficient prices for every use rather than only the politically favored ones.

Even taken on its own terms, the Ramsey defense has force that varies by use in ways that matter. The case is strongest for EV charging, which can typically be shifted to genuinely off-peak hours when both marginal cost is low and the substitute—gasoline—carries a substantial environmental externality. The case is weaker, arguably reversed, for heat pumps. Heat pump demand for electricity is concentrated in winter cold spells, when the marginal cost of serving load is high and capacity is most constrained. A flat heat-pump discount during exactly those hours moves consumption in the wrong direction relative to system cost. Once the heat pump is installed, moreover, the relevant substitute for running it on a cold night is not the gas furnace, which may no longer exist, but lower indoor temperatures—a much weaker substitute that makes short-run demand for heat-pump electricity less elastic than the long-run electrification decision suggests. The Ramsey case for an EV discount is therefore not the same as the Ramsey case for a heat pump discount, and the latter can recommend a higher rate during constrained hours rather than a lower one. This reinforces rather than weakens the basic point: the right instrument is a price that varies by time and location, not by appliance.

Electric-vehicle rates are the more defensible version of appliance-specific pricing because EV charging can often be separately metered or observed through managed charging equipment. EV charging also has a useful behavioral property: it is widely perceived as happening in the background. A vehicle plugged in overnight charges itself without the household noticing, and shifting EV charging to off-peak hours through automatic scheduling imposes essentially no inconvenience. Other appliances do not always share this property. A dishwasher or laundry cycle is started intentionally, and shifting space-heating or cooling load can be noticeable during temperature extremes—though most of the time even thermostatic loads can be shifted with little effect on comfort. Outside of those extremes, resistance to shifting flexible loads is often more a matter of utility caution than of customer preference. But even granting the unique convenience of EV charging, the better logic is not "EVs deserve a lower rate." The better logic is "flexible load should be rewarded for charging when the system is cheap and avoiding hours when the system is constrained." If an EV charges during low-cost hours, it should face a low price. If it charges during peak hours, it should not receive a discount simply because it is an EV. The same principle should apply to any flexible load.

Heat-pump rates are often more problematic because utilities usually do not separately meter

the heat pump. Without submetering, the utility cannot observe how many kilowatt-hours the heat pump actually uses or when it uses them. The rate therefore cannot truly be a lower price for heat-pump electricity. Instead, it becomes an administrative transfer to the household based on assumed usage. The regulator or utility estimates how much electricity the heat pump is expected to use and applies a discount based on that assumed quantity. In simplified form, the transfer is:

$$\text{Transfer} = \text{Assumed heat-pump usage} \times (\text{standard rate} - \text{heat-pump rate}).$$

That may still encourage adoption, but it is important to be clear about what is happening. The household is not actually facing a lower price for the electricity used by the heat pump unless the heat pump is separately metered. **The household is receiving a fixed or formula-based subsidy for owning the appliance; they are not receiving a lower rate to run their heat pump.** Once the subsidy is calculated, the customer's actual decision about when to run the heat pump may still be governed by the standard household electricity rate. If the standard rate is flat and inflated by fixed-cost recovery, the customer still does not see the true cost of running the heat pump in low-cost versus high-cost hours.

This distinction is subtle but important. A heat-pump discount without submetering can help with adoption, but it does not fix the price signal. It does not tell the household to preheat when electricity is cheap, reduce consumption during system peaks, or respond to renewable availability. It also creates new administrative questions: what assumed usage should be used, how should it vary by climate zone and home efficiency, should larger homes receive larger transfers, and how should the policy treat households that install heat pumps but use them only as backup?

The better approach is to fix the underlying electricity tariff. If electricity is cheap to serve in a given hour, all customers and all appliances should see a low price. If the system is constrained, all flexible loads should have a reason to reduce or shift consumption. If policymakers want to accelerate EV or heat-pump adoption, they should use transparent rebates, tax credits, low-income installation support, or financing programs. Appliance-specific rates can be politically attractive, but they risk adding another layer of complexity to a tariff structure that is already sending the wrong signals.

8 Smart Meters, Dumb Prices

The second major distortion is that most customers pay one energy price even though the cost of serving electricity varies sharply over time. This has always been true, but it becomes more important as renewable penetration rises.

In a conventional fossil-dominated system, wholesale prices vary with load and fuel costs. In a renewable-heavy system, prices can vary even more dramatically (Schittekatte et al., 2024). During hours with abundant solar or wind, the system may be able to serve additional electricity at very low cost. During hours when renewable output is low and demand is high, the system may need expensive generation, scarce stored energy, or costly grid capacity.

Flat retail prices hide these differences. They tell customers that a kilowatt-hour at 3 am, 3 pm, and 7 pm is worth the same. This is the essence of the problem: smart meters, dumb prices. We now have the technology to measure when electricity is used, but many tariffs still pretend that

timing does not matter. In a decarbonized grid, that is increasingly false. Figure 3 illustrates the mismatch for a stylized high-solar grid: the system value of electricity drops below 5 cents during midday solar surplus but spikes above 40 cents during the evening peak, while the flat retail rate remains at 25 cents throughout.

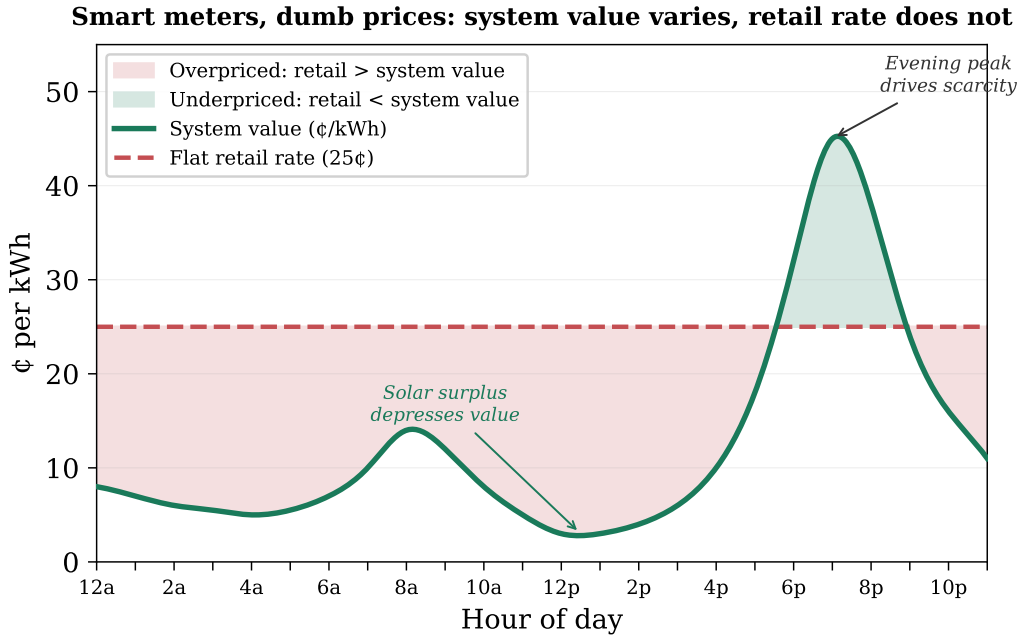

Figure 3: Smart meters, dumb prices. The system value of electricity varies sharply across hours in a renewable-heavy grid, but a flat retail rate hides these differences from customers. During solar-rich midday hours, the system is overpriced (red shading), discouraging efficient consumption. During the evening peak, the system is underpriced (green shading), providing no incentive to reduce load.

This matters for both demand and supply (Allcott, 2011; Wolak, 2011; Jessoe and Rapson, 2014). A customer with an electric vehicle should have an incentive to charge when energy is cheap, and the grid is unconstrained. A customer with a battery should have an incentive to charge during low-cost hours and discharge during high-cost hours. A data center that can shift workloads should have an incentive to reduce demand during constrained periods. A rooftop solar customer should receive higher compensation when exports are valuable and lower compensation when exports arrive during low-value or congested hours.

Time-varying prices can take several forms:

- **Time-of-use rates** divide the day into predefined blocks, such as off-peak, mid-peak, and peak.
- **Critical peak pricing** applies much higher prices during a limited number of system-stress events.
- **Peak-time rebates** pay customers for reducing consumption during called events relative to a baseline.
- **Real-time pricing** passes through hourly or sub-hourly wholesale prices, sometimes with hedging or bill-protection features.

- **Dynamic export compensation** pays distributed generators based on the time-varying value of their exports.

The appropriate design depends on metering technology, customer sophistication, automation, political feasibility, and risk tolerance. But the direction of travel should be clear: the more the cost of electricity varies across hours, the more costly it is to hide that variation from customers.

A common objection in utility proceedings is that customers will face bill shock or will not bother to shift consumption to save modest amounts. This argument was more persuasive when electricity was cheap and meaningful savings required active behavioral change. In an affordability crisis, with electricity bills under sustained upward pressure and substantial annual savings now available through automation, the case for withholding cost-control tools from customers has weakened considerably. Smart thermostats, EV charging schedulers, water-heater controls, and home energy management systems make time-varying rates much closer to "set and forget" than to constant vigilance. Bill-shock concerns themselves admit a less drastic remedy than flat rates: real-time pricing combined with conservative price caps recovers most of the available efficiency gains while bounding customer exposure during the extreme hours that drive the concern (Hinchberger et al., Forthcoming). Customers should retain the option to choose fixed-price contracts, and some will value that hedge enough to pay a modest premium for it. But the default presumption that customers cannot or will not respond to better price signals—and therefore must be protected from those signals by regulators—is increasingly difficult to defend on either efficiency or equity grounds.

A related but separate concern is that customers care about month-to-month bill predictability even when they do not care about hour-to-hour price volatility. A household budgeting for a bill of around $150 per month is rightly distressed if a single hot week pushes that to $300. But the appropriate response is not to eliminate hourly price variation. It is to recognize that the bill-predictability problem is really two distinct services. The first is bill smoothing: spreading actual costs evenly across months so that the timing of payment matches the household's budgeting cycle. Bill smoothing is essentially free and changes nothing about the marginal price the customer faces. Budget billing, already widely offered, levelizes monthly payments and reconciles annually, while the customer can still see real-time prices when deciding whether to run an appliance. The second service is genuine price insurance: protecting the customer against the realized variation in their costs. Insurance is not free—someone has to bear the risk that prices spike, and that risk-bearing has to be compensated. **A fixed-rate contract is, in economic terms, a real-time-price contract bundled with an insurance policy, and the price of the insurance is the markup of the flat rate above the expected real-time bill.**

In a well-designed system, the customer chooses how much insurance to buy. A household that prefers a fixed price can purchase a fixed-rate contract at a modest premium above the expected real-time bill. A household that prefers to face marginal prices, capturing the savings of consuming more in low-cost hours, can opt out of the insurance. The current tariff system gets this exactly backwards. Every customer is enrolled by default in a mandatory bundled insurance product, and the cost of providing that insurance—the risk premium embedded in the flat rate—is socialized across customers, including those who never wanted the hedge and would have preferred to face marginal prices. The default should be reversed: marginal pricing for those who want it, explicit insurance for those who pay for it, and bill smoothing for anyone who wants predictable monthly

outlays without the risk transfer.

8.1 Does real-time pricing destabilize utility revenue?

A separate objection to real-time pricing comes from the utility side rather than the customer side. If retail prices vary hour by hour, how does a utility build a budget? Cost-of-service ratemaking requires the utility to project costs over a future period and design rates that recover them. A tariff in which the price the customer pays moves with wholesale conditions, and in which automated devices may shift consumption in response to those prices, can look like a tariff that makes revenue impossible to forecast.

The objection conflates two distinct prices that happen to be measured in the same units. The price the *customer* pays for energy varies hour by hour under real-time pricing. The price the *utility* needs to recover its fixed and residual costs does not. They are separate instruments in the tariff, and the framework developed in Section 3 keeps them separate by design. Energy charges—volumetric, time-varying, and approximately equal to marginal cost—pass wholesale price variation through to customers. Fixed and capacity-based charges—per-customer fixed charges for residential customers, contract-demand and reservation charges for large customers—recover the costs that do not vary with hourly consumption. The first set is volatile because wholesale conditions are volatile; the second set is exactly as predictable as the cost categories it is sized to recover.

A utility moving from flat volumetric rates to the framework in this primer therefore sees its revenue become *more* predictable on the fixed and residual side, not less. Capacity and residual costs that previously had to be recovered through a kilowatt-hour charge—and whose recovery therefore depended on weather, economic conditions, and customer behavior—are moved to instruments tied to capacity reservations or customer counts, which are stable on the timescale of a rate case. What the utility loses is the implicit hedge that flat rates provided against wholesale price spikes for the energy passed through. But that hedge was being sold to every customer regardless of whether they wanted it, and its cost was built into the flat rate as a risk premium above expected wholesale cost.[1] Marginal-cost pricing on the energy side does not destabilize that side of the business so much as it makes the existing risk transfer visible and optional.

The balancing concern—that production must exactly equal demand, and there is no storage backstop large enough to absorb mismatches—cuts in favor of real-time pricing rather than against it. Price-responsive demand is itself a balancing resource. A heat pump that pre-heats during low-price hours, a fleet of EVs that shifts charging away from a constrained evening, and an industrial process that ramps with system conditions all reduce the gap the utility must close through generation dispatch. The harder balancing problem in a renewable-heavy system is precisely the one that flat retail rates make worse, by hiding from flexible loads the hour-by-hour conditions they could otherwise help solve.

[1] A utility that still wants to offer hedged products can do so explicitly, the way competitive retailers in restructured markets do: a fixed-rate contract priced as the expected real-time bill plus a transparent risk premium. The hedge becomes a product customers can choose rather than a default everyone pays for.

9 The Tariff Menu

There is no single electricity tariff used around the world. Regulators use different combinations of volumetric charges, fixed charges, demand charges, time-varying rates, block rates, export compensation, and policy riders. The main structures are summarized below.

9.1 Flat volumetric rates

A flat volumetric rate charges the same cents per kilowatt-hour regardless of when electricity is consumed. It is simple and familiar, but it is poorly aligned with system cost. It also becomes problematic when fixed costs are embedded in the volumetric charge.

9.2 Inclining and declining block rates

Inclining block rates charge a higher price as monthly consumption increases. They are often justified on conservation or equity grounds. Declining block rates charge a lower price at higher consumption levels and have historically been used in some industrial contexts where delivery costs fall with scale. Block rates are blunt instruments: they vary with total monthly consumption, not with hourly system conditions or cost causation. Evidence suggests that many consumers respond to average rather than marginal prices, which can further blunt the incentive effects of nonlinear tariff schedules (Ito, 2014).

9.3 Time-of-use rates

Time-of-use rates charge different prices during pre-specified periods. They are easier to understand than real-time pricing and can provide useful incentives, especially for electric vehicles and other flexible loads. Their weakness is that the true high-cost hours may not always match the predetermined windows.

9.4 Critical peak pricing and peak-time rebates

Critical peak pricing imposes high prices during a small number of system-stress events. Peak-time rebates pay customers for reductions during those events relative to a baseline. Both designs target the hours that drive system costs while limiting exposure to frequent price changes. They differ fundamentally, however, in their incentive structure and in the problems that arise in practice. Recent evidence quantifies the efficiency gap each instrument leaves on the table: across US wholesale markets, simple time-of-use rates and well-designed critical peak pricing each recover roughly ten percent of the mispricing wedge between flat retail rates and real-time pricing, with a combined recovery of about seventeen percent because the two instruments target largely complementary sources of price variation (Hinchberger et al., Forthcoming).

Critical peak pricing is symmetric: a customer who consumes during a peak event pays the high price, and a customer who avoids consumption avoids paying. The economic signal is the same as in a price-based system, just concentrated in a small number of hours. Peak-time rebates are asymmetric: a customer who reduces consumption relative to a baseline gets paid, but a customer who increases consumption above the baseline pays nothing extra. This asymmetry is the source

of three well-known problems with rebate-style demand response programs (Bushnell, Hobbs, and Wolak, 2009; Borenstein, 2014).

The first problem is moral hazard in baseline construction. Because the baseline is typically estimated from the customer's consumption in non-event periods, particularly other high-demand days, customers have an incentive to consume more during baseline-setting hours so as to inflate the baseline against which event-period reductions are measured. The most extreme documented case involved a baseball stadium turning on its lights during electricity shortages in order to be paid for turning them off; less dramatic versions of the same incentive arise routinely. The Anaheim baseline-rebate experiment provides direct empirical evidence: treatment-group customers consumed significantly more during the peak periods of non-event weekdays than control-group customers, consistent with an effort to raise the reference level against which rebates would later be paid, and roughly half of the apparent demand response on event days reflected this inflated baseline rather than actual consumption reduction (Wolak, 2007). The mechanism also undermines investments in energy efficiency, since a more efficient appliance reduces consumption at baseline-setting times as well as at event times and therefore reduces the rebate the customer can claim.

The second problem is adverse selection in participation. Customers whose consumption would have been low during event hours for reasons unrelated to the program—a vacation, a holiday, normal schedule variation—collect rebates for reductions they would have made anyway, while customers whose consumption is genuinely peak-coincident and could be moved by the program may not enroll. Empirical work has found that a large share of paid demand reducers are free riders in this sense: in the Anaheim experiment, even control-group customers, who received no information about events and faced no rebate incentive, would have qualified for rebates roughly six times larger than the predicted reduction in consumption attributable to the program, indicating that most of what such a program pays for would have happened anyway (Wolak, 2007).

The third problem is the asymmetry itself. The effective marginal cost of consumption changes discontinuously at the baseline: below baseline, the cost of an additional kilowatt-hour is the retail rate plus the foregone rebate; above baseline, it is just the retail rate. Random variation in consumption is rewarded, since the customer collects rebates on low days but is not charged for high days. A customer with predictable consumption patterns who averages 45 kWh on event days collects a smaller rebate than a customer whose consumption is volatile but averages the same.

These problems do not mean peak-time rebates have no role. They can be useful as a transitional tool, and they may be the most politically feasible vehicle in jurisdictions where higher prices during peak events are unacceptable. But they should be understood as a substantially weaker instrument than critical peak pricing, and weaker still than time- or location-varying prices that apply symmetrically to all consumption.

9.5 Real-time pricing

Real-time pricing passes through hourly or sub-hourly price variation. It is the closest retail approximation to pricing electricity based on what it costs to serve one more kilowatt-hour at that time and place (Borenstein, 2005). It is most useful when customers have automation or when large customers can actively manage load. For households, real-time pricing may need bill protection,

hedging products, default automation, or opt-in structures to address risk and complexity. Recent estimates suggest that real-time pricing with conservative price caps—binding in roughly five percent of hours—can recover more than ninety percent of the efficiency gain available from uncapped real-time pricing, with caps that average between roughly $28 and $72 per megawatt-hour across US wholesale markets (Hinchberger et al., Forthcoming). These caps are an order of magnitude below typical wholesale market caps, suggesting that the customer-protection rationale for flat rates is much weaker than commonly assumed once modest price caps are available.

9.6 Customer charges and connection charges

Customer charges recover costs associated with connection, metering, billing, customer service, and residual cost recovery. A higher fixed charge allows a lower volumetric rate, improving marginal incentives. But uniform fixed charges can be regressive because low-income customers tend to consume less electricity and therefore benefit less from reductions in volumetric prices (Borenstein, 2012; Borenstein, Fowlie, and Sallee, 2021).

9.7 Demand charges

Demand charges bill customers based on maximum kilowatt demand during a period. Traditional non-coincident demand charges are common for commercial and industrial customers but less common for residential customers. They can be problematic if they penalize a customer's individual peak even when that peak does not coincide with system or local network peaks (Hledik and Greenstein, 2016). Much of the historical use of demand charges has not been about cost causation at all: a non-coincident demand charge offered a workable way to recover residual costs from large customers, where a flat per-customer fixed fee would have been meaningless across customers of widely different sizes. Borenstein (2019) argues that demand charges as traditionally implemented have done this badly, imposing costs on customers based on the timing of their individual peak regardless of whether the system was constrained at that moment, and that dynamic pricing would do the cost-causation work better. This critique applies most forcefully to non-coincident demand charges used as scalable fixed-cost recovery. Coincident peak demand charges are a distinct instrument, better aligned with cost causation and capturing local distribution capacity costs that wholesale dynamic prices do not.

9.8 Subscription and capacity-based tariffs

Some tariffs charge customers for a subscribed level of capacity, similar to a broadband plan. Customers pay for the right to draw up to a specified amount of power, with penalties or higher charges if they exceed it. These designs may become more relevant for customers with EVs, batteries, heat pumps, or large flexible loads.

9.9 Net metering, net billing, and buy-all/sell-all

Net metering credits exports at the retail rate, often allowing exports to offset imports one-for-one over a billing period. Net billing credits exports at an administratively determined export rate, often closer to avoided cost. Buy-all/sell-all arrangements require the customer to buy all

consumption at the retail rate and sell all generation at a separate compensation rate. The key economic question is whether the export compensation reflects the value of exports to the system, not whether the electrons are generated on the customer's roof.

9.10 Standby, backup, and reservation charges

Standby charges apply to customers with self-generation who remain connected to the grid. They are intended to recover the cost of backup service, reserved capacity, and network availability. These charges are especially relevant for large customers such as data centers, campuses, hospitals, and industrial facilities with behind-the-meter generation.

Figure 4 compares these tariff designs across five dimensions. A time signal indicates whether the tariff communicates that the cost of electricity varies across hours. A peak signal indicates whether the tariff charges customers more for consumption during system or local peak periods. Cost reflective captures whether the tariff recovers each cost category through an instrument that matches its cost driver. Simplicity reflects how easy the tariff is for customers and regulators to understand and administer. DER compatible measures whether the tariff sends appropriate price signals to customers with rooftop solar, batteries, electric vehicles, or on-site generation.

Tariff design comparison
●●● = strong ●● = moderate ● = weak — = absent or poor

Tariff design	Time signal	Peak signal	Cost reflective	Simplicity	DER compatible
Flat volumetric	—	—	—	●●●	—
Inclining block	—	—	●	●●	—
Time-of-use	●●	●	●●	●●	●●
Critical peak pricing	●●	●●●	●●	●●	●●
Real-time pricing	●●●	●●●	●●●	●	●●●
Demand charge (NCP)	—	●	●	●●	●
Coincident peak	●	●●●	●●●	●	●●●
Net metering (retail)	—	—	—	●●●	—
Net billing (avoided)	●●	●	●●●	●●	●●●
Standby/reservation	—	●●	●●●	●●	●●●

Figure 4: Tariff design comparison. Each design is rated on whether it provides a time-varying price signal, a peak-demand signal, cost-reflective pricing, administrative simplicity, and compatibility with distributed energy resources. No single design scores highest on all dimensions; the efficient tariff combines time-varying energy prices with coincident peak charges and transparent fixed-charge cost recovery.

The table reveals a core tradeoff: the simplest tariffs—flat volumetric rates and retail-rate net metering—score poorly on every efficiency dimension. The most cost-reflective designs—real-time pricing, coincident peak charges, avoided-cost net billing—are more complex but send signals that become increasingly important as the grid decarbonizes and distributed resources proliferate.

No single instrument scores highest on all dimensions. The efficient tariff is not one row in the table; it is a combination: time-varying energy prices for the energy component, coincident peak or capacity charges for network and generation capacity, avoided-cost export compensation for distributed generation, and transparent fixed or income-adjusted charges for residual costs.

10 Other Countries Are Not Waiting

Electricity tariff reform is not only a US issue. Countries with high renewable penetration, widespread smart-meter deployment, or rapid growth in distributed energy resources are confronting the same basic problem: traditional retail tariffs recover too much shared system cost through flat volumetric charges and provide too little information about when and where electricity is valuable. The institutional details differ, but the direction of reform is similar across many jurisdictions.

10.1 European Union: dynamic pricing as a consumer right

The European Union has moved furthest in making dynamic retail pricing part of the formal consumer-rights framework. Directive (EU) 2019/944 gives eligible consumers the right to a dynamic electricity price contract from at least one supplier and from suppliers with more than 200,000 customers, while also requiring that customers be informed about the opportunities and risks of such contracts (European Parliament and Council of the European Union, 2019; Council of European Energy Regulators, 2020). This does not mean all customers must take real-time prices. Rather, the EU model treats dynamic pricing as an available option within a broader consumer-protection regime that also includes fixed-price contracts, comparison tools, switching rights, smart-meter provisions, and protections for vulnerable customers.

The EU experience is useful because it separates two questions that are often conflated in US debates. First, should retail prices be capable of reflecting wholesale market conditions? The EU answer is increasingly yes. Second, should households be forced to bear unmanaged wholesale-price risk? The EU answer is no. Dynamic prices can coexist with consumer protection, contract choice, risk disclosure, and targeted support for vulnerable customers.

10.2 United Kingdom: settlement reform as the foundation for time-varying tariffs

The United Kingdom provides an example of a different route to tariff reform: changing settlement first. Ofgem's market-wide half-hourly settlement reform is intended to move electricity settlement toward measured half-hourly consumption, rather than relying on estimated load profiles for many customers (Office of Gas and Electricity Markets, 2016, 2021). This matters because suppliers have stronger incentives to offer time-varying retail products when they themselves are settled based on customers' actual time-varying consumption.

The U.K. case highlights an important implementation lesson. Dynamic retail tariffs require more than smart meters. They also require settlement systems, data access rules, supplier incentives, customer communication, and consumer protections. Without those institutional changes, the value of advanced metering may not flow through to retail prices.

10.3 Australia: cost-reflective network tariffs and two-way pricing

Australia is one of the most important laboratories for tariff reform because it has high rooftop solar adoption and significant distribution-network challenges. The Australian Energy Regulator describes network tariff reform as a process of moving distribution tariffs toward prices that more accurately reflect the cost of serving customers, including time-of-use tariffs and other cost-reflective structures (Australian Energy Regulator, 2024). The motivation is closely related to the argument in this primer: distributed energy resources such as solar, batteries, and electric vehicles can be integrated more efficiently if tariffs better reflect network costs.

Australia has also been at the frontier of debates over export pricing and two-way tariffs. In areas with high rooftop solar penetration, exports during sunny low-load hours can create local network costs, while exports during constrained or high-value periods can provide benefits. A cost-reflective tariff, therefore, may need to price both imports and exports (Simshauser, 2016). This is the logical extension of the principle that customers should be paid for the value they provide and charged for the costs they impose.

A distinctive feature of the Australian model worth flagging is architectural: cost-reflective network tariffs are sent to retailers rather than translated directly into customer-facing rates. Retailers then decide how to package those signals into products their customers will receive and respond to—some passing through the variation in full, others offering flatter products and absorbing the residual risk. The arrangement is designed to address a tension this primer has emphasized throughout: complex prices are difficult for customers to interpret, but simplifying them at the regulatory level discards information that the system needs. Delegating the translation to a competitive retail layer is a plausible institutional answer to that tension. It is also the same arrangement implicit in the fixed-price-contract discussion in Section 8: a retailer offering a flat rate to a customer paying network tariffs that vary with time is, in effect, selling an insurance product on top of the underlying signal. Whether this architecture works in practice depends critically on the strength of retail competition. Where competition is robust, retailers have incentives to design products that match customers' actual preferences for risk and complexity, and the insurance premium embedded in flat rates is disciplined by entry. Where it is not—which describes most of the United States outside the few states with active retail markets—there is no competitive process to deliver the same result, and the design problem reverts to the regulator. The Australian experiment is most valuable as a demonstration of what becomes possible when the institutional preconditions are in place; it is not a turnkey solution where they are absent.

10.4 California: moving from retail-rate net metering toward net billing

California illustrates the politics and economics of distributed-generation compensation. The state moved from earlier net energy metering structures toward a Net Billing Tariff adopted by the California Public Utilities Commission in December 2022 and applicable to new interconnection applications beginning in April 2023 (California Public Utilities Commission, 2022). Under net billing, exports are compensated separately from imports rather than simply offsetting consumption at the full bundled retail rate.

The conceptual rationale is straightforward. The value of an exported kilowatt-hour is not necessarily equal to the bundled retail rate. California's avoided-cost framework includes components

such as generation energy, generation capacity, ancillary services, transmission and distribution capacity, and decarbonization policy compliance (California Public Utilities Commission, 2025). That is the right conceptual object: exports should be credited based on the costs they actually avoid, ideally varying by time and location.

California also shows why tariff reform is politically difficult. Customers and firms made investments under earlier rules. Solar advocates worry that lower export compensation will slow adoption. Consumer advocates worry about affordability. Utilities and regulators worry about cost shifts and grid reliability. These concerns are real, which is why transitions need to be gradual, transparent, and paired with protections for vulnerable households. California's experience also highlights that high volumetric rates are themselves a major barrier to equitable electrification (Borenstein, Fowlie, and Sallee, 2022).

10.5 Lessons from international experience

The international experience suggests five lessons for US policymakers.

1. **Tariff reform is a system reform, not just a rate change.** Dynamic tariffs require smart meters, data systems, settlement rules, customer education, and supplier incentives.
2. **Dynamic prices and consumer protection can coexist.** Customers can be offered time-varying prices while retaining access to fixed-price contracts, bill protection, risk disclosure, or targeted assistance.
3. **Network tariffs matter as much as energy prices.** In high-DER systems, the key constraint may be distribution capacity, hosting capacity, or local peak demand rather than bulk energy cost (Pérez-Arriaga et al., 1995; Schittekatte and Meeus, 2020).
4. **Exports should be priced, not simply netted.** Netting imports and exports at the bundled retail rate hides the fact that imports and exports can have different system values.
5. **High rooftop solar penetration forces the issue.** Jurisdictions with substantial distributed generation eventually confront the mismatch between volumetric cost recovery and grid reliance.

The broad pattern is that jurisdictions are moving away from simple flat volumetric tariffs and toward more granular pricing of energy, network use, exports, and flexibility. The pace differs, and the politics are difficult everywhere. But the economic pressure is common.

11 Equity Means Protecting People, Not Bad Prices

Efficient rate design is not simply a matter of moving all fixed costs into uniform fixed charges (Borenstein, 2012; Burger et al., 2019; Borenstein, Fowlie, and Sallee, 2021). That would improve marginal incentives but could harm low-income customers. Because lower-income households typically consume less electricity, they benefit less from reductions in volumetric rates (Borenstein, 2012; Burger et al., 2019). A uniform fixed charge can therefore increase their bills even if it improves economic efficiency.

This creates a genuine policy tradeoff. The current system often uses distorted marginal prices to achieve implicit redistribution. That redistribution may have been politically convenient, but it

is increasingly poorly targeted. It helps high-consuming low-income households less than intended, burdens low-consuming households in ways that may not match ability to pay, and becomes less progressive when high-income customers can install rooftop solar or otherwise reduce net purchases. Evidence on energy poverty shows that the geography of energy burden is shifting, and that federal assistance programs have not kept pace with these changes (Batlle et al., 2024).

The better approach is to address equity explicitly (Burger et al., 2019; Borenstein, Fowlie, and Sallee, 2021). Several designs are available:

- **Income-graduated fixed charges.** Customers with higher incomes pay higher fixed charges, while low-income customers pay lower fixed charges. This preserves better marginal energy prices while addressing ability to pay (Borenstein, Fowlie, and Sallee, 2021, 2022).
- **Low-income discounts or credits.** Eligible households receive bill credits, fixed-charge discounts, or percentage discounts funded transparently.
- **Lifeline allowances.** A basic level of consumption can be protected through targeted credits, while marginal prices above that level remain cost-reflective.
- **Geographic or housing-based proxies.** Where income data are not available, regulators can use census geography, housing type, or program participation as imperfect but administratively feasible proxies.
- **Tax funding for broad policy costs.** Costs associated with climate policy, energy efficiency, low-income assistance, or electrification can be funded through general revenues rather than electricity rates.

The principle is that distributional policy should be visible and targeted. Hiding redistribution in the volumetric electricity rate is increasingly costly because it distorts electrification, overcompensates self-generation, and shifts costs to customers without the ability to avoid net purchases.

12 Nine Rules for Better Tariffs

A politically durable tariff reform should be gradual, transparent, and paired with protections for vulnerable households. The following principles can guide reform.

Principle 1: Energy prices should reflect what it costs to serve electricity at that time and place

Retail energy prices should move toward time-varying system-value signals. For some customers, this may mean real-time pricing. For others, it may mean time-of-use rates, critical peak pricing, or automated demand response. The important step is to stop pretending that all kilowatt-hours have the same system value.

Principle 2: Grid costs should be recovered according to cost causation where feasible

Costs caused by system peaks, local constraints, or reserved capacity should be recovered through charges tied to those drivers. Coincident peak charges are generally preferable to simple non-

coincident demand charges because they better reflect the customer's contribution to system costs.

Principle 3: Residual costs should not be hidden in volumetric rates

Residual costs should be recovered transparently through fixed charges, connection charges, subscription charges, income-adjusted charges, or non-rate funding. The cents-per-kilowatt-hour price should not be used as the default bucket for every cost that needs to be recovered.

Principle 4: Distributed generation exports should be compensated at avoided cost

Export compensation should reflect the value of exports to the system. That value varies by time, location, and system conditions. Full retail net metering is usually not cost-reflective when the retail rate includes fixed grid and policy costs (Pérez-Arriaga, Jenkins, and Batlle, 2017; California Public Utilities Commission, 2025).

Principle 5: Grid connection has option value

Customers with behind-the-meter generation should pay for the value of grid connection, backup service, and reserved capacity (Pérez-Arriaga, Jenkins, and Batlle, 2017; Knittel, Senga, and Wang, 2025). This applies to rooftop solar households, commercial campuses, industrial facilities, and data centers. Net-zero energy consumption is not the same as zero grid reliance.

Principle 6: Flexibility should be rewarded, primarily through prices

If energy and capacity prices fully reflect time and location, flexibility is rewarded automatically. A customer who shifts consumption away from a high-price hour avoids the high price. A customer whose usage avoids the system's coincident peak avoids the corresponding capacity charge. A customer who exports during a scarcity event receives the scarcity price. These efficiency-enhancing behaviors are incentivized through location-specific real-time prices. Layering separate flexibility payments on top of correctly priced energy and capacity risks double-counting the savings the customer already captures and can distort investment in flexible technologies. This principle is often missed, with industry observers believing customers must be separately paid for the consumption they forgo when the system is constrained. But that expectation has no parallel in other markets. When a freeze tightens the orange supply, prices rise, and consumers buy fewer oranges; no one suggests we should also pay consumers for the oranges they choose not to buy. The price increase is itself the reward for reducing consumption when supply is tight.

Supplementary compensation for flexibility is appropriate in the case that prices do not accurately reflect costs on either a temporal or geographic basis. While retail prices remain flat or weakly time-varying, interruptible tariffs, demand response programs, and peak-time rebates can partially substitute for the missing price signal. It is important to note, however, these quantity-based instruments carry their own incentive problems—particularly moral hazard in baseline construction, adverse selection in participation, and asymmetric incentives at the baseline—that make them substantially weaker substitutes than they may appear. For system services that energy prices

structurally do not capture—frequency regulation, voltage support, fast ramping—ancillary services markets remain the right vehicle. Where local distribution-level constraints create value that wholesale prices do not reflect, location-specific incentives may also be warranted. The first-best answer is to fix the prices and reserve quantity-based programs for the gaps that prices cannot fill. This applies as much to data centers and other large flexible loads as to residential customers (Knittel, Senga, and Wang, 2025).

Principle 7: Equity should be explicit

Rate reform should not ignore affordability. But affordability is better addressed through targeted support than through distorted marginal prices (Borenstein, Fowlie, and Sallee, 2021). Regulators should pair tariff reform with income-based protections, bill credits, and transition assistance.

Principle 8: Transitions should be gradual and predictable

Sudden tariff changes create political backlash and investment risk. Regulators can use shadow bills, opt-in pilots, default automation, transition credits, grandfathering with sunset dates, and phased implementation.

Principle 9: Regulators should measure avoided costs, not guess at them

If exports are going to be compensated at avoided cost, regulators need a transparent way to estimate avoided cost (California Public Utilities Commission, 2025). Public utility commissions should regularly conduct avoided-cost or value-of-DER studies that estimate the value of distributed exports and flexible load by time, location, and system condition. These studies should include energy, losses, congestion, generation capacity, transmission and distribution capacity, reliability, emissions, and other policy-relevant costs where appropriate.

The goal is not to find a number that is favorable or unfavorable to rooftop solar (California Public Utilities Commission, 2025). The goal is to identify what costs the system actually avoids when a customer exports electricity or reduces load. California's avoided-cost framework provides one useful model: not necessarily the exact answer for every state, but the kind of institutional machinery regulators need if export credits, battery incentives, demand-response payments, and flexible-load tariffs are going to be based on system value rather than on the bundled retail rate or political bargaining.

13 How to Get From Here to There

The question is not whether to implement the theoretical first-best tariff tomorrow (Burger et al., 2019; California Public Utilities Commission, 2025). The question is how to move from today's tariff toward a more efficient and equitable structure.

A practical sequence might look like this:

1. **Unbundle the bill for analytical purposes.** Regulators should require utilities to report how much revenue is collected for energy, capacity, transmission, distribution, policy programs, customer costs, and residual costs.

2. **Quantify marginal and residual costs.** Utilities should estimate which costs vary with energy consumption, which are peak-driven, which are customer-specific, and which are residual.
3. **Introduce or expand time-varying rates.** Begin with opt-in time-of-use or critical peak pricing, then move toward default time-varying rates where smart meters and customer protections are in place.
4. **Conduct regular avoided-cost or value-of-DER studies.** Regulators should not move from retail-rate net metering to arbitrary export credits. They should estimate the value of distributed exports and flexible load using a transparent methodology. These studies should report how value varies by hour, season, location, and system condition, and should separate energy value from capacity value, transmission and distribution value, emissions value, and other policy values. The result should inform export compensation, battery incentives, demand-response payments, and flexible-load tariffs.
5. **Reform export compensation.** Replace full retail net metering with avoided-cost export compensation, ideally varying by time and location. Preserve reasonable transition rules for existing customers but avoid indefinite grandfathering.
6. **Create explicit fixed-cost recovery mechanisms.** Shift residual costs out of volumetric rates and into transparent charges, paired with low-income protections.
7. **Design standby and reservation charges for self-supplying large loads.** Data centers and other large customers with behind-the-meter generation should pay for firm service, backup rights, and transmission reservation if they require them.
8. **Reward flexibility through prices, with programs for the gaps.** Time- and location-varying energy and capacity prices reward flexibility automatically. Interruptible tariffs, demand response programs, flexible interconnection, and performance-based compensation are most valuable as transitional substitutes while retail prices remain flat, and as long-run instruments for services energy prices do not capture, such as ancillary services and local distribution-level value.
9. **Move broad public-policy costs to broader funding sources where possible.** If a program serves a public objective rather than an electricity-system objective, policymakers should consider tax funding or dedicated public funds rather than volumetric electricity charges.

14 Conclusion: The Grid Is More Valuable Than Ever

The rise of rooftop solar, batteries, electric vehicles, heat pumps, and data centers does not make the grid obsolete. The grid is not becoming less important. It is becoming more important—and harder to pay for correctly. A household with solar still needs the grid when the sun is down, when clouds pass, when equipment fails, and when it wants to export. A data center with on-site generation may still need the grid as backup, as an export platform, and as a source of reliability insurance. A renewable-heavy power system needs flexible demand, storage, and price-responsive load more than a conventional system does.

The tariff problem is that current prices often treat the grid as if it were simply a commodity supplier of kilowatt-hours. It is more than that. It is a reliability platform, a balancing mechanism, a

delivery network, an insurance product, and increasingly a market platform for distributed resources (Joskow, 2007; MIT Energy Initiative, 2016).

A modern electricity tariff should therefore charge customers for what they use, pay them for what they provide, and fund shared infrastructure in a way that is transparent, efficient, and fair. That requires moving away from a tariff built around flat volumetric charges and toward a tariff that distinguishes energy, capacity, grid access, exports, flexibility, and residual cost recovery.

A Glossary

Avoided cost

The cost the system avoids when a customer reduces consumption or exports electricity. Components can include energy, capacity, losses, congestion relief, emissions, and deferred transmission or distribution investment, with values that vary by time and location. Pricing exports at avoided cost rather than at the bundled retail rate is the central reform separating net billing from traditional net metering. (See Section 9.)

Behind-the-meter generation

Generation located on the customer's side of the retail meter, such as rooftop solar, a campus generator, or on-site data-center generation. Customers with behind-the-meter generation may reduce their measured net purchases from the grid substantially while continuing to rely on it for backup, balancing, and exports. (See Section 6.)

Capacity charge

A charge intended to recover the cost of generation or network capacity built to meet peak demand. Capacity charges can be levied per kilowatt of contracted or coincident peak demand and are common in commercial and industrial tariffs and increasingly in emerging large-load tariffs. (See Section 3.2.)

Coincident peak

A customer's usage during the time when the system or local network reaches peak demand. Coincident peak is the cost-causation-relevant measure of a customer's contribution to capacity costs, in contrast to the customer's own individual non-coincident maximum demand. (See Section 3.2.)

Cost causation

The principle that customers should pay for the costs their consumption decisions impose on the system, and that each tariff instrument should be matched to the cost driver it is meant to recover. Cost causation is the organizing concept behind separating energy charges, peak-coincident charges, and residual fixed charges rather than bundling everything into the volumetric rate. (See Section 3.)

Cost shift The reallocation of grid, policy, or other residual costs from customers who reduce their measured net purchases to those who do not, even though the utility's underlying revenue requirement is unchanged. Cost shifts arise mechanically when residual costs are recovered through volumetric charges and some customers can reduce volumetric purchases (for example, through rooftop solar under retail-rate net metering) without proportionally reducing their reliance on the grid. (See Section 6.)

Critical peak pricing (CPP)

A tariff design that applies a substantially higher energy price during a small number of system-stress events each year, with prices and event hours typically announced day-ahead. CPP is symmetric in the sense that any consumption during an event

hour faces the high price, in contrast to peak-time rebates, which pay for reductions relative to a baseline. (See Section 9.4.)

Demand charge

A charge based on kilowatt demand, often the customer's maximum demand during a billing period. Non-coincident demand charges measure the customer's own peak regardless of when it occurs; coincident demand charges are tied to system or local peak periods and are better aligned with cost causation. (See Section 9.)

Demand response (DR)

Programs in which customers reduce or shift consumption during stressed system conditions in exchange for payment, lower rates, or other compensation. Demand response can substitute for time-varying prices when retail prices remain flat, but quantity-based DR programs carry incentive problems—moral hazard in baseline construction, adverse selection in participation, and asymmetric incentives—that price-based instruments avoid. (See Section 9.4.)

Distributed energy resource (DER)

A resource located on the distribution system or customer premises, such as rooftop solar, battery storage, demand response, or a controllable load. The growth of DERs creates pressure on tariff structures designed for one-way power flows and passive customers.

Distribution-level locational marginal price (DLMP)

A locational marginal price computed at the distribution-system level, reflecting the marginal cost of energy plus distribution-level congestion and capacity constraints local to a particular feeder or substation. DLMPs are largely theoretical at the residential and small-commercial level today; peak-coincident demand charges exist as the administratively feasible approximation to a DLMP regime. (See Section 3.2.)

Fixed charge

A charge that does not vary with kilowatt-hour consumption during the billing period. Fixed charges are well suited to recovering residual costs that do not vary with marginal consumption, though uniform fixed charges can be regressive because low-income customers tend to consume less electricity and therefore benefit less from any corresponding reduction in volumetric rates. (See Section 3.3.)

Income-graduated fixed charge (IGFC)

A fixed charge that varies with household income, set higher for high-income households and lower for low-income households. IGFCs are intended to reconcile the efficiency benefits of moving residual costs out of volumetric rates with concerns that uniform fixed charges burden low-consuming households. (See Section 11.)

Locational marginal price (LMP)

The marginal cost of serving one more unit of load at a specific location in an organized wholesale market, accounting for energy, congestion, and losses on the transmission

system. LMPs vary across hours and across pricing nodes, but they do not generally reflect distribution-level constraints. (See Section 3.2.)

Long-run marginal cost (LRMC)

The cost of meeting an additional unit of demand once capacity expansion is included, in contrast to short-run marginal cost, which holds capacity fixed. LRMC is the economically correct level for capacity-based charges intended to signal the cost of system expansion driven by peak demand. (See Section 3.2.)

Missing-money problem

The shortfall between the revenues that wholesale energy and ancillary services markets generate and the revenues required to fund the capacity the system commits to in advance. The shortfall arises because actual prices are suppressed below efficient scarcity levels by bid caps, flattened operating-reserve demand curves, out-of-market interventions, and political resistance to extreme prices. Capacity markets and standby or reservation charges exist in part to close this gap. (See Section 6.)

Net billing

A compensation structure in which exports from behind-the-meter generation are credited at a separate export rate, typically based on the system's avoided cost rather than at the full bundled retail rate. (See Section 9.)

Net metering

A compensation structure in which exports offset imports one-for-one at the full bundled retail rate over a billing period. Because the bundled retail rate typically includes more than the marginal cost of energy, net metering at the retail rate overcompensates exports relative to the value they deliver and contributes to cost shifts onto non-exporting customers. (See Section 6.)

Peak-time rebate (PTR)

A demand-response program in which customers receive a payment for reducing consumption during called events relative to a baseline estimated from their prior consumption. PTRs are asymmetric—customers are paid for reductions below the baseline but face no penalty for consumption above it—which creates moral hazard, adverse selection, and free-rider problems that pure price-based instruments avoid. (See Section 9.4.)

Ramsey-Boiteux pricing

A second-best pricing principle that, when residual fixed costs must be recovered through markups over marginal cost, allocates higher markups to demands with lower price elasticity. Ramsey-Boiteux logic is sometimes invoked to justify discounted rates for electrification end-uses with substitute fuels (such as EV charging), but the case for it weakens once residual costs are recovered through fixed or capacity-based instruments that allow energy prices to track marginal cost. (See Section 7.)

Real-time pricing (RTP)

A tariff design in which the retail energy price varies hour by hour, or more frequently, with wholesale market conditions, losses, and congestion. RTP is the closest retail approximation to pricing electricity based on what it costs to serve at the relevant time and place; in practice it is often paired with bill protection, hedging products, or price caps to limit residential exposure to extreme prices. (See Section 9.)

Reservation charge

A charge paid by a customer for the right to draw a specified amount of capacity from the grid, regardless of actual usage. Reservation charges are especially relevant for customers with behind-the-meter generation who rely on the grid as backup, since they price the option value of grid access rather than only metered consumption. (See Section 6.)

Residual costs

Costs that cannot be efficiently recovered through marginal energy prices or direct cost-causation charges. These include customer-specific costs, sunk historical investment, policy-program costs, and the portion of network revenue not recovered through cost-reflective peak charges. The central residential question is whether residual costs are recovered through volumetric charges, fixed charges, or some combination. (See Section 3.3.)

Scarcity pricing

Wholesale prices that rise sharply during periods when generation capacity or operating reserves are scarce, providing both short-run signals to flexible load and long-run signals supporting capacity investment. Suppression of scarcity prices through bid caps and out-of-market interventions is the main source of the missing-money problem. (See Section 6.)

Standby charge

A charge for customers with behind-the-meter generation who remain connected to the grid and require backup or reserved service. Standby charges price the insurance value of grid connection, which a customer who reduces measured net purchases would otherwise avoid paying for. (See Section 6.)

Subscription / capacity-based tariff

A tariff in which the customer pays a recurring charge for the right to draw up to a specified level of capacity, often combined with overage charges or restrictions for consumption above that level. Subscription tariffs may become more relevant as customers with EVs, batteries, heat pumps, and other large flexible loads put pressure on local distribution capacity. (See Section 9.)

Time-of-use (TOU) rate

A tariff with prices that vary across predetermined time blocks, such as off-peak, mid-peak, and peak. TOU rates are simpler than real-time pricing and easier to

communicate, but their predetermined blocks may miss the actual hours when system costs are highest. (See Section 9.)

Value of DER (VoDER)

A regulatory framework for valuing distributed energy resources by summing the components of avoided cost—energy, capacity, transmission and distribution, ancillary services, losses, emissions, and policy-relevant values—typically varying by time and location. VoDER methodologies (such as California's avoided-cost calculator) provide the analytical basis for setting export compensation under net billing. (See Section 9.)

References

Allcott, Hunt. 2011. "Rethinking Real-Time Electricity Pricing." *Resource and Energy Economics*, 33(4): 820–842.

Argosino, Fischer J. Espiritu, and Christopher R. Knittel. 2025. "Renewables and Electricity Affordability: Untangling Correlation from Causation." MIT Center for Energy and Environmental Policy Research Working Paper 2025-22.

Australian Energy Regulator. 2024. "Network Tariff Reform." *Australian Energy Regulator*, Accessed 2026-05-23.

Batlle, Carlos, Peter Heller, Christopher R. Knittel, and Tim Schittekatte. 2024. "U.S. Federal Resource Allocations Are Inconsistent with Concentrations of Energy Poverty." *Science Advances*, 10(41): 81–83.

Boiteux, Marcel. 1956. "Sur la gestion des monopoles publics astreints à l'équilibre budgétaire." *Econometrica*, 24(1): 22–40.

Bonbright, James C. 1961. *Principles of Public Utility Rates.* New York:Columbia University Press.

Borenstein, Severin. 2005. "The Long-Run Efficiency of Real-Time Electricity Pricing." *The Energy Journal*, 26(3): 93–116.

Borenstein, Severin. 2012. "The Redistributional Impact of Nonlinear Electricity Pricing." *American Economic Journal: Economic Policy*, 4(3): 56–90.

Borenstein, Severin. 2014. "Money for Nothing?" *Energy Institute Blog, UC Berkeley*, Available at https://energyathaas.wordpress.com/2014/05/12/money-for-nothing/.

Borenstein, Severin. 2019. "Are Demand Charges Fair?" *Energy Institute Blog, UC Berkeley*, Available at https://energyathaas.wordpress.com/2019/07/08/rethinking-demand-charges/.

Borenstein, Severin. 2025. "Should the Price of Electricity Depend On What You Use It For?" *Energy Institute Blog, UC Berkeley*, Available at https://energyathaas.wordpress.com/2025/03/10/should-the-price-of-electricity-depend-on-what-you-use-it-for/.

Borenstein, Severin, and James B. Bushnell. 2022*a*. "Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency." *American Economic Journal: Economic Policy*, 14(4): 80–110.

Borenstein, Severin, and James B. Bushnell. 2022*b*. "Headwinds and Tailwinds: Implications of Inefficient Retail Energy Pricing for Energy Substitution." *Environmental and Energy Policy and the Economy*, 3: 37–70.

Borenstein, Severin, Meredith Fowlie, and James Sallee. 2021. "Designing Electricity Rates for an Equitable Energy Transition." Energy Institute at Haas, UC Berkeley Working Paper 314.

Borenstein, Severin, Meredith Fowlie, and James Sallee. 2022. "Paying for Electricity in California: How Residential Rate Design Impacts Equity and Electrification." Energy Institute at Haas, UC Berkeley Working Paper 330.

Brown, Stephen J., and David S. Sibley. 1986. *The Theory of Public Utility Pricing.* Cambridge:Cambridge University Press.

Brown, Toby, Ahmad Faruqui, and Neil Lessem. 2018. "Electricity Distribution Network Tariffs: Principles and Analysis of Options." The Brattle Group. Prepared for the Victorian Distribution Businesses.

Burger, Scott P., Christopher R. Knittel, Ignacio J. Pérez-Arriaga, Ian Schneider, and Frederik vom Scheidt. 2019. "The Efficiency and Distributional Effects of Alternative Residential Electricity Rate Designs." MIT Center for Energy and Environmental Policy Research Working Paper 2019-002.

Bushnell, James, Benjamin F. Hobbs, and Frank A. Wolak. 2009. "When It Comes to Demand Response, Is FERC Its Own Worst Enemy?" *The Electricity Journal*, 22(8): 9–18.

California Public Utilities Commission. 2022. "Net Billing Tariff." *California Public Utilities Commission*, Decision D.22-12-056 adopted December 15, 2022; accessed 2026-05-23.

California Public Utilities Commission. 2025. "Distributed Energy Resources Cost-Effectiveness and Avoided Cost Calculator." *California Public Utilities Commission*, Accessed 2026-05-23.

Clinton, Bentley C., Christopher R. Knittel, and Konstantinos Metaxoglou. 2020. "Electrifying Transportation: Issues and Opportunities." MIT Center for Energy and Environmental Policy Research Working Paper 2020-010.

Council of European Energy Regulators. 2020. "Recommendations on Dynamic Price Implementation." *CEER Policy Paper, Ref. C19-IRM-020-03-14*, Accessed 2026-05-23.

European Parliament and Council of the European Union. 2019. "Directive (EU) 2019/944 of the European Parliament and of the Council of 5 June 2019 on Common Rules for the Internal Market for Electricity and Amending Directive 2012/27/EU." *Official Journal of the European Union*, Accessed 2026-05-23.

Hinchberger, Andrew J., Mark R. Jacobsen, Christopher R. Knittel, James M. Sallee, and Arthur A. van Benthem. Forthcoming. "The Efficiency of Dynamic Electricity Prices." *American Economic Journal: Economic Policy.*

Hledik, Ryan, and Gus Greenstein. 2016. "The Distributional Impacts of Residential Demand Charges." *The Electricity Journal*, 29(6): 33–41.

Houthakker, Hendrik S. 1951. "Electricity Tariffs in Theory and Practice." *The Economic Journal*, 61(241): 1–25.

Ito, Koichiro. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review*, 104(2): 537–563.

Jessoe, Katrina, and David Rapson. 2014. "Knowledge Is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review*, 104(4): 1417–1438.

Joskow, Paul L. 2007. "Regulation of Natural Monopoly." In *Handbook of Law and Economics*. Vol. 2, , ed. A. Mitchell Polinsky and Steven Shavell, 1227–1348. Elsevier.

Knittel, Christopher R., Juan Ramon L. Senga, and Shen Wang. 2025. "Flexible Data Centers and the Grid: Lower Costs, Higher Emissions?" National Bureau of Economic Research Working Paper 34065.

MIT Energy Initiative. 2016. "Utility of the Future: An MIT Energy Initiative Response to an Industry in Transition." Massachusetts Institute of Technology.

Office of Gas and Electricity Markets. 2016. "Electricity Settlement: Moving to Half-Hourly Settlement." *Ofgem guidance, last updated March 16, 2016*, Accessed 2026-05-23.

Office of Gas and Electricity Markets. 2021. "Market-Wide Half-Hourly Settlement: Full Business Case." *Ofgem*, Accessed 2026-05-23.

Pérez-Arriaga, Ignacio J., Francisco J. Rubio, Javier F. Puerta, Javier Arceluz, and Javier Marin. 1995. "Marginal Pricing of Transmission Services: An Analysis of Cost Recovery." *IEEE Transactions on Power Systems*, 10(1): 546–553.

Pérez-Arriaga, Ignacio J., Jesse D. Jenkins, and Carlos Batlle. 2017. "A Regulatory Framework for an Evolving Electricity Sector: Highlights of the MIT Utility of the Future Study." *Economics of Energy and Environmental Policy*, 6(1): 71–92.

Ramsey, Frank P. 1927. "A Contribution to the Theory of Taxation." *The Economic Journal*, 37(145): 47–61.

Schittekatte, Tim, and Leonardo Meeus. 2020. "Least-Cost Distribution Network Tariff Design in Theory and Practice." *The Energy Journal*, 41(5): 119–156.

Schittekatte, Tim, Dharik Mallapragada, Paul L. Joskow, and Richard Schmalensee. 2024. "Electricity Retail Rate Design in a Decarbonizing Economy: An Analysis of Time-of-use and Critical Peak Pricing." *The Energy Journal*, 45(3): 25–56.

Schittekatte, Tim, Ilan Momber, and Leonardo Meeus. 2018. "Future-Proof Tariff Design: Recovering Sunk Grid Costs in a World Where Consumers Are Pushing Back." *Energy Economics*, 70: 484–498.

Schweppe, Fred C., Michael C. Caramanis, Richard D. Tabors, and Roger E. Bohn. 1988. *Spot Pricing of Electricity.* Boston:Kluwer Academic Publishers.

Simshauser, Paul. 2016. "Distribution Network Prices and Solar PV: Resolving Rate Instability and Wealth Transfers Through Demand Tariffs." *Energy Economics*, 54: 108–122.

Turk, Graham, Tim Schittekatte, Pablo Duenas-Martinez, Paul L. Joskow, and Richard Schmalensee. 2025. "Designing Distribution Network Tariffs in the US with an Application to Increased Electric Vehicle Adoption." *The Energy Journal*, 46(6): 143–165.

Wolak, Frank A. 2007. "Residential Customer Response to Real-Time Pricing: The Anaheim Critical-Peak Pricing Experiment." Stanford University, Department of Economics Working paper.

Wolak, Frank A. 2011. "Do Residential Customers Respond to Hourly Prices? Evidence from a Dynamic Pricing Experiment." *American Economic Review*, 101(3): 83–87.

Contact.

The MIT CEEPR Research Commentary Series is published by the MIT Center for Energy and Environmental Policy Research from submissions by affiliated researchers.

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

General inquiries: ceepr@mit.edu

Media inquiries: ceepr-media@mit.edu

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

MASSACHUSETTS INSTITUTE OF TECHNOLOGY