Designing second-best price zones in electricity markets*

Jonas Boeschemeier† Sebastian Schwenen‡

Last updated: February 17, 2026

Abstract

In electricity markets, marginal costs vary substantially across space and time, implying welfare losses under spatially or temporally uniform pricing. In practice, prices are typically aggregated into large zones with spatially uniform prices. This paper develops an empirical framework to quantify the welfare loss of zonal pricing and to design welfare-improving price zones. We propose the spatial R^2 as a measure of spatial market efficiency and show that economically motivated clustering methods recover efficiency-maximizing zones. Applying the framework to three U.S. wholesale electricity markets, we demonstrate its computational tractability and find that existing zones are misaligned with current network conditions.

JEL-Classification: D40, L11, Q41, Q48

Keywords: Imperfect pricing, spatial pricing, electricity, networks

*We thank Alida Johannsen, Yannik Perez, Andrei Ovsiannikov, John Parsons, Paul Joskow, Christopher Knittel, Jasdeep Mandia, Karsten Neuhoff, Karen Pittel, Mar Reguant and Teodora Dobos for helpful comments and discussions. We are also grateful for valuable feedback from participants at YEEES 35 in Paris and at research seminars at the Massachusetts Institute of Technology, the ifo Institute and the Technical University of Munich.

†Technical University of Munich, Germany. Email: jonas.boeschemeier@tum.de

‡Technical University of Munich, German Institute for Economic Research DIW Berlin, and Mannheim Institute for Sustainable Energy Studies, Germany. Email: sebastian.schwenen@tum.de

1 Introduction

Designing efficient electricity markets is difficult. In North America, markets have undergone restructuring for over three decades, continuously adapting to emerging inefficiencies and technological change (e.g. Borenstein, 2002; Borenstein and Bushnell, 2015; Cramton, 2017; Wolak, 2019). One central design choice has been how to address network congestion, leading regulators to adopt either a nodal or zonal pricing framework; a discussion that is ongoing in other parts of the world as well (e.g. Green, 2007; Weibelzahl, 2017; Graf et al., 2020; Eicke and Schittekatte, 2022; Neuhoff et al., 2025).

Most liberalized wholesale markets in North America adopted nodal pricing, which delivered short-term efficiency gains relative to their zonal predecessors (e.g. Wolak, 2011; Zarnikau et al., 2014; Triolo and Wolak, 2022). Under this framework, generators are paid locational prices that reflect the marginal cost of delivering electricity at each location in the network. In theory, this ensures efficient spatial allocation and pricing (Bohn et al., 1984; Schweppe et al., 2013). In practice, however, nodal pricing is often combined with price zones, under which consumer prices at multiple network locations are aggregated at the zonal level, spanning broader geographic regions. Such aggregation can generate welfare losses when consumers respond to zonal prices that do not reflect local scarcity—a concern that may become more salient with the growth of flexible demand from data centers and other large consumers (e.g. Knittel et al., 2025; Norris et al., 2025). While price aggregation may be justified by technological or institutional constraints, many zonal configurations have remained unchanged since the initial adoption of nodal pricing, may no longer align with current network conditions, and are therefore likely to entail nontrivial welfare losses.

The existing literature has devoted considerable attention to measuring the welfare losses from temporally uniform prices, such as fixed retail rates (e.g. Borenstein and Holland, 2005; Hinchberger et al., 2024). Tractable methods to assess the efficiency of spatial price signals and their welfare consequences remain scarce. This paper develops an empirical framework to measure and minimize the welfare loss of price zones. Our approach builds on Jacobsen et al. (2020), who provide a theoretical framework for empirically estimating the welfare effects of imperfect pricing. In our setting, welfare losses arise because zonal prices do not fully reflect local heterogeneity in marginal costs and willingness to pay within a given price zone. As a result, inefficient consumption patterns may emerge, with overconsumption in some locations and underconsumption in others, generating deadweight losses and reducing overall market efficiency. Our empirical framework is designed to measure this *spatial market efficiency*, with higher efficiency corresponding to situations in which zonal prices obscure less of the underlying local variation in supply and demand.

Building on this intuition and drawing from Jacobsen et al. (2020), we introduce a

metric of spatial market efficiency, the *spatial* R^2, derived from the goodness-of-fit of linear regressions of locational prices on zonal indicators. To identify alternative spatial price configurations that minimize inefficiency, we employ economically motivated clustering methods to construct counterfactual, second-best zonal configurations. These zones yield second-best outcomes as they maximize market efficiency given an imposed number of zones. Our empirical approach is highly flexible, allowing counterfactual price zones to be constructed under a variety of institutional or geographic constraints and at modest computational cost.

We illustrate our methodology using data from three major U.S. wholesale electricity markets: New York (NYISO), New England (ISO-NE), and California (CAISO). First, we use our framework to assess the efficiency of currently implemented price zones. Using a large sample of hourly locational prices from 2020 to 2024, we find substantial variation in spatial market efficiency across these markets. Our results show that NYISO's, and to a lesser extent ISO-NE's, existing load zones perform relatively well in capturing spatial price variation. By contrast, the CAISO market displays low spatial efficiency, with a marked decline in recent years. This decline in CAISO is largely attributable to the small number of highly aggregated zones, which exhibit growing within-zone price differences.

In a second step, we identify counterfactual second-best zonal configurations. We employ a clustering algorithm that constructs price zones so as to minimize the deadweight loss. Put differently, the algorithm re-assigns locations while holding the number of zones in each market fixed. We find that adopting efficiency-maximizing zonal configurations would increase spatial market efficiency by 19%, 22% and 95% in NYISO, CAISO and ISO-NE, respectively.

Notably, a parsimonious approach that clusters locations solely on the basis of local prices—without incorporating explicit geographic information—nonetheless yields geographically coherent zones, in the sense that the resulting zonal delineations correspond to meaningful contiguous regions. Comparable efficiency gains can also be achieved by optimally subdividing the currently least efficient zones and thereby marginally increasing the total number of zones in a given market. However, the efficiency gains from additional zone splitting exhibit diminishing marginal returns: in practice, a relatively small number of zones is sufficient to capture the majority of price variation in networks comprising thousands of nodes.

Our findings relate to the broader economic literature on the welfare consequences of inefficient pricing in electricity networks. A central focus of this literature has been time-invariant retail tariffs, which set fixed retail prices for long periods while marginal costs of electricity typically differ by hour and sometimes by minute. The benefits of time-varying prices have been examined both theoretically (Borenstein and Holland, 2005; Joskow and Tirole, 2006) and empirically (Allcott, 2011; Fabra et al., 2021). More recently, Jacobsen et al. (2020) as well as Hinchberger et al. (2024) estimate the efficiency gains associated

with the introduction of time-of-use pricing, i.e., periods during which retail prices are held fixed, while Astier (2021) uses cluster analysis to study the optimal design of these pricing rules. The same intuition extends naturally to the spatial dimension that we study in this paper: just as aggregated retail tariffs obscure temporal cost variation, zonal prices obscure locational cost variation. We contribute to the literature by providing a framework for estimating the welfare losses from spatially inefficient pricing.

In so doing, our analysis contributes to the literature on the design of price zones in electricity markets. Studying European and North American markets, Moretto and Pollitt (2026) show that the efficiency of price zones can deteriorate over time as new demand and supply sources connect to the network, and argue for periodic and evidence-based reassessment to reflect evolving grid conditions. This concern has given rise to a large literature on the evaluation of price zone configurations (e.g., Burstedde, 2012; Felling and Weber, 2018; Dobos et al., 2025; ENTSO-E, 2025; Bushnell et al., 2015; Potomac Economics, 2025).[1] We contribute to this strand of research by proposing a micro-founded economic framework that delivers theory-driven efficiency estimates for both current and counterfactual zonal configurations while requiring minimal computational resources.

The trade-off between spatially differentiated prices and aggregating prices across larger regional zones is not unique to electricity markets. Analogous challenges arise in other network industries in which congestion pricing is relevant, most notably natural gas markets, telecommunications, and transportation.[2] The price aggregation approach we propose is flexible and readily adaptable across applications and market settings. In particular, we show that price zones can be efficiently clustered on the basis of local scarcity, while also allowing for the incorporation of geographic and institutional constraints to yield tractable and policy-relevant configurations tailored to the industry under consideration.

2 Context and data

This section presents the institutional setting and data. We introduce the spatial pricing regimes applied in U.S. electricity markets and summarize the data used for our empirical analysis.

2.1 Spatial pricing in U.S. electricity markets

In recent decades, most wholesale electricity markets in the United States have adopted a locational marginal pricing (LMP) market framework. Within LMP, market participants face locational prices that internalize the physical constraints of the electricity grid

[1] See Miraftabzadeh et al. (2023) for an overview of cluster analysis applications in power system analysis.

[2] See Vickrey (1969) for an early theoretical treatment of congestion pricing.

(Hogan, 2002). LMP, also termed nodal pricing, functions fundamentally different from zonal electricity markets and can result in up to thousands of different locational prices depending on the number of nodes in the grid. LMP is bid-based, and the market clears such that locational prices represent the marginal cost (if the market is competitive) of providing one additional unit of power at a specific location in the grid. In other words, LMP represents the first-best price of electricity at a certain location.[3]

To fix ideas, consider a given location j in the grid, where power is consumed and/or produced. Define the market price at this node by LMP_{jt} where t indexes time (say the hourly market price at location j). With LMP-based nodal pricing, the price and consequently the marginal costs of having an additional unit of power at location j during hour t are defined by

$$\text{LMP}_{jt} = energy_t + congestion_{jt} + losses_{jt}, \tag{1}$$

where the right-hand side specifies the three components of LMPs (Schweppe et al., 2013). The *energy* component does not have a location index and provides an estimate for the system-wide energy cost at that point in time.[4] In the LMP framework, any spatial differences in prices are thus due to two factors, namely congestion and transmission losses. The *losses* component of one location will be positive due to power losses on the transmission line if electricity travels long distances. The *congestion* component captures the shadow price of a congested transmission line at its capacity limit. When located at the start of a congested transmission line, i.e., the point that is feeding into the congestion, the congestion component will be negative to reflect the relative abundance of electricity at that location. Vice versa, when located at the end, i.e., on the constrained side where not all needed power can be delivered through the line, the congestion component will be positive to reflect the relative scarcity of electricity.[5] As we show below, congestion (rather than losses) drives most of the spatial price variation in U.S. electricity markets.

In theory, LMPs yield efficient prices for both producers and consumers. However, in practice most U.S. markets employ a hybrid model instead, in which producers face locational prices, while consumer (load) prices are aggregated across larger regions, which we refer to as price zones. Historically, this aggregation of LMPs into larger price zones has been implemented to protect consumers from locational price risks and to avoid price discrimination. With changing geographical load patterns, the discussion on potential reconfigurations of load zones is frequently reoccurring in many markets.[6]

Yet, most price zones still follow the same institutional borders from when LMP

[3]See Bohn et al. (1984) for an early theoretical analysis of LMP.

[4]In essence, the energy component of LMP is the price at an arbitrarily chosen reference node (reference bus) in the grid, and all congestion components are defined relative to that node.

[5]When transmission congestion is present, the congestion component of the LMP generally reflects the difference in marginal costs across the locations caused by the inability to freely deliver electricity.

[6]For the Texas electricity market see, e.g., Potomac Economics (2025).

was first implemented up to three decades ago. In single-state power markets, such as California or New York, price zones were typically determined based on the responsibility area of existing utility companies. In multi-state power markets, such as in New England, price zones largely follow state borders.

As the number of zones is much smaller than the total number of nodes in the grid, aggregating consumer prices at the zonal level may not reflect locational differences in marginal cost, hence leading to inefficiencies. The institutional setting of U.S. price zones serves as a useful application to assess the efficiency of zonal pricing and to examine alternative, more efficient zonal configurations.

2.2 Markets and data

To examine the welfare implications of price zones, we analyze three major power markets in North America, namely the New York (NYISO) market, the New England (ISO-NE) market, and the California (CAISO) market. We chose these markets because rich data are available on (i) locational marginal prices, (ii) the zonal assignment of nodes, (iii) zonal prices, and (iv) coordinates for individual nodes. The latter allow us to assess and illustrate the geographic distribution of zones. At the same time, each market has distinct characteristics that aid our analysis. NYISO and CAISO are both single-state markets, while ISO-NE covers multiple states. Moreover, the markets in our sample differ substantially in their number of nodes and zones. CAISO has the highest number of nodes and the fewest price zones, while the opposite holds for NYISO. For each market, we collect data covering the five-year period from 2020 to 2024. We next introduce each market and illustrate its number of nodes and zonal composition.

NYISO represents the power market of New York state and is divided into eleven price zones, termed load zones. For our dataset that spans five years, we gathered price and load zone information for 602 distinct nodes.[7] When considering only nodes that persist throughout the entire sample period and have associated geographical coordinates, the sample size reduces to 547 nodes. The price zones are shown in panel (a) of Figure 1 and largely follow the institutional borders of regional transmission owners. In sum, our final sample for the NYISO market consists of 547 nodes with eleven associated price zones.

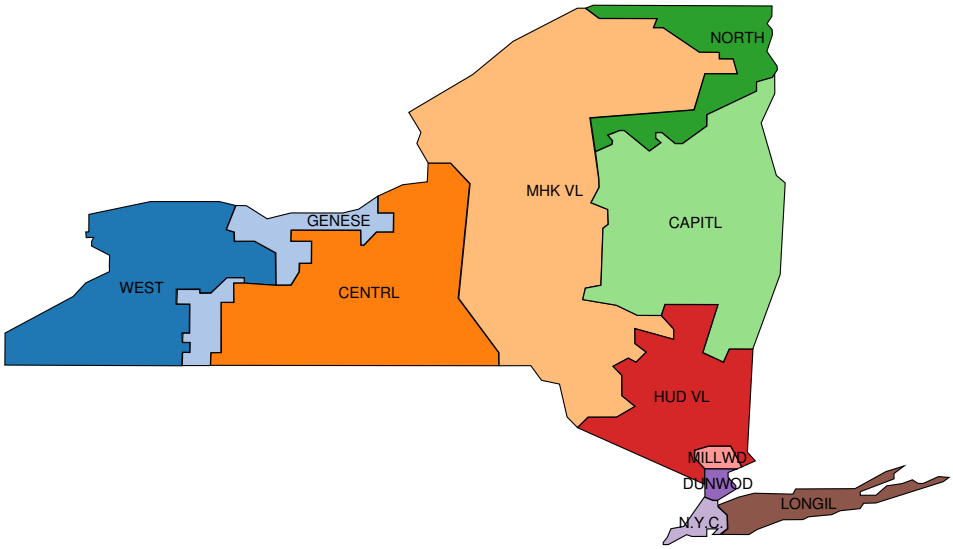
ISO-NE is the power market of New England and covers all six of its states. Each state represents one price zone, except for Massachusetts which is split into three, resulting in a total of eight price zones, also termed load zones. Our dataset includes 1,193 unique nodes and 1,005 nodes with price and coordinates data for the entire sample period.[8] Panel

[7]We match nodes to zones using the generator node list, downloaded from the NYISO website (accessed: July 2025).

[8]We match nodes to zones using the most recent pricing-node tables of that respective year, downloaded from the ISO-NE website (accessed: February 2025).

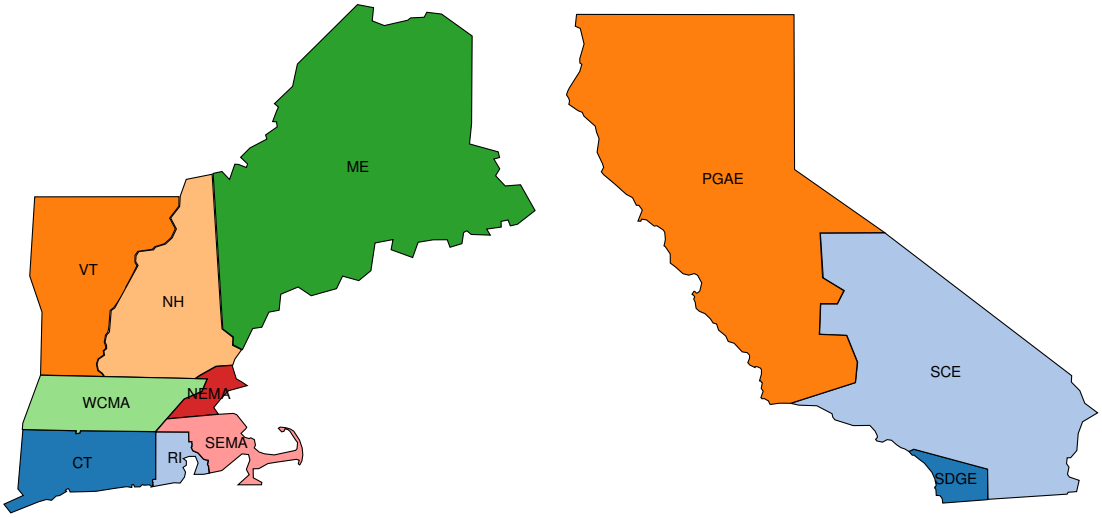
Figure 1: Price Zones of NYISO, ISO-NE and CAISO

(a) NYISO



(b) ISO-NE

(c) CAISO



Note: Zones are drawn based on illustrations and coordinates of nodes and state borders. The VEA zone in CAISO is omitted.

(b) of Figure 1 shows the price zones of ISO-NE. As shown, Massachusetts is split into three sub-zones: Boston, West Massachusetts, and South-East Massachusetts. In sum, our final sample for the ISO-NE market consists of 1,005 nodes with eight associated price zones.

CAISO represents the Californian power market and is the largest of the three in terms of the number of nodes. Our dataset consists of 2,258 pricing nodes, which are aggregated to four different price zones, termed default load aggregation points (DLAP) in

the CAISO market.[9] These price zones represent the geographic area covered by major utility companies in California. The coordinates are not directly provided by CAISO and are obtained by web-scraping the open source data platform of CAISO.[10] We drop pricing nodes with no information on their coordinates. Among others, all 19 nodes that belong to the price zone of the Valley Electric Association (VEA) do not have publicly available coordinates. Our final sample includes 1,398 nodes. Panel (c) of Figure 1 displays the three remaining price zones covered in our analysis: Pacific Gas & Electric Company (PGAE), Southern California Edison Company (SCE) and San Diego Gas & Electric (SDGE). In sum, our final sample for the CAISO market consists of 1,398 nodes with three associated price zones.[11]

We collect LMPs (i.e, locational hourly day-ahead prices) from 2020 to 2024 across the three markets introduced above.[12] In total, our final sample consists of about 24 million LMP observations for the NYISO market, about 44 million LMP observations for the ISO-NE market, and about 61 million LMP observations for the CAISO market.

2.3 Measuring spatial price variation

A key metric in our analysis is the degree to which zonal prices explain—or fail to explain—variation in locational marginal prices. We measure total variation in LMPs across both space and time using the total sum of squares (SST), defined as

$$\text{SST} = \sum_t \sum_j (p_{jt} - \bar{p})^2, \tag{2}$$

where p_{jt} denotes the LMP at location j and time t, and $\bar{p}$ is the overall mean LMP across all locations and times.

To isolate the *spatial* variation in LMPs, we decompose the SST into its spatial and temporal components:

$$\text{SST} = \text{SST}_{\text{spatial}} + \text{SST}_{\text{temporal}} = \underbrace{\sum_t \sum_j (p_{jt} - \bar{p}_t)^2}_{\text{SST}_{\text{spatial}}} + \underbrace{\sum_t \sum_j (\bar{p}_t - \bar{p})^2}_{\text{SST}_{\text{temporal}}}. \tag{3}$$

The spatial SST captures within-period spatial variation in LMPs around the time-t spatial mean, $\bar{p}_t$. The temporal SST captures between-period variation in these spatial

[9] We match nodes to zones using the node to DLAP mapping, downloaded from the CAISO OASIS platform (accessed: June 2025).

[10] Our web-scraping code is based on the code provided by Eric Munsing.

[11] When including nodes that have no coordinates data available our sample consists of 2,177 nodes, excl. VEA. We run robustness tests using this larger sample. Our results remain virtually unchanged.

[12] We collect the data using the `Gridstatus` API in Python. When daylight-saving time causes repeated hours, two prices are recorded for the same hour; we aggregate these as averages. Price data for CAISO on 14 April 2021 are not available.

means around the overall mean LMP, thereby abstracting from spatial heterogeneity.

Finally, to assess the efficiency of price zones in capturing spatial price variation we assign each location j to a zone n and further decompose the spatial SST as

$$\text{SST}_{\text{spatial}} = \underbrace{\sum_t \sum_j (p_{jt} - \bar{p}_{nt}(j))^2}_{\text{within-zone variation}} + \underbrace{\sum_t \sum_j (\bar{p}_{nt}(j) - \bar{p}_t)^2}_{\text{between-zone variation}}. \tag{4}$$

The first component captures within-zone variation of LMPs by measuring the deviation of individual LMPs p_{jt} from their corresponding zonal average $\bar{p}_{nt}(j)$, where $\bar{p}_{nt}(j)$ is the simple spatial average of all LMPs belonging to the same zone as j at time t. This term reflects spatial variation of LMPs within zones at a given time. The second term represents between-zone variation by comparing the zonal averages $\bar{p}_{nt}$ to the market-wide average $\bar{p}_t$ at each time t. Together, these two terms account for the spatial dimension of volatility. Appendix A.1 derives the above decompositions in full. When introducing our empirical strategy further below, we demonstrate that the SST decomposition is useful not only for identifying sources of price volatility but also for providing a straightforward interpretation of the welfare losses associated with price zones.

Table 1 reports statistics on the spatial and temporal variation of LMPs in 2024, the final year of our sample. To compare levels of variation, we compute the standard deviation of each SST component introduced above.[13] In Panel A, we provide statistics on the NYISO market. The average LMP across all nodes and hours is 37.55 USD/MWh. The overall standard deviation of 25.6 USD/MWh indicates substantial price variation. Following equation (3), we decompose this volatility into a temporal component of 23.38 USD/MWh and a spatial component of 10.43 USD/MWh. Hence, a large share of the total price variation originates from temporal differences rather than spatial differences. Using equation (4), the spatial component can be further divided into between-zone and within-zone variation, with respective standard deviations of 8.56 and 5.95 USD/MWh. This suggests that the majority of the spatial price variation in NYISO arises between zones, while prices within individual zones remain more homogeneous. Re-examining the individual components of LMP, given by equation (1) above, the congestion term exhibits a spatial standard deviation of 9.49 USD/MWh, compared to only 2.13 USD/MWh for losses. Therefore, congestion remains the dominant source of spatial price differences in NYISO during 2024.

Panel B reports the corresponding statistics for the ISO-NE market. The average LMP is slightly higher at 41.1 USD/MWh, with an overall standard deviation of 30.1 USD/MWh. Similar to NYISO, most volatility stems from temporal differences, reflected in a temporal standard deviation of roughly 30 USD/MWh. Spatial variation, by contrast,

[13] Standard deviations are obtained by dividing each SST component by the number of observations (hours $\times$ nodes) and taking the square root.

is relatively minor at 2.69 USD/MWh, indicating that prices across locations are much more uniform in ISO-NE. The remaining spatial dispersion is evenly distributed between and within zones, with both standard deviations around 2 USD/MWh. Looking at the underlying components, congestion accounts for most of this spatial variation, with a spatial standard deviation of 2.15 USD/MWh, compared to 1.28 USD/MWh for losses. Hence, although spatial price differences exist, they remain limited in magnitude and largely driven by localized congestion.

Panel C shows the results for CAISO. Average prices are again comparable at around 40 USD/MWh, yet the market exhibits considerably higher volatility overall, with a standard deviation of 36.32 USD/MWh. As in the other markets, most volatility is temporal (31.09 USD/MWh), but spatial variation is also substantial at 18.77 USD/MWh. Unlike NYISO and ISO-NE, most of this spatial variation arises within zones rather than between them, as indicated by the within-zone standard deviation of 17.49 USD/MWh. This points to a relatively high degree of locational heterogeneity even inside existing zones. In terms of components, the pattern is consistent with the other markets: congestion explains most of the spatial variation, whereas losses play only a minor role.

Table 1: Mean and standard deviation of LMP and its components in 2024

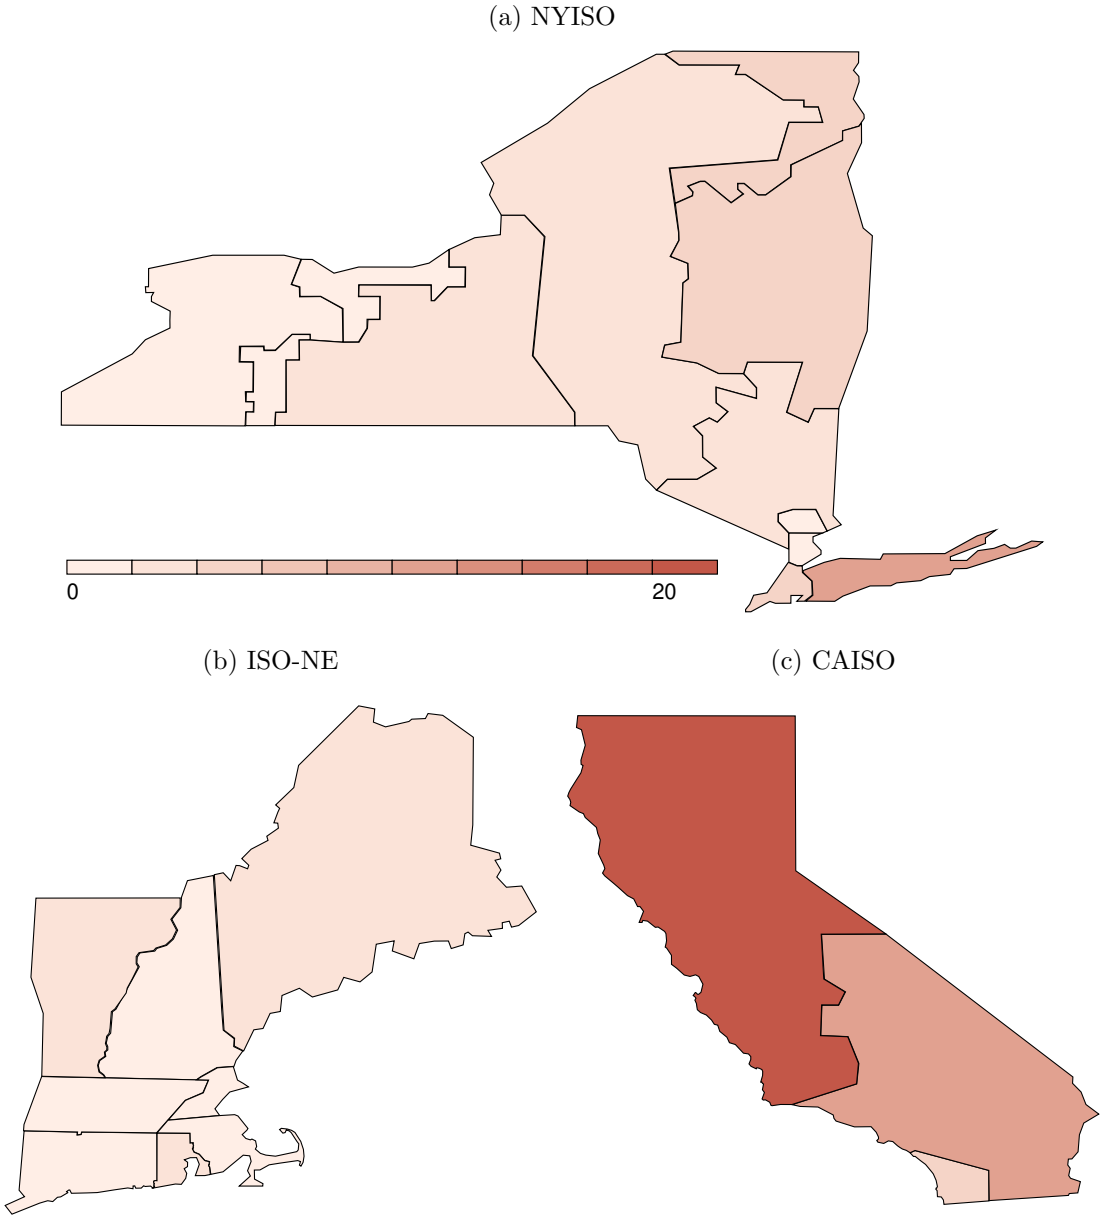
		Standard Deviation				
	Mean	Overall	Temporal	Spatial	Between	Within
Panel A: NYISO						
LMP	37.55	25.60	23.38	10.43	8.56	5.95
Congestion	2.40	11.99	7.34	9.49	7.49	5.82
Losses	1.31	2.31	0.90	2.13	2.02	0.69
Panel B: ISO-NE						
LMP	41.10	30.10	29.98	2.69	1.79	2.01
Congestion	-0.07	2.17	0.28	2.15	1.23	1.76
Losses	-0.25	1.31	0.28	1.28	1.01	0.78
Panel C: CAISO						
LMP	40.78	36.32	31.09	18.77	6.80	17.49
Congestion	2.54	19.23	5.54	18.42	6.55	17.21
Losses	0.11	2.01	0.43	1.97	0.79	1.80

Note: The sample includes all locations with price, zone, and coordinate data in 2024. The standard deviations are obtained by dividing each SST component in equations (3) and (4) by the number of observations (hours x nodes) and taking the square root.

Figure 2 illustrates the geographic distribution of within-zone price variation across the three markets, where darker colors indicate higher within-zone standard deviations. Panel (a) displays results for NYISO, where the overall within-zone standard deviation is 5.95 USD/MWh, as reported in Table 1. Most zones exhibit similar and relatively

moderate levels of internal price variation. By contrast, the Long Island zone and parts of northeastern New York display substantially higher within-zone LMP dispersion. In Long Island, the within-zone standard deviation reaches 11.8 USD/MWh—approximately twice the NYISO-wide average, suggesting persistent local congestion.

Figure 2: Within-zone price variation across sample markets in 2024



Note: Colors indicate the standard deviation of within-zone prices in USD/MWh. The sample includes all locations with price, zone, and coordinate data in 2024. The zone VEA in CAISO is omitted.

Panel (b) presents results for the ISO-NE market, which generally exhibits very low spatial price volatility. The highest within-zone standard deviation occurs in Maine at 3.4 USD/MWh, a level that remains modest relative to the other markets. Panel (c) presents the CAISO market, where within-zone price variation is both substantial and spatially

uneven. Northern California (the PG&E zone) exhibits a high within-zone standard deviation of 19.5 USD/MWh in 2024—nearly double that of SCE (10.8 USD/MWh) and more than triple that of SDG&E (5.4 USD/MWh). These results indicate pronounced locational heterogeneity within CAISO's zonal design and suggest that some zones may be poorly aligned with underlying grid constraints. Consistent with this interpretation, the next section shows that efficiency losses of zonal pricing are directly proportional to the magnitude of within-zone price deviations. Consequently, higher within-zone volatility signals lower spatial efficiency.

3 Methodology

This section describes our methodology. We introduce a model framework in which prices of different products are aggregated into bundles and formulate the planner's problem to derive the second-best optimal product bundles and bundle prices. We then show how the solution to the planner's problem can be estimated empirically in the context of price zones in electricity markets.

3.1 Theoretical setup

Consider an economy that consists of J different products. The quantity demanded of each product $j \in (1, ..., J)$ is denoted by x_j. The marginal costs of supplying product j, denoted $mc(x_j)$, may differ across products and are increasing in quantity, i.e., $\partial mc_j/\partial x_j > 0$. While producers are remunerated based on their product-specific marginal cost, consumer prices are aggregated across products due to institutional or technological constraints. In our electricity market setting, the different products correspond to electricity produced at different locations, aggregated product bundles correspond to price zones, and producers receive location-specific prices, while consumers pay an average zonal price.

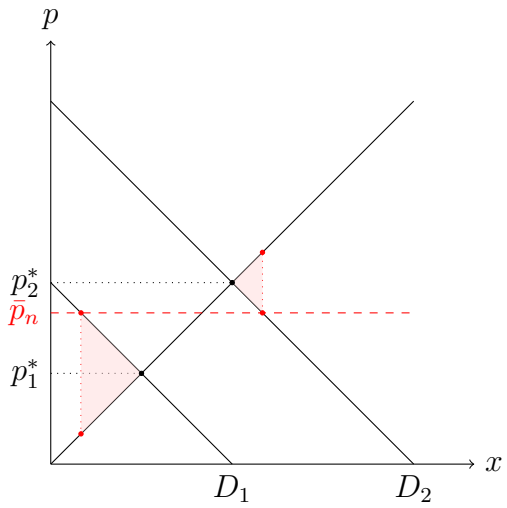
To ensure market clearing, there exists a market operator that buys products at marginal cost from producers and sells them at an aggregated price to consumers. The market operator hence functions as an intermediary between producers and consumers and can, broadly, also be interpreted as a retailer. The aggregated prices are set by the market operator for a group of products, which we refer to as bundles.

We index bundles by $n \in (1, ..., N)$ with $1 \leq N \leq J$. Each product j belongs to exactly one non-empty bundle n and we denote the set of products belonging to bundle n by $\mathcal{Z}_n$. Rather than setting product-specific prices, the market operator aggregates all products in N bundles and sets a uniform price within each bundle, so that the price of product j, p_j, equals the price of its corresponding bundle, $p_j = \bar{p}_n$ for all $j \in \mathcal{Z}_n$, where $\bar{p}_n$ denotes the bundle-n price.

First-best welfare is achieved if consumer prices equal marginal costs, i.e., if prices are set at the intersection of product-specific supply and demand. Yet, unless $N = J$, market clearing with bundle prices typically results in imperfect prices and welfare loss. We define the resulting deadweight loss from bundle pricing as the sum of Harberger triangles across all products (Harberger, 1964).

Figure 3 illustrates the deadweight loss for $J = 2$ products and $N = 1$ bundle. As shown, the bundle price $\bar{p}_n$ exceeds the first-best price for product $j = 1$, denoted as p_1^*. For product 1, this bundle price implies undersupply, and potential gains from trade remain unrealized. The opposite holds for product $j = 2$, where the marginal cost exceeds the marginal benefit, and market efficiency would improve if the traded quantity were reduced. In both cases, the magnitude of the efficiency loss from aggregate pricing depends on two factors: the price wedge $\bar{p}_n - mc_j(x_j(\bar{p}_n))$, i.e., the difference between marginal benefit and cost, and the corresponding deviation in quantity $x_j(\bar{p}_n) - x_j(p_j^*)$.

Figure 3: Illustration of deadweight loss of arbitrary bundle price



To formalize the deadweight loss, we follow Jacobsen et al. (2020) and assume locally linear supply and demand functions. Under this assumption, Harberger triangles are locally triangular areas, as illustrated in Figure 3. Furthermore, we assume independent products with zero cross-price elasticities.[14]

Using the above setup, the total deadweight loss from bundle pricing for J products and given some arbitrary first-best price vector $(p_1^*, \ldots, p_J^*)$ can be written as

$$\text{DWL}(\bar{p}_n \mid p_1^*, \ldots, p_J^*) = -\frac{1}{2} \sum_j \big(\bar{p}_n(j) - mc_j(\bar{p}_n)\big)\big(x_j(\bar{p}_n) - x_j(p_j^*)\big), \tag{5}$$

where $\bar{p}_n(j)$ is the bundle price of product j and $mc_j(\bar{p}_n)$ is shorthand notation for

[14] In our application, this assumption is not overly restrictive, as products correspond to a large number of grid locations and producers are technically constrained to inject and sell electricity at their respective locations.

$mc_j(x_j(\bar{p}_n))$. For two products and one bundle, equation (5) reduces to summing up the Harberger triangles in Figure 3.

As shown in Appendix A.2, the deadweight loss in equation (5) can be rearranged to

$$\text{DWL}(\bar{p}_n \mid p_1^*, \ldots, p_J^*) = -\frac{1}{2}\sum_j (\bar{p}_n(j) - p_j^*)^2 \omega_j, \tag{6}$$

where the deadweight loss is computed as the squared difference between the bundle price and the first-best prices weighted by $\omega_j = \frac{\partial x_j}{\partial p_j}\left(1 - \frac{\partial mc_j}{\partial x_j}\frac{\partial x_j}{\partial p_j}\right) \leq 0$. Put differently, equation (6) expresses the deadweight loss as a function of the difference between the bundle price and the first-best price weighted by the slopes of demand and supply. To provide intuition, consider the case of constant marginal costs where ω_j reduces to the price derivative of demand, $\partial x_j / \partial p_j$.[15] In this scenario, the deadweight loss corresponds to the price wedge multiplied by the resulting change in quantity demanded. When marginal costs are increasing, however, the wedge between marginal benefit and marginal cost widens; the weights ω_j capture the fact that a given price deviation generates a larger welfare loss.

3.2 The planner's problem

We consider a market operator that seeks to implement the second-best outcome by minimizing the deadweight loss, given an exogenously imposed number of product bundles. The market operator's problem involves two decisions. First, the bundle choice: the operator must assign all J products into N different bundles. Second, the pricing choice: the operator sets an optimal price for each bundle.

Pricing choice. The market operator sets bundle prices to minimize the resulting deadweight loss from equation (6). Formally, given a predetermined set of bundles, the market operator's minimization problem becomes

$$\min_{\bar{p}_1,\ldots,\bar{p}_N} -\frac{1}{2}\sum_j (\bar{p}_n(j) - p_j^*)^2 \omega_j. \tag{7}$$

Solving for the bundle prices $\{\bar{p}_n\}_{n=1}^N$, the first-order condition takes the form

$$\sum_{j \in \mathcal{Z}_n} \left(\bar{p}_n - p_j^*\right) \omega_j = 0. \tag{8}$$

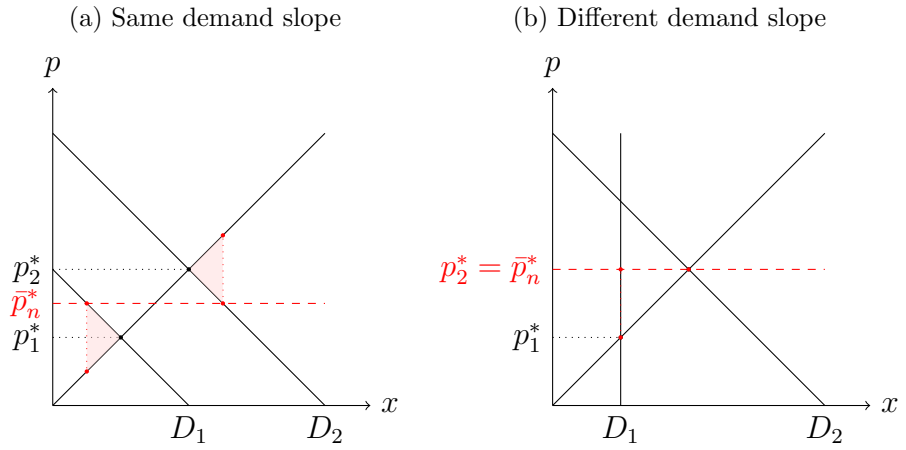
As shown, optimal bundle prices net out the marginal welfare losses associated with aggregate pricing, which arise from overconsumption in some products and underconsumption in others.

[15] This is the case studied in Astier (2021) and Jacobsen et al. (2020).

Figure 4 illustrates the intuition behind the optimal bundle price using a two-product example. In panel (a), marginal costs are identical across both products, and demand differs only by a constant while maintaining the same slope. Because both goods are therefore weighted equally, the optimal bundle price is simply the average of the first-best prices, i.e., the mean of p_1^* and p_2^*. This results in two Harberger triangles of equal size—one due to underconsumption and one due to overconsumption. Any deviation from this bundle price would reduce overall welfare, as the aggregate deadweight loss across both products would increase.

Panel (b) illustrates the case of different demand slopes, with demand for product 1 being perfectly inelastic. In this scenario, imperfect bundle pricing does not cause efficiency losses, because the quantity consumed of product 1 remains unchanged regardless of the bundle price. In fact, the product weights ω_j show that perfectly inelastic demand corresponds to a weight of zero. Hence, in the two-product example of panel (b), the optimal bundle price is simply equal to the price of product 2. In this case, the bundle price coincides with the first-best outcome.

Figure 4: Illustration of deadweight loss in pricing choice



Finally, rearranging equation (8) yields the optimal bundle price

$$\bar{p}_n^* = \frac{\sum_{j \in \mathcal{Z}_n} \omega_j p_j^*}{\sum_{j \in \mathcal{Z}_n} \omega_j}, \tag{9}$$

which is the weighted average of first-best prices of all products belonging to the bundle. As such, the market operator's pricing problem is a typical application of Ramsey-Bouteux pricing (Ramsey, 1927; Boiteux, 1956), in which a product with high price sensitivity, either through supply or demand, is given a relatively higher weight in calculating the aggregate price.[16]

[16] Borenstein and Holland (2005) and Joskow and Tirole (2006) derive similar expressions for optimal retail electricity prices, where the optimal time-invariant retail price is determined as the weighted average of wholesale prices across time periods.

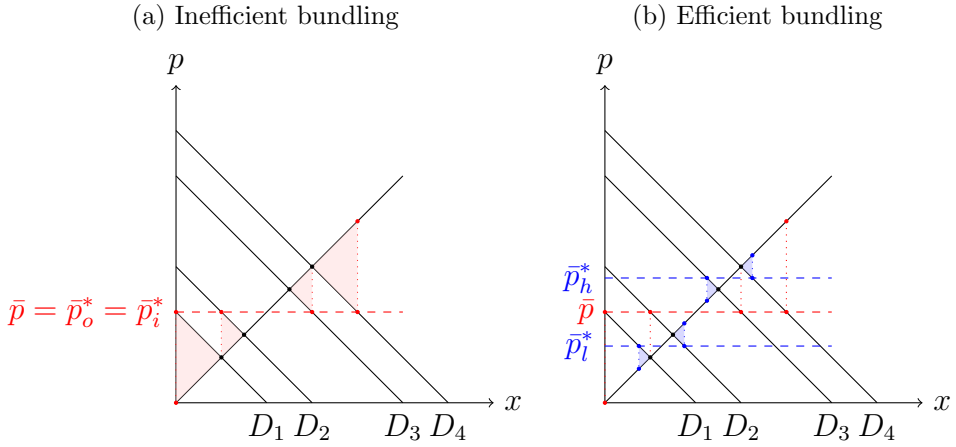
Bundle choice. So far, we considered the pricing problem of the market operator given a predetermined bundle structure. Next, we study the bundle choice, where the market operator is constrained by the number of bundles but is free to assign products to bundles. Formally, the market operator solves

$$\min_{\mathcal{Z}_1,\ldots,\mathcal{Z}_N} -\frac{1}{2}\sum_{n=1}^{N}\sum_{j\in\mathcal{Z}_n}(\bar{p}_n^*(j)-p_j^*)^2\omega_j \tag{10}$$

by assigning each product j to a set of products $\mathcal{Z}_n$ to which a common bundle price $\bar{p}_n$ applies. Notice that this assignment endogenously determines the optimal bundle price in equation (9): by choosing bundles the market operators minimizes the squared difference between the optimal bundle prices and the first-best product prices. In essence, the market operator bundles products with similar first-best prices together to keep the intra-bundle price variations—and their corresponding welfare losses—minimal.

Figure 5 serves as an illustrative example of how the bundle choice affects welfare losses. Figure 5 plots four different products with different demand that need to be grouped in two different bundles. In panel (a), the operator bundles the two 'outer' and the two 'inner' products together, $\mathcal{Z}_o = \{1,4\}$ and $\mathcal{Z}_i = \{2,3\}$. This results in two bundle prices, $\bar{p}_o$ for the outer bundle and $\bar{p}_i$ for the inner bundle. In this example, these prices coincide with the uniform price $\bar{p}$ that would arise if all four products were bundled together. Clearly, this bundling is inefficient, as splitting into two bundles provides no efficiency gain relative to a single aggregate price.

Figure 5: Illustration of deadweight losses due to allocation choice



Now consider an alternative pricing approach where the two products with low demand and the two with high demand are bundled together, $\mathcal{Z}_l = \{1,2\}$ and $\mathcal{Z}_h = \{3,4\}$. The deadweight loss under this pricing strategy is highlighted by the blue-shaded Harberger triangles in panel (b) and is substantially smaller than in panel (a). The welfare gain from this bundling, relative to the alternative bundle structure in panel (a), is indicated

by the red-dashed area. Consistent with this intuition, the joint price and bundle choice shown in equation (10) identifies the second-best outcome that minimizes the deadweight loss, constrained by the exogenously given number of bundles.

3.3 Empirical strategy

Next, we implement the market operator's problem empirically. We first estimate optimal bundle prices and their associated welfare losses using ordinary least squares regression, and then apply a clustering approach to implement the operator's optimal bundle choice, which minimizes welfare losses for a given number of bundles.

Pricing Choice. To empirically estimate optimal bundle prices and their corresponding efficiency, we draw from the methods developed in Jacobsen et al. (2020). To fix ideas, consider the regression

$$p_j^* = \sum_{n=1}^{N} \beta_n D_j^n + \epsilon_j, \tag{11}$$

where p_j^* are the first-best prices and D_j^n is an indicator variable set to one if product j is part of bundle n and zero otherwise. When estimating equation (11) without a constant and using ordinary least squares, the estimated coefficients for β_n immediately yield a simple average price for each bundle. Assuming that the weights ω_j in equation (9) are identical across products or uncorrelated with the residuals, the estimated coefficients represent the empirical estimate of the optimal bundle price $\bar{p}_n^*$ in equation (9). In case of heterogeneous and correlated weights, optimal bundle prices can still be empirically estimated using weighted least squares (Jacobsen et al., 2020), which we illustrate as a robustness test in the results section.

Note that the residual term of this regression can be expressed as $\bar{p}_n^*(j) - p_j^*$, so that $\sum_j (\bar{p}_n^*(j) - p_j)^2$ corresponds to the residual sum of squares. Similarly, if we write the total sum of squares as $\sum_j (\bar{p} - p_j)^2$, where $\bar{p}$ denotes the overall average price (i.e, the uniform price if all products were bundled together), the standard expression for the R^2 of regression (11) is $R^2 = 1 - \frac{\sum_j (\bar{p}_n^*(j) - p_j^*)^2}{\sum_j (\bar{p} - p_j^*)^2}$. Using the deadweight loss expression in equation (6) and assuming that the weights ω_j cancel out, the R^2 can be re-written as

$$R^2 = 1 - \frac{\sum_j (\bar{p}_n^*(j) - p_j^*)^2}{\sum_j (\bar{p} - p_j^*)^2} = 1 - \frac{\text{DWL}(\bar{p}_1^*, \ldots, \bar{p}_N^*)}{\text{DWL}(\bar{p})}. \tag{12}$$

The usefulness of this approach lies in the interpretation of the goodness-of-fit measure. As shown in equation (12), the standard R^2 formula can quantify the relative efficiency gains when implementing imperfect bundle prices $(\bar{p}_1^*, \ldots, \bar{p}_N^*)$ relative to a uniform price $\bar{p}$. An R^2 equal to one indicates that the deadweight loss under a given pricing

rule is zero; that is, the numerator is zero, bundle prices perfectly explain the individual product prices, and replicate the first-best outcome. Conversely, an R^2 of zero suggests that the deadweight loss of bundle prices is identical to that under a uniform price. In this case, the bundle prices are just as inefficient as assigning a single aggregate price to all products.

The R^2 can thus serve as a relative measure for short-term market efficiency, which can be readily obtained without requiring any information on the underlying demand structure (Jacobsen et al., 2020). Intuitively, because the fitted values (i.e., the bundle prices) are obtained by minimizing the sum of squared residuals between the bundle prices and the first-best prices, they also minimize the associated welfare loss for a given number of bundles. The degree to which they minimize welfare loss is proportional to the R^2.

In our empirical application, the first-best prices are the locational marginal prices (LMPs), which efficiently signal local scarcity. The bundle prices correspond to zonal prices that aggregate LMPs across grid locations and may therefore introduce inefficiencies. Our analysis seeks to identify the optimal bundle prices, that is, the second-best zonal prices and their welfare effects.

LMPs vary across both time and space. We therefore extend the above framework and regress LMPs on interactions between zone and time dummies. The zone dummies serve as the bundle dummies in equation (11), yielding a zonal price equal to the average LMP within each zone. The time dummies disentangle temporal from spatial variation and ensure that we obtain optimal zonal prices conditional on time t. Our main specification therefore becomes

$$\text{LMP}_{jt} = \sum_{s=1}^{T} \sum_{n=1}^{N} \beta_{nt} D_{jt}^{s} \times D_{jt}^{n} + \epsilon_{jt}, \tag{13}$$

where D_{jt}^{s} is set to one if $t = s$ and D_{jt}^{n} is set to one if location j belongs to zone n. The estimated coefficient β_{nt} corresponds to the average of all LMPs in zone n at a given time period, hourly in our case. Thus, β_{nt} represents the empirical estimate of the second-best zonal price at hour t, $\bar{p}_{nt}^{*}$.

Note that the R^2 of this regression becomes

$$R^2 = 1 - \frac{\sum_t \sum_j (\bar{p}_{nt}^{*}(j) - p_{jt}^{*})^2}{\sum_t \sum_j (\bar{p} - p_{jt}^{*})^2} = 1 - \frac{\text{DWL}(\bar{p}_{11}^{*}, \ldots, \bar{p}_{NT}^{*})}{\text{DWL}(\bar{p})}, \tag{14}$$

which measures the efficiency gain of zonal prices relative to a uniform price $\bar{p}$ which is the same for all locations across space *and* time. The numerator denotes the residual sum of squares of regression (13) and is equivalent to the within-zone variation from the decomposition of LMP variation in equation (4). The denominator denotes the total sum of squares, which entails both spatial and temporal sum of squares. Hence, the R^2 approaches one if within-zone price variation is small compared to total variation

and, consequently, most price variation occurs between time periods and between zones. Conversely, if all price variation occurs within zones, the R^2 becomes zero.

In terms of market efficiency, the R^2 reflects the welfare gain from implementing zonal *time-variant* prices, as captured by the numerator, relative to a uniform *time-invariant* price. Therefore, the welfare gains measured in equation (14) may in part stem from an increase in the temporal granularity of prices.

To assess how well price zones capture *spatial* variation, we introduce a *spatial* R^2 measure that quantifies the extent to which zonal prices capture spatial inefficiencies net of any temporal effects. Formally, our main measure for the efficiency of zonal prices is

$$R^2_{spatial} := 1 - \frac{\sum_t \sum_j (\bar{p}^*_{nt}(j) - p^*_{jt})^2}{\sum_t \sum_j (\bar{p}^*_t - p^*_{jt})^2} = 1 - \frac{\text{DWL}(\bar{p}^*_{11}, \ldots, \bar{p}^*_{NT})}{\text{DWL}(\bar{p}^*_1, \ldots, \bar{p}^*_T)}. \tag{15}$$

The spatial R^2 measures the efficiency gains from using time-varying (hourly) zonal prices $(\bar{p}^*_{11}, \ldots, \bar{p}^*_{NT})$ relative to a time-varying (hourly) market-wide price that is uniform across locations. Recalling equations (3) and (4), the numerator captures the within-zone price variation, while the denominator captures total spatial price variation. Hence, if most spatial variation occurs between zones, within-zone variation is small relative to total spatial variation, the zonal price reflects within-zone scarcity well, and the spatial R^2 approaches one. The spatial R^2 equals one if hourly zonal prices coincide with first-best prices, in which case the numerator becomes zero. In Appendix A.3, we show that the spatial R^2 in equation (15) can be obtained by evaluating our main regression in specification (13), together with a second regression of LMPs exclusively on time dummies.

Finally, note that estimating equation (13) using ordinary least squares assumes that the weights ω_j in equation (9) are identical across locations or, alternatively, that they are uncorrelated with the residuals (Jacobsen et al., 2020). In our results section, we also present findings using heterogeneous weights, where we estimate weighted least squares regressions with load factors as weights ω_j. ISOs often aggregate prices using location-specific load factors, which effectively weight locations by demand and can serve as proxies for heterogeneity in demand and supply slopes. Using such information from the CAISO market, we illustrate the application of heterogeneous weights as a robustness test.

Bundle Choice. Next, we implement the empirical counterpart of the market operator's bundle choice in equation (10). The optimal bundle choice corresponds to the optimal assignment of LMPs to different pricing zones and identifies the second-best zonal configuration for a given number of zones. To do so, we use clustering techniques that group similar LMPs based on shared characteristics. As shown by Astier (2021) in the context of real-time pricing, the bundle-choice problem in (10) parallels the objective function of k-means clustering.

The k-means clustering algorithm partitions observations into k clusters to minimize the within-cluster sum of squares. The bundle choice problem in equation (10) implements an equivalent algorithm which partitions all LMPs into $k = N$ clusters, i.e., zones in our context. The number of zones N is exogenously determined by the market operator. The algorithm then assigns LMPs to zones so as to minimize the deadweight loss.

We implement the k-means algorithm as follows. The algorithm begins by randomly initializing N zones following the k-means++ procedure (Arthur and Vassilvitskii, 2007); each zone has distinct hourly zonal prices $\bar{p}_{nt}$ for every hour t. The algorithm then assigns each location to the zone with the lowest squared Euclidean distance $\sum_t (p^*_{jt} - \bar{p}_{nt})^2$, minimizing the within-zone price deviation. After each re-assignment, zonal prices $\bar{p}_{nt}$ are recalculated as the simple average of all locations linked to the zone, following equation (9) on the optimal bundle price. This procedure continues iteratively until (i) no locations are reassigned and (ii) the algorithm has converged to the minimal deadweight loss given (iii) second-best zones with optimal zonal prices $\bar{p}^*_{nt}$ for each zone n and hour t. In Appendix A.4, we provide the pseudo-code for this algorithm.[17]

Our main empirical analysis employs an unweighted k-means algorithm. However, similar to using a weighted least squares regression to obtain weighted optimal zonal prices as discussed above, a weighted k-means algorithm with heterogeneous weights can be applied (Astier, 2021). In our results section, we discuss the results of a weighted k-means algorithm for the CAISO market.

To analyze different potential objectives of the market operator, we implement three clustering algorithms that differ in their treatment of the underlying market's topology:

Clustering by price. First, we cluster by price and solve the market operator's problem in equation (10) by minimizing the within-zone variation of LMPs (through k-means clustering). Because all following approaches attach further constraints to the minimization problem, clustering exclusively by price serves as the upper bound for the attainable spatial R^2 and identifies the empirical efficiency-maximizing zonal configuration for any given number of zones.

Clustering by price and coordinates. Clustering solely on prices can yield zonal allocations that are infeasible to implement due to geographic or institutional constraints. Our second approach therefore incorporates both prices and spatial information by clustering on LMPs alongside their geographic coordinates (latitude and longitude). The aim is to construct zones that are geographically coherent, as relying on prices alone may group together nodes that are spatially distant. This method still maximizes the spatial R^2, but imposes an additional requirement that counterfactual zones remain ge-

[17] Because the k-means clustering algorithm is sensitive to its initial starting points that may lead to local minima, we employ the k-means++ algorithm of Arthur and Vassilvitskii (2007), which ensures well-distributed initial starting points. Moreover, we execute multiple runs of the algorithm to reduce the likelihood of local minima.

ographically meaningful, measured in terms of spatial proximity among locations. The full procedure and corresponding pseudo-code are provided in Appendix A.4.

Clustering by price and borders. This approach incorporates the institutional boundaries of the existing zones. Any counterfactual zone generated under this method must adhere to the current zonal borders. In the case of aggregation (reducing the number of zones) each new counterfactual zone must consist of one or more entire existing zones. Conversely, for disaggregation (expanding the number of zones) existing zones may be subdivided. In both cases, aggregation and disaggregation decisions rely exclusively on LMPs as input features. For aggregation, the pair of zones whose merger yields the smallest reduction in spatial R^2 is selected; for disaggregation, the split that produces the largest increase in spatial R^2 is chosen. The aim of this approach is to identify zonal configurations that are institutionally feasible, as utilities and transmission owners may be unwilling to accept arbitrary node reassignments. Thus, while this method retains the original objective of maximizing spatial R^2, it imposes the additional constraint that counterfactual zones must respect existing institutional boundaries. A detailed description and corresponding pseudo-code are provided in Appendix A.4.

4 Results

We start by presenting the deadweight loss results for the currently implemented zones. We then report results for counterfactual zones.

4.1 Deadweight loss of current zones

Table 2 shows the R^2 values, that is, our deadweight loss estimates, for the current zonal design in the three markets. To obtain these estimates, we run regressions using the entire five year sample of LMPs from 2020 until 2024.

The first row of Table 2 shows our main results for the spatial R^2. As shown in Panel A, the price zones in the NYISO market achieve a spatial market efficiency of 0.786. Within our theoretical framework, this estimate indicates that the currently implemented zonal prices reduce the deadweight loss by 78.6% relative to the deadweight loss that would arise under a *time-varying* market-wide price that is uniform across locations.

To better evaluate our estimates, the second row of Table 2 reports the standard R^2 from regressing LMPs on time and zone dummies, as specified in equation (13). This regression yields an R^2 of 0.959, indicating that the combination of time and zone dummies explains approximately 95% of the total variation in LMPs. Interpreted in terms of deadweight loss, hourly zonal prices therefore reduce welfare losses by about 95% relative to a uniform, *time-invariant* price.

Table 2: Market efficiency of current price zones in 2020-2024

	Panel A: NYISO			*Panel B*: ISO-NE			*Panel C*: CAISO		
	LMP	Congestion	Losses	LMP	Congestion	Losses	LMP	Congestion	Losses
R^2_{spatial}	0.786	0.767	0.880	0.373	0.287	0.566	0.289	0.287	0.153
R^2	0.959	0.876	0.917	0.996	0.298	0.585	0.928	0.410	0.202
R^2_{temporal}	0.808	0.470	0.309	0.994	0.016	0.043	0.899	0.172	0.059
Nodes		547			1,005			1,398	
Obs.		23,982,121			44,062,215			61,258,962	

Note: The sample includes all locations with price, zone, and coordinate data from 2020 to 2024.

The third row reports what we label the *temporal* R^2, which measures how well hourly time dummies alone explain variation in LMPs. As shown, time dummies—corresponding to hourly time-varying but spatially uniform prices—already explain 80.8% of the total variation in LMPs. Consequently, 19.2% of the variation remains unexplained by time-varying prices alone. An alternative way to assess the spatial efficiency of the NYISO zonal market design is to note that the inclusion of zonal prices explains 78.6% of this remaining variation.[18] Overall, these results indicate that the NYISO's zonal market design effectively mitigates welfare losses, as zonal prices capture a substantial share of the spatial variation in LMPs, that is, variation left unexplained by hourly-varying prices.

Panel A of Table 2 also reports analogous results for the individual components of LMPs for the NYISO market, namely congestion and losses. We obtain these results by estimating the same set of regressions, but replacing LMPs with their congestion and loss components. As shown, the spatial R^2 of congestion is 0.767, indicating that 76.7% of the spatial variation in congestion is captured by the current zonal prices. This result is not surprising, given that congestion accounts for most of the spatial variation in LMPs in NYISO (see Table 1). Accordingly, the results for LMPs and their congestion component closely align.

The remaining variation in LMPs arises from losses. The spatial R^2 for losses indicates that 88% of the spatial variation in transmission losses is captured by zonal dummies. Overall, the zonal structure in the NYISO market performs well in minimizing deadweight loss. This is because LMPs and their congestion component exhibit limited within-zone variation, so a common zonal price closely reflects supply and demand conditions. Moreover, losses also vary little within zones, implying that a common zonal price does not conceal substantial price differences arising from transmission losses.

Panel B of Table 2 reports the corresponding results for the ISO-NE market. The spatial R^2 indicates that zonal prices explain approximately 37% of the spatial variation in LMPs; equivalently, zonal pricing eliminates about 37% of the deadweight loss that

[18] We compute this as $\frac{R^2 - R^2_{temporal}}{1 - R^2_{temporal}} = \frac{0.959 - 0.808}{1 - 0.808} = 0.786$; see also Appendix A.3.

remains under spatially uniform prices. Spatial market efficiency is therefore substantially lower than in the NYISO market.

At the same time, the temporal R^2 of 0.994 in ISO-NE shows that hourly prices account for nearly all variation in LMPs, implying that there is relatively little spatial variation in the ISO-NE market to begin with. Analyses of the congestion and loss components yield similar conclusions to those for the NYISO market: results for congestion closely mirror those for overall LMPs, and 56% of the spatial variation in losses is explained by zonal prices, indicating that losses are again well represented by ISO-NE's current zonal structure.

In stark contrast, our results in Panel C show that price zones in the CAISO market achieve a spatial market efficiency of only 28.9%. Thus, the price zones in California do not adequately capture spatial price dynamics. Moreover, the temporal R^2 of 0.899 indicates substantial spatial variation in LMPs that is not fully explained by hourly prices. Taken together, this suggests that a more efficient zonal pricing structure is feasible in the CAISO market. Note also that the current price zones do not adequately explain variation in transmission losses, as zone dummies account for only about 15% of this variation. This finding, together with the relatively low spatial efficiency of zonal prices in the CAISO market, points to substantial deadweight loss that arises when zonal prices conceal high transmission losses in the CAISO market.

In sum, our results show that the price zones in the NYISO market and, to some extent, in the ISO-NE market perform reasonably well. In contrast, the results for the CAISO market suggest that a reconfiguration of price zones could improve efficiency and reduce the deadweight loss associated with zonal pricing.

Robustness. Our results remain robust when including all nodes in our sample, including those without available coordinate data. We present the results for this larger sample in Table B1 in Appendix B. For further robustness tests, we estimate the deadweight loss at the annual level, separately for each of the three markets and for each year between 2020 and 2024. These results are shown in Figure B1 in Appendix B and again corroborate the presence of considerable deadweight loss in the CAISO market. Furthermore, we find that inefficiencies in the CAISO market increase in the later years of the sample. We conjecture that this increase in the deadweight loss is related to the growth of new decentralized generation and load assets (CAISO, 2025), which are not adequately captured by the current price zone configuration.

Finally, we perform our baseline analysis using node-specific weights and a weighted least squares regression. We illustrate this approach using load factors from CAISO. Each location's load factor represents the share of total system load at that location.[19] Market

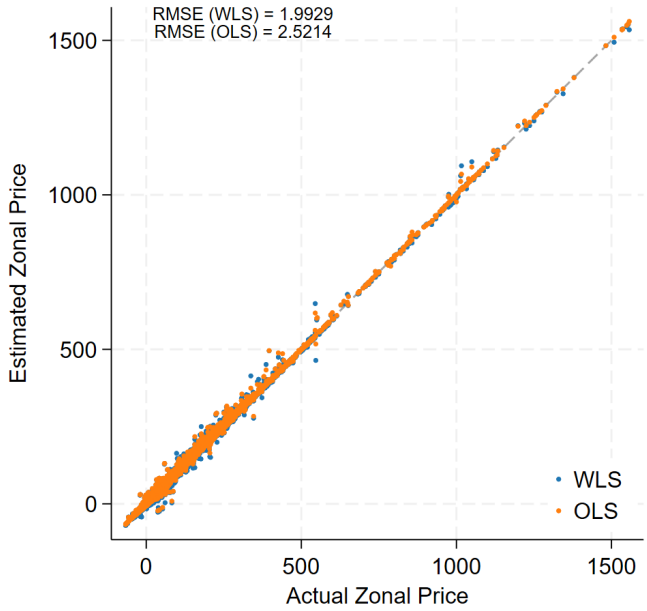
[19]Because load factors may vary over time as nodes are added to the grid, we compute the weight for each location j as the time average of its load factor.

operators commonly use load factors to compute load-weighted zonal prices in order to ensure revenue neutrality.[20] Accordingly, our analysis using weighted LMPs provides a benchmark consistent with practical market operations by incorporating the market operator's revenue-neutrality constraint.

Table B2 in Appendix B reports the results from the weighted least-squares estimation. When spatial deviations at high-load nodes are given greater weight, the spatial R^2 in the CAISO market increases to 0.52, compared to 0.29 in the baseline case, where all locations are weighted equally. Yet, the efficiency of zonal pricing in CAISO remains substantially below par relative to the NYISO market.

For an additional comparison of weighted and unweighted deadweight loss results, Figure 6 plots the estimated zonal prices from both approaches against the actual zonal prices observed in the CAISO market. As shown, both weighted and unweighted least squares regressions produce very similar estimated zonal prices, and in each case the estimates closely track the observed zonal prices plotted on the horizontal axis. Moreover, the root mean squared errors (RMSEs) relative to actual zonal prices are comparable in magnitude, at approximately 2 for the weighted least squares specification and about 2.5 for the unweighted least squares specification. We interpret these findings as evidence that assuming homogeneous weights does not materially distort the results.

Figure 6: Estimated zonal prices (weighted and unweighted) against actual zonal prices



Note: The sample includes all locations with price, zone, and coordinate data from 2020 to 2024.

[20] Analogously, Borenstein and Holland (2005) show that for retailers, offering time-invariant retail prices, the equilibrium retail price is demand-weighted.

4.2 Counterfactual second-best zones

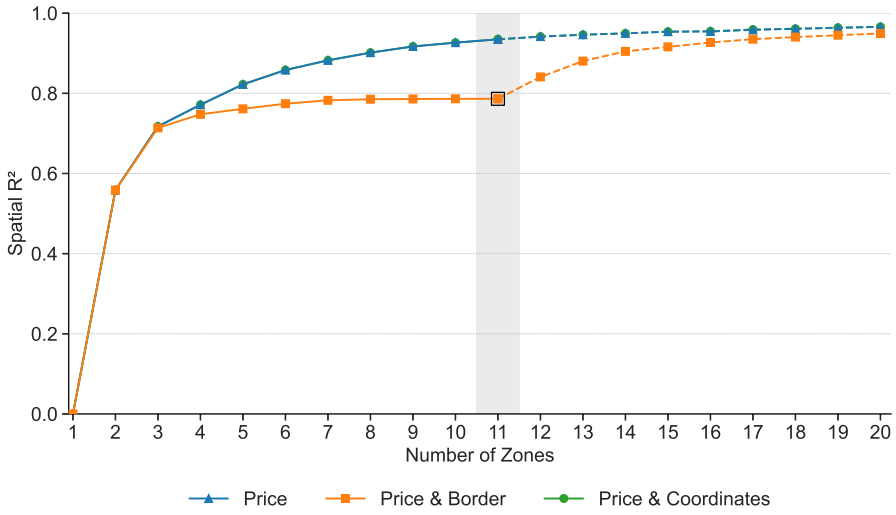
Does market efficiency increase for counterfactual price zones? We implement the bundle choice in equation (10) to construct counterfactual zones and evaluate their spatial efficiency. Figure 7 displays the main results and shows the spatial R^2 of counterfactual zones when exogenously varying the number of zones from $N = 1$ to $N = 20$.

Panel (a) of Figure 7 shows the results for the NYISO market. The results indicate that the currently implemented eleven price zones could be aggregated without incurring any major efficiency loss. This is illustrated by the orange line, which displays the spatial market efficiency of the price-and-border approach. Recall that the price-and-border approach merges and splits existing zones; therefore, when the number of zones equals the currently implemented number of zones (eleven, indicated by the grey bar), the orange line represents the zonal efficiency of the existing price zone regime. As shown, the spatial R^2 remains almost unchanged when the number of zones is marginally reduced to $N = 10$ (or even $N = 7$). Similarly, splitting the most inefficient existing zone—Long Island—and thereby increasing the number of zones to twelve improves spatial efficiency considerably, from 0.79 to 0.84. Consequently, if the costs of administering price zones are high, the current number of zones could be reduced without incurring a substantial efficiency loss. Conversely, if the institutional costs of administering price zones are low, increasing the number of zones could yield sizable efficiency gains for the NYISO market.

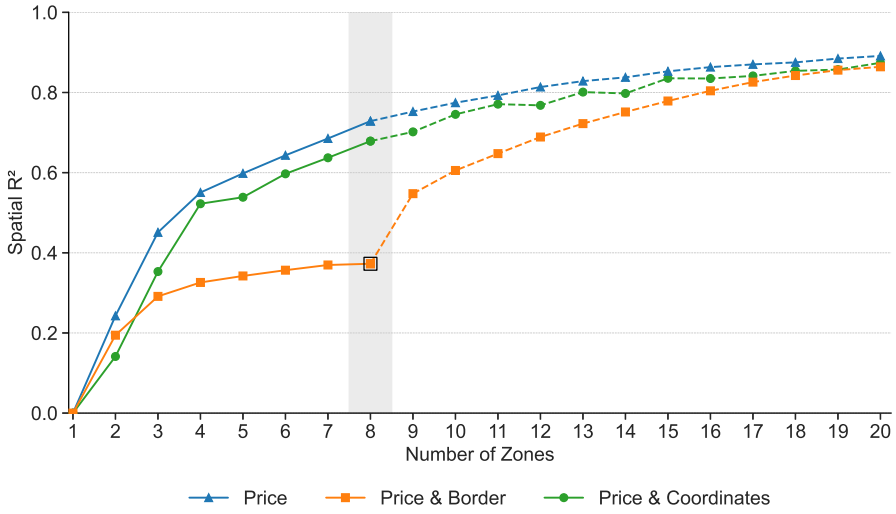
A second main finding is that accounting for geographical borders in the design of price zones comes with significant efficiency losses. To see this, note that the blue line depicts the spatial market efficiency for zones determined by clustering solely on prices. For a small number of zones, e.g., two or three, price clustering provides equivalent welfare results as the price-and-border approach. However, for an intermediate number of zones, the price approach achieves higher degrees of spatial market efficiency. If the currently implemented eleven zones were restructured and assigned by price clustering, the spatial market efficiency increases to 0.93, a 19% increase over the status quo. This marked difference illustrates that the efficiency of the current zones is less than second-best optimal. For a larger number of zones, efficiency values converge again. Finally, note that the spatial R^2 of the price-and-coordinate approach is almost identical to those of the price clustering. This is because the variation in locational prices already follows a strong geographical pattern. Incorporating coordinates into the design of price zones therefore does not change their efficiency.

Figure 7: Spatial market efficiency for counterfactual zones

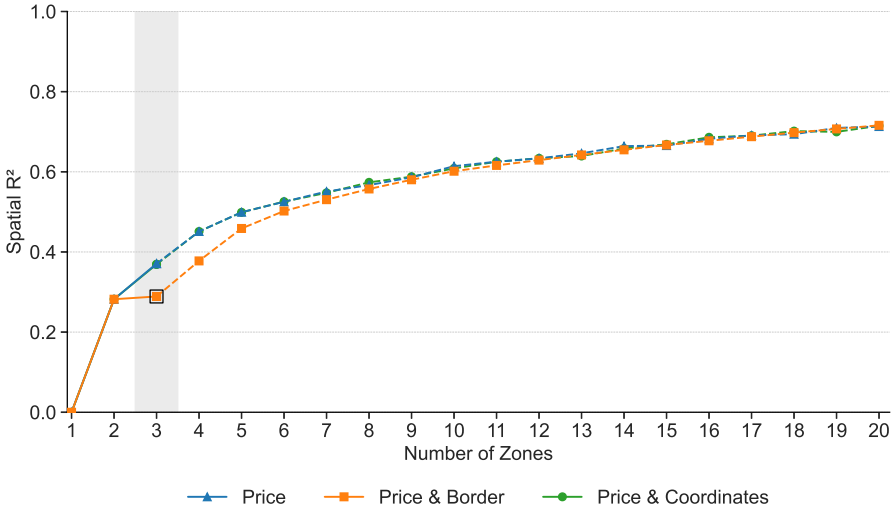
(a) NYISO



(b) ISO-NE



(c) CAISO



Note: The sample includes all locations with price, zone, and coordinate data from 2020 to 2024.

Panel (b) presents the corresponding results for the ISO-NE market. As in the NYISO case, market efficiency increases with the number of zones, again with diminishing marginal returns, where the first few zones capture most of the relevant spatial variation in LMPs. Moreover, the spatial R^2 of the price-and-border approach follows a pattern similar to that observed for NYISO. In particular, splitting the most inefficient zone in ISO-NE—Maine—and increasing the number of zones from eight to nine raises spatial efficiency from 0.37 to 0.54. By contrast, merging the most similar zones in New England and going from $N = 8$ to $N = 7$ zones leads only to a marginal reduction in spatial market efficiency.

Notably, ISO-NE also exhibits substantial welfare differences relative to the price-based second-best allocation. In magnitude, spatial market efficiency nearly doubles, reaching 0.73, when eight zones are optimally assigned solely on the basis of prices (rather than following the current institutional borders). Taken together, the results for both ISO-NE and NYISO indicate that, first, modest reductions in the number of zones entail only negligible efficiency losses; second, splitting zones and adding a single, optimally designed zone can yield sizable efficiency gains; and third, the price-and-border approach, which adheres to institutional geographic boundaries, imposes considerable costs in terms of market efficiency.

Finally, panel (c) presents the results for the CAISO market. While the qualitative patterns observed for NYISO and ISO-NE largely persist, overall efficiency levels are markedly lower. Even with as many as twenty zones, only about 70% of the spatial price variation in California is captured. Again, splitting the existing three zones or reassigning zones solely based on price variation leads to immediate improvements in spatial market efficiency, raising it by about 22% to approximately 0.4.

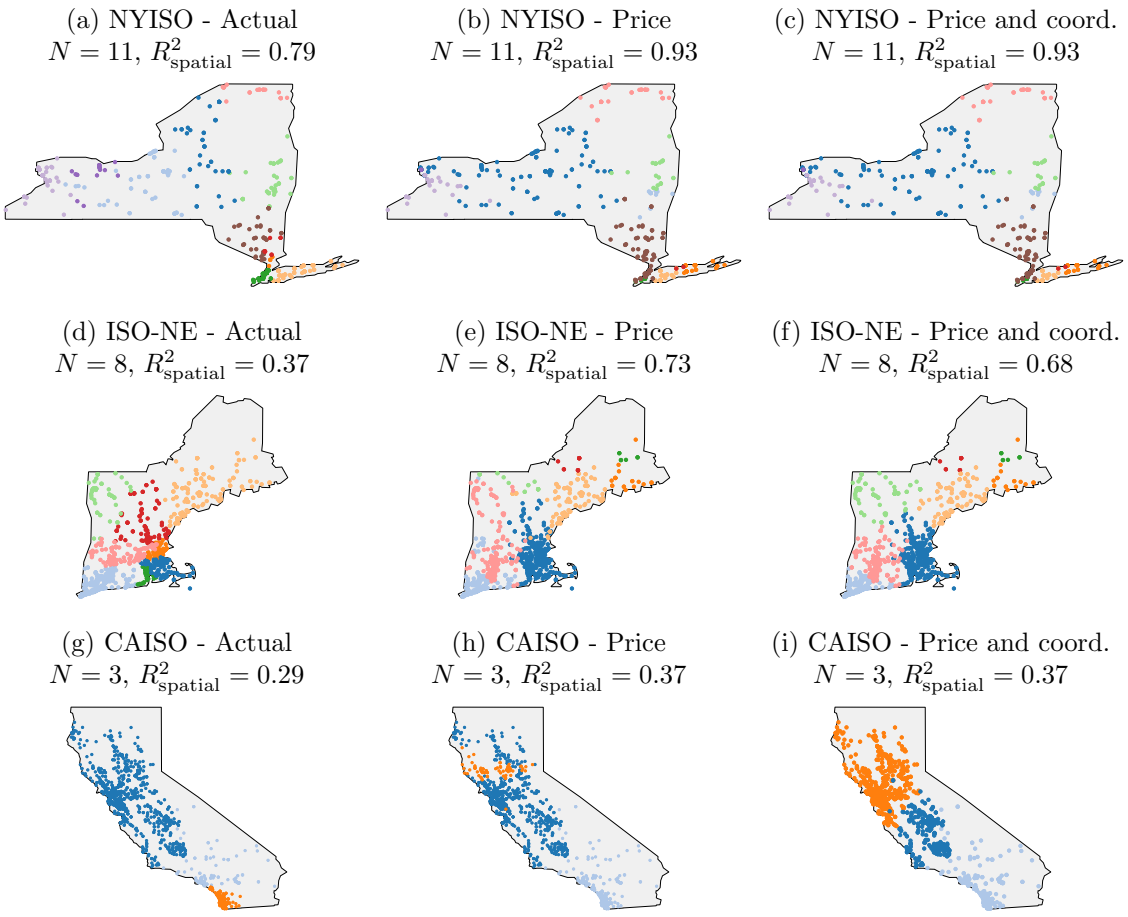
Notably, across all clustering methods, market efficiency remains comparatively low. Consequently, substantial deadweight losses persist under the current price zone implementation in CAISO, arguably reflecting changes in decentralized load and generation patterns that are not accounted for by the current zones. These results remain robust when again including nodes without geographical data available or when applying a load-weighted clustering approach, as shown in Figure B2 in the Appendix.

Geography of counterfactual zones. Our results so far raise the question of how counterfactual zones with higher spatial efficiency are shaped geographically. Figure 8 illustrates the geography of the currently implemented zones alongside counterfactual zones determined by our clustering approaches, where we have kept the number of zones constant for each market.

The top-left map in panel (a) displays the eleven currently implemented price zones of NYISO, where LMPs are color-coded by their corresponding zone. If eleven zones were assigned using price-clustering instead, we obtain the price zones shown in map

(b). Compared to the actual configuration, map (b) shows that Long Island would be split into three subzones—an intuitive result given the high within standard deviation observed in that region (see Figure 2). Importantly, the counterfactual zones obtained from price-clustering are geographically plausible, that is, optimal price zones constructed without explicit geographical information nonetheless yield geographically contiguous regions. This result is confirmed in map (c), which plots the optimal price zones obtained using the price-and-coordinate approach, which explicitly includes geographical information. Comparing maps (b) and (c), the geography of optimal price zones remains nearly unchanged, reflecting the fact that variation in locational prices already closely aligns with geographic patterns in the NYISO market. Consequently, adding coordinates does not significantly alter the market operator's optimization problem. Taken together, our results show that defining price zones based solely on price yields geographically meaningful zones that at the same time yield considerable higher spatial efficiency.

Figure 8: Geographical distribution of selected counterfactual zones



Note: The sample includes all locations with price, zone, and coordinate data from 2020 to 2024. The zone VEA in CAISO is omitted.

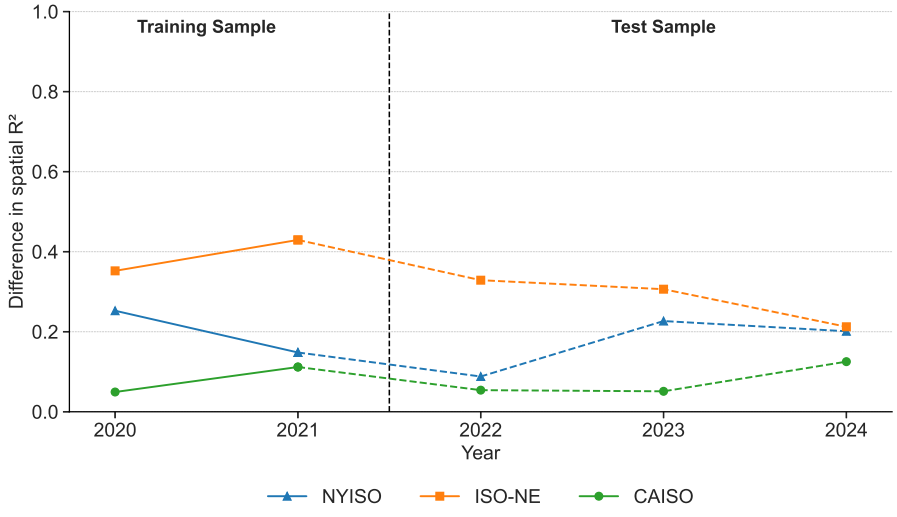
We find similar results for the ISO-NE and CAISO markets. Results for the ISO-NE market are shown in maps (d) through (f). As before, spatial market efficiency increases

significantly from the actual configuration shown in map (d) to the configuration obtained using price-clustering in map (e), while still producing geographically meaningful zone configurations. To achieve higher market efficiency, the price-based configuration in map (e) keeps the total number of zones constant ($N = 8$), but merges Rhode Island and Eastern Massachusetts into a single zone, while splitting the former Maine-area zone into four smaller distinct zones.

Finally, for the CAISO market depicted in maps (g) through (i), we again find that counterfactual zones obtained using price clustering yield meaningful zone configurations with higher market efficiency. Specifically, even when keeping the number of zones unchanged ($N = 3$), higher efficiency is achieved by splitting the Northern California load zone (PGAE) while merging the remaining two zones. As before, our price-clustering approach splits the zone with the largest within-zone standard deviation. We observe a similar split when counterfactual zones are determined by a load-weighted clustering approach, shown in Figure B3) in the Appendix.

Out-of-sample performance. In closing, we investigate the out-of-sample performance for a set of counterfactual zones. We identify these zones using price-clustering applied to a training sample consisting of the first two years (2020 and 2021) of our five-year sample. We then compute the yearly spatial R^2 for these zones over the three subsequent years, which constitute our test sample running from 2022 to 2024. Figure 9 shows the results by plotting the difference in the spatial R^2 relative to the currently implemented zonal configuration. As shown, across all three markets, the counterfactual zones achieve higher spatial market efficiency than the actual zones. Notably, this result holds for both the training and test samples. In terms of magnitude, counterfactual zones exhibit approximately 10 to 20 percentage points higher market efficiency.

Figure 9: Out-of-sample spatial market efficiency by year



Note: The sample includes all locations with price, zone, and coordinate data from 2020 to 2024.

5 Conclusion

In many network-bound markets, prices are aggregated over large geographic regions even though scarcity varies sharply across locations in the network. This paper develops a modeling framework to evaluate the welfare consequences of such spatially coarse pricing. We focus on electricity markets, where regulators commonly define price zones that span multiple nodes in the power grid, abstracting from local congestion and scarcity.

Building on Jacobsen et al. (2020), we show that the welfare effects of coarse price aggregation can be assessed using simple regression-based statistics. Specifically, we introduce the spatial R^2, which measures the efficiency gains of more granular pricing relative to the naive benchmark of a spatially uniform price across the entire market. We then characterize the regulator's problem of selecting an optimal configuration of price zones and use a clustering-based approach to solve for efficiency-maximizing zonal configurations.

We illustrate our methodology using data from three electricity wholesale markets in the United States, all of which aggregate nodal prices into larger zones for retail pricing. Across these markets, we find that alternative zonal configurations, especially those that abstract from institutional constraints, can deliver substantial welfare gains. In turn, our results suggest that the current spatial pricing regimes leave economically meaningful welfare improvements unrealized by failing to align price zones more closely with underlying network conditions.

More broadly, the framework developed in this paper offers a computationally tractable approach for evaluating spatial pricing design in network industries. Although our application focuses on electricity markets, the approach naturally extends to other settings characterized by binding network constraints and spatial heterogeneity of supply and demand, such as natural gas transmission, transportation systems, and telecommunications.

References

Allcott, H. (2011). Rethinking real-time electricity pricing. *Resource and Energy economics 33*(4), 820–842.

Arthur, D. and S. Vassilvitskii (2007). k-means++ the advantages of careful seeding. In *Proceedings of the Annual ACM-SIAM symposium on Discrete algorithms*, pp. 1027–1035.

Astier, N. (2021). Second-best pricing for incomplete market segments: Application to electricity pricing. *Journal of Public Economic Theory 23*(6), 1287–1311.

Bohn, R. E., M. C. Caramanis, and F. C. Schweppe (1984). Optimal pricing in electrical networks over space and time. *The RAND Journal of Economics*, 360–376.

Boiteux, M. (1956). Sur la gestion des monopoles publics astreints à l'équilibre budgétaire. *Econometrica, Journal of the Econometric Society*, 22–40.

Borenstein, S. (2002). The trouble with electricity markets: understanding California's restructuring disaster. *Journal of Economic Perspectives 16*(1), 191–211.

Borenstein, S. and J. Bushnell (2015). The US electricity industry after 20 years of restructuring. *Annual Review of Economics 7*(1), 437–463.

Borenstein, S. and S. Holland (2005). On the Efficiency of Competitive Electricity Markets with Time-Invariant Retail Prices. *The RAND Journal of Economics 36*(3), 469–493.

Burstedde, B. (2012). From nodal to zonal pricing: A bottom-up approach to the second-best. In *2012 9th International Conference on the European Energy Market*, pp. 1–8. IEEE.

Bushnell, J., S. M. Harvey, and B. F. Hobbs (2015). Opinion on Load Granularity Refinements. Members of the Market Surveillance Committee of CAISO.

CAISO (2025). 2024 Annual Report on Market Issues & Performance.

Cramton, P. (2017). Electricity market design. *Oxford Review of Economic Policy 33*(4), 589–612.

Dobos, T., M. Bichler, and J. Knörr (2025). Challenges in finding stable price zones in european electricity markets: Aiming to square the circle? *Applied Energy 382*, 125315.

Eicke, A. and T. Schittekatte (2022). Fighting the wrong battle? A critical assessment of arguments against nodal electricity prices in the European debate. *Energy Policy 170*, 113220.

ENTSO-E (2025). Bidding Zone Review of the 2025 Target Year. Technical report, European Network of Transmission System Operators, Brussels, Belgium.

Fabra, N., D. Rapson, M. Reguant, and J. Wang (2021). Estimating the elasticity to real-time pricing: evidence from the Spanish electricity market. In *AEA Papers and Proceedings*, Volume 111, pp. 425–429. American Economic Association.

Felling, T. and C. Weber (2018). Consistent and robust delimitation of price zones under uncertainty with an application to Central Western Europe. *Energy Economics 75*, 583–601.

Graf, C., F. Quaglia, and F. A. Wolak (2020). Simplified electricity market models with significant intermittent renewable capacity: Evidence from italy. NBER Working Paper 27262, National Bureau of Economic Research.

Green, R. (2007). Nodal pricing of electricity: how much does it cost to get it wrong? *Journal of Regulatory Economics 31*(2), 125–149.

Harberger, A. C. (1964). The measurement of waste. *The American Economic Review 54*(3), 58–76.

Hinchberger, A. J., M. R. Jacobsen, C. R. Knittel, J. M. Sallee, and A. A. van Benthem (2024). The Efficiency of Dynamic Electricity Prices. NBER Working Paper 32995, National Bureau of Economic Research.

Hogan, W. W. (2002). Electricity market restructuring: reforms of reforms. *Journal of Regulatory Economics 21*(1), 103–132.

Jacobsen, M. R., C. R. Knittel, J. M. Sallee, and A. A. Van Benthem (2020). The use of regression statistics to analyze imperfect pricing policies. *Journal of Political Economy 128*(5), 1826–1876.

Joskow, P. and J. Tirole (2006). Retail electricity competition. *The RAND Journal of Economics 37*(4), 799–815.

Knittel, C. R., J. R. L. Senga, and S. Wang (2025). Flexible data centers and the grid: Lower costs, higher emissions? NBER Working Paper 34065, National Bureau of Economic Research.

Miraftabzadeh, S. M., C. G. Colombo, M. Longo, and F. Foiadelli (2023). K-means and alternative clustering methods in modern power systems. *IEEE Access 11*, 119596–119633.

Moretto, M. and M. G. Pollitt (2026). How many zones should an electricity market have? A cross-country perspective on bidding zone design. *Energy Policy 210*, 115030.

Neuhoff, K., F. Klaucke, L. Olmos, L. Ryan, S. Vitiello, A. Papavasiliou, and K. Staschus (2025). EU power market reform toward locational pricing: Rewarding flexible consumers for resolving transmission constraints. *Energy Policy 207*, 114808.

Norris, T., T. Profeta, D. Patino-Echeverri, and A. Cowie-Haskell (2025). Rethinking load growth: assessing the potential for integration of large flexible loads in US power systems. Technical report, Nicholas Institute for Energy, Environment & Sustainability, Duke University.

Potomac Economics (2025). 2024 State of the Market Report for the ERCOT Electricity Market.

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

Schweppe, F. C., M. C. Caramanis, R. D. Tabors, and R. E. Bohn (2013). *Spot pricing of electricity*. Springer Science & Business Media.

Triolo, R. C. and F. A. Wolak (2022). Quantifying the benefits of a nodal market design in the Texas electricity market. *Energy Economics 112*, 106154.

Vickrey, W. S. (1969). Congestion theory and transport investment. *The American Economic Review 59*(2), 251–260.

Weibelzahl, M. (2017). Nodal, zonal, or uniform electricity pricing: how to deal with network congestion. *Frontiers in Energy 11*(2), 210–232.

Wolak, F. A. (2011). Measuring the benefits of greater spatial granularity in short-term pricing in wholesale electricity markets. *American Economic Review 101*(3), 247–252.

Wolak, F. A. (2019). The role of efficient pricing in enabling a low-carbon electricity sector. *Economics of Energy & Environmental Policy 8*(2), 29–52.

Zarnikau, J., C. K. Woo, and R. Baldick (2014). Did the introduction of a nodal market structure impact wholesale electricity prices in the Texas (ERCOT) market? *Journal of Regulatory Economics 45*(2), 194–208.

A Proofs and methods

A.1 SST decomposition

Consider J locations, T time periods and N zones. The locational price of location j at time t is given by p_{jt}, The spatial mean at time t is

$$\bar{p}_t = \frac{1}{J}\sum_j p_{jt}.$$

We write the total sum of squares, SST, as

$$\sum_t \sum_j (p_{jt} - \bar{p})^2 = \sum_t \sum_j \big((p_{jt} - \bar{p}_t) + (\bar{p}_t - \bar{p})\big)^2.$$

Expanding individual terms yields

$$\sum_t \sum_j (p_{jt} - \bar{p}_t)^2 + \sum_t \sum_j (\bar{p}_t - \bar{p})^2 + 2\sum_t \sum_j (p_{jt} - \bar{p}_t)(\bar{p}_t - \bar{p}).$$

To obtain equation (3) in the main text, we show that the last term is equal to 0:

$$\begin{aligned}\sum_t \sum_j (p_{jt} - \bar{p}_t)(\bar{p}_t - \bar{p}) &= \sum_t \sum_j \bar{p}(\bar{p}_t - p_{jt}) + \sum_t \sum_j \bar{p}_t(p_{jt} - \bar{p}_t) \\ &= \bar{p}\sum_t \sum_j (\bar{p}_t - p_{jt}) + \sum_t \bar{p}_t \sum_j (p_{jt} - \bar{p}_t),\end{aligned}$$

where $\sum_j \bar{p}_t = J\bar{p}_t = \sum_j p_{jt}$. Hence, the term simplifies to

$$= \bar{p}\sum_t \underbrace{\sum_j (\bar{p}_t - p_{jt})}_{0} + \sum_t \bar{p}_t \underbrace{\sum_j (p_{jt} - \bar{p}_t)}_{0} = 0.$$

Equation (4) follows analogously. Let the zonal price of location j at time t be $\bar{p}_{nt}(j)$. The zonal price of every location belonging to zone n is

$$\bar{p}_{nt} = \frac{1}{J_n}\sum_{j \in \mathcal{Z}_n} p_{jt},$$

where $J_n = |\mathcal{Z}_n|$. We write the spatial sum of squares, $\text{SST}_{\text{spatial}}$, as

$$\sum_t \sum_j (p_{jt} - \bar{p}_t)^2 = \sum_t \sum_j \big((p_{jt} - \bar{p}_{nt}(j)) + (\bar{p}_{nt}(j) - \bar{p}_t)\big)^2.$$

Again, expanding individual terms yields

$$= \sum_t \sum_j (p_{jt} - \bar{p}_{nt}(j))^2 + \sum_t \sum_j (\bar{p}_{nt}(j) - \bar{p}_t)^2 + 2\sum_t \sum_j (p_{jt} - \bar{p}_{nt}(j))(\bar{p}_{nt}(j) - \bar{p}_t).$$

To obtain equation (4) in the main text, we show that the last term is equal to 0.

Rearrange

$$\sum_t \sum_j (p_{jt} - \bar{p}_{nt}(j))(\bar{p}_{nt}(j) - \bar{p}_t) = \sum_t \sum_j \bar{p}_t(\bar{p}_{nt}(j) - p_{jt}) + \sum_t \sum_j \bar{p}_{nt}(j)(p_{jt} - \bar{p}_{nt}(j))$$
$$= \sum_t \sum_n \sum_{j \in \mathcal{Z}_n} \bar{p}_t(\bar{p}_{nt}(j) - p_{jt}) + \sum_t \sum_n \sum_{j \in \mathcal{Z}_n} \bar{p}_{nt}(j)(p_{jt} - \bar{p}_{nt}(j))$$

and recall that $\bar{p}_{nt}(j) = \bar{p}_{nt}$ for all $j \in \mathcal{Z}_n$ so that the term further reduces to

$$= \sum_t \bar{p}_t \sum_n \sum_{j \in \mathcal{Z}_n} (\bar{p}_{nt} - p_{jt}) + \sum_t \sum_n \bar{p}_{nt} \sum_{j \in \mathcal{Z}_n} (p_{jt} - \bar{p}_{nt}),$$

where $\sum_{j \in \mathcal{Z}_n} \bar{p}_{nt} = J_n \bar{p}_{nt} = \sum_{j \in \mathcal{Z}_n} p_{jt}$. Hence, the term simplifies to

$$= \sum_t \bar{p}_t \sum_n \underbrace{\sum_{j \in \mathcal{Z}_n} (\bar{p}_{nt} - p_{jt})}_{0} + \sum_t \sum_n \bar{p}_{nt} \underbrace{\sum_{j \in \mathcal{Z}_n} (p_{jt} - \bar{p}_{nt})}_{0} = 0.$$

A.2 Weighting deadweight loss by supply and demand slopes

As in the main text, let $(p_1^*, \ldots, p_J^*)$ denote the vector of first-best prices and $(p_1, \ldots, p_J)$ the vector of (potentially imperfect) product prices, i.e., bundle prices. First, note that due to local linearity of supply and demand, we can write $mc_j(x_j(p_j))$ as

$$mc_j(x_j(p_j)) = mc_j^* + \frac{\partial mc_j}{\partial x_j}(x_j(p_j) - x_j(p_j^*))$$
$$= mc_j^* + \frac{\partial mc_j}{\partial x_j}\frac{\partial x_j}{\partial p_j}(p_j - p_j^*),$$

where the first equality follows from local linearity of supply and the second from local linearity of demand, and mc_j^* denotes the first-best marginal costs. Using $p_j^* = mc_j^*$ and further re-arranging we obtain

$$p_j^* + \frac{\partial mc_j}{\partial x_j}\frac{\partial x_j}{\partial p_j}(p_j - p_j^*) = (1 - \frac{\partial mc_j}{\partial x_j}\frac{\partial x_j}{\partial p_j})p_j^* + \frac{\partial mc_j}{\partial x_j}\frac{\partial x_j}{\partial p_j}p_j.$$

Inserting this expression into equation (5) in the main text yields

$$\text{DWL}_j = -\frac{1}{2}(p_j - p_j^*)\left(1 - \frac{\partial mc_j}{\partial x_j}\frac{\partial x_j}{\partial p_j}\right)[x_j(p_j) - x_j(p_j^*)].$$

From local linearity of demand, we have

$$\text{DWL}_j = -\frac{1}{2}(p_j - p_j^*)\left(1 - \frac{\partial mc_j}{\partial x_j}\frac{\partial x_j}{\partial p_j}\right)\frac{\partial x_j}{\partial p_j}(p_j - p_j^*),$$
$$= -\frac{1}{2}(p_j - p_j^*)^2\frac{\partial x_j}{\partial p_j}\left(1 - \frac{\partial mc_j}{\partial x_j}\frac{\partial x_j}{\partial p_j}\right),$$

$$= -\frac{1}{2}(p_j - p_j^*)^2 \omega_j, \quad \text{where}$$
$$\omega_j = \frac{\partial x_j}{\partial p_j}\left(1 - \frac{\partial mc_j}{\partial x_j}\frac{\partial x_j}{\partial p_j}\right) \leq 0.$$

A.3 Derivation of spatial R^2

Consider the main specification (13) with zonal and time dummies

$$\text{LMP}_{jt} = \sum_{s=1}^{T}\sum_{n=1}^{N} \beta_{nt} D_{jt}^s \times D_{jt}^n + \epsilon_{jt},$$

whose goodness-of-fit measure we denote as R^2. Moreover, consider the alternative specification with time dummies

$$\text{LMP}_{jt} = \sum_{s=1}^{T} \beta_t D_{jt}^s + \epsilon_{jt},$$

whose goodness-of-fit measure we denote as R^2_{temporal}. Then, our spatial R^2 measure is obtained from a normalization of these R^2 values. Specifically,

$$R^2_{\text{spatial}} := \frac{R^2 - R^2_{\text{temporal}}}{1 - R^2_{\text{temporal}}} = \frac{1 - \frac{\text{SSR}}{\text{SST}} - (1 - \frac{\text{SSR}_{\text{temporal}}}{\text{SST}})}{1 - (1 - \frac{\text{SSR}_{\text{temporal}}}{\text{SST}})} = \frac{\frac{\text{SSR}_{\text{temporal}}}{\text{SST}} - \frac{\text{SSR}}{\text{SST}}}{\frac{\text{SSR}_{\text{temporal}}}{\text{SST}}},$$

where SST denotes total sum of squares and SSR and $\text{SSR}_{\text{temporal}}$ the sum of squared residuals of the respective regression. Further simplifying yields

$$= \frac{\text{SSR}_{\text{temporal}} - \text{SSR}}{\text{SSR}_{\text{temporal}}} = 1 - \frac{\text{SSR}}{\text{SSR}_{\text{temporal}}}.$$

Following equation (4), SSR represent the within-zone variation (fitted values are hourly zonal averages), while $\text{SSR}_{\text{temporal}}$ is equal to total spatial variation (fitted values are hourly market-wide averages). Writing out the SSR expressions yields

$$= 1 - \frac{\sum_t \sum_j (\bar{p}^*_{nt}(j) - p^*_{jt})^2}{\sum_t \sum_j (\bar{p}^*_t - p^*_{jt})^2} = 1 - \frac{\text{DWL}(\bar{p}^*_{11}, \ldots, \bar{p}^*_{NT})}{\text{DWL}(\bar{p}^*_1, \ldots, \bar{p}^*_T)},$$

which is identical to equation (15) in the main text. Our spatial R^2 measure can therefore also be interpreted as the relative (percent-wise) improvement of the goodness-of-fit through zonal pricing compared to one hourly market-wide price.

A.4 Clustering algorithms

k-means++. All clustering in this paper is conducted based on the k-means++ algorithm. Our analysis uses the `scikit-learn` package in Python. The pseudo-code below (Algorithm 1) describes the general implementation of the algorithm in this paper. For more details regarding the initialization process of k-means++, see Arthur and Vassilvitskii (2007).

Algorithm 1 General (weighted) k-means++ implementation

1: **Input** data vector $\{\boldsymbol{x}_j\}_{j=1}^J \in \mathbb{R}^T$, weight vector $\{w_j\}_{j=1}^J$, number of clusters N
2: **Initialize** N cluster centroids $\{\boldsymbol{\mu}_n\}_{n=1}^N \in \mathbb{R}^T$ randomly following k-means++ procedure
3: **Repeat** until convergence:
4: For every j, **set** ▷ Assign data points to clusters

$$j \in \mathcal{Z}_n \text{ if } n = \arg\min_n \|\boldsymbol{x}_j - \boldsymbol{\mu}_n\|^2.$$

5: For every n, **set** ▷ Recompute centroids as weighted average

$$\mu_n := \frac{\sum_{j \in \mathcal{Z}_n} \boldsymbol{x}_j w_j}{\sum_{j \in \mathcal{Z}_n} w_j}.$$

Clustering by price. This approach employs the clustering algorithm with hourly locational prices as input features. We repeat the clustering 20 times and choose the repetition that minimizes within-cluster sum of squares. The algorithm is described in the pseudo-code below (Algorithm 2).

Algorithm 2 Price approach

1: **Apply** Algorithm 1 with `20` repetitions (indexed i) and inputs:
2: Standardized locational hourly prices as data vector $\{\boldsymbol{x}_j\}_{j=1}^J \in \mathbb{R}^T$
3: ▷ J nodes, T time periods
4: Number of desired zones as clusters N
5: (optional) Locational weights as weight vector $\{w_j\}_{j=1}^J$
6: **Determine** repetition with set of clusters that minimizes the (weighted) within-cluster sum of squares

$$\{\mathcal{Z}_n^*\}_{n=1}^N = \arg\min_{\mathcal{Z}_1^i,\ldots,\mathcal{Z}_N^i} \sum_{n=1}^N \sum_{j \in \mathcal{Z}_n^i} \|\boldsymbol{x}_j - \boldsymbol{\mu}_n\|^2 w_j$$

Clustering by price and coordinates. This approach employs the clustering algorithm with hourly locational prices and coordinates as input features. Because the vector of hourly prices can be large (more than 40000 entries for a five year sample), we scale the coordinate variables to ensure they receive sufficient weight relative to the price vector.

The algorithm is described in the pseudo-code below (Algorithm 3). Line 1-12 describe the procedure of determining the optimal feature scale for coordinates, which are depicted in Table A1. Line 13-17 determine the clusters based on hourly locational prices and (scaled) coordinates as input features.

Algorithm 3 Price and coordinates approach

1: **Compute** the within-zone sum of squares of standardized coordinates (longitude and latitude) with respect to current zones and **define** it as threshold τ
2: **Initiate** the feature scale of coordinates as $c = 1$
3: **Multiply** standardized coordinate features by the feature scale c
▷ Scaling of coordinates relative to price vector
4: **Apply** Algorithm 1 with `5` repetitions (indexed i) and inputs:
5: Standardized locational hourly prices and scaled coordinates as data vector $\{\boldsymbol{x}_j\}_{j=1}^{J} \in \mathbb{R}^{T+2}$
▷ J nodes, T time periods, 2 coordinates
6: Number of existing zones as clusters N
7: **Compute** the within-clusters sum of squares of coordinates for every repetition and **define** it as t_i

$$t_i := \sum_{n=1}^{N} \sum_{j \in \mathcal{Z}_n^i} (x_j^{\text{long}} - \mu_n^{\text{long}})^2 + (x_j^{\text{lat}} - \mu_n^{\text{lat}})^2$$

8: **if** $\{t_i\}_{i=1}^{5} << \tau$ **then** ▷ Geographic sum of squares lower than threshold
9: **Set** $c^* = c$ and **continue** with Step 13
10: **else**
11: **Increase** feature scale $c = c + \epsilon$ and **continue** with Step 3
12: **end if**
13: **Multiply** standardized coordinate features by the determined feature scale c^*
14: **Apply** Algorithm 1 with `20` repetitions (indexed i) and inputs:
15: Standardized locational hourly prices and scaled coordinates as data vector $\{\boldsymbol{x}_j\}_{j=1}^{J} \in \mathbb{R}^{T+2}$
16: Number of desired zones as clusters N
17: **Determine** repetition with set of clusters that minimizes the within-cluster sum of squares

$$\{\mathcal{Z}_n^*\}_{n=1}^{N} = \arg \min_{\mathcal{Z}_1^i, \ldots, \mathcal{Z}_N^i} \sum_{n=1}^{N} \sum_{j \in \mathcal{Z}_n^i} \|\boldsymbol{x}_j - \boldsymbol{\mu}_n\|^2$$

Table A1: Feature scale for coordinates for five year sample

Market	Weight
NYISO	1
ISO-NE	13
CAISO	15

Note: The feature scale is determined based on Algorithm 3. A scale of x means that both (standardized) latitude and longitude is equivalent to x hours of prices.

Clustering by price and borders. This approach takes into account the existing zonal structure, which may be determined by institutional borders. Thus, any counterfactual zones from this approach are a merger of existing zones (in case of aggregation) or splits within one zone (in case of disaggregation). The procedure of this approach differs between aggregation and disaggregation.

Algorithm 4 Price and borders approach (aggregation)

1: **Inputs:** A set of M zones $\mathcal{M} = \{\mathcal{Z}\}_{n=1}^{M}$ with geographic adjacency relation $\sim \subseteq \mathcal{M} \times \mathcal{M}$
2: Identify all eligible (adjacent) pairs $\mathcal{P} = \{(m_1, m_2) | \mathcal{Z}_{m_1} \sim \mathcal{Z}_{m_2}\}$
3: For every pair $p = (m_1, m_2) \in \mathcal{P}$, **define** the merged zone $\mathcal{Z}_p = \mathcal{Z}_{m_1} \bigcup \mathcal{Z}_{m_2}$
4: **Determine** the pair whose set of zones minimizes the (weighted) within-zone sum of squares

$$\{\mathcal{Z}_n^*\}_{n=1}^{M-1} = \arg \min_{\mathcal{Z}_1^p, \ldots, \mathcal{Z}_p^p, \ldots, \mathcal{Z}_{M-1}^p} \sum_{n=1}^{M-1} \sum_{j \in \mathcal{Z}_n^p} \|\boldsymbol{x}_j - \boldsymbol{\mu}_n\|^2 w_j$$

5: **Replace** the set of zones $\mathcal{M}$ with $\{\mathcal{Z}_n^*\}_{n=1}^{M-1}$, reducing the number of zones to $M-1$
6: **Repeat** Step 1-6 until $M \leq N$, with N being the number of desired zones

Algorithm 5 Price and borders approach (disaggregation)

1: **Inputs:** A set of M zones $\mathcal{M} = \{\mathcal{Z}\}_{n=1}^{M}$
2: For each zone m, **determine** its optimal split $\mathcal{Z}_m \to \mathcal{Z}_{m_1} \bigcup \mathcal{Z}_{m_2}$ by **applying** Algorithm 2 with inputs:
3: Standardized locational hourly prices in zone m as data vector $\{\boldsymbol{x}_j\}_{j \in \mathcal{Z}_m} \in \mathbb{R}^T$
4: 2 as number of clusters
5: (optional) Locational weights as weight vector $\{w_j\}_{j \in \mathcal{Z}_m}$
6: **Determine** the split whose set of zones minimizes the (weighted) within-zone sum of squares

$$\{\mathcal{Z}_n^*\}_{n=1}^{M+1} = \arg \min_{\mathcal{Z}_1^m, \ldots, \mathcal{Z}_{m_1}^m, \mathcal{Z}_{m_2}^m, \ldots, \mathcal{Z}_{M+1}^m} \sum_{n=1}^{M+1} \sum_{j \in \mathcal{Z}_n^m} \|\boldsymbol{x}_j - \boldsymbol{\mu}_n\|^2 w_j$$

7: **Replace** the set of zones $\mathcal{M}$ with $\{\mathcal{Z}_n^*\}_{n=1}^{M+1}$, increasing the number of zones to $M+1$
8: **Repeat** Step 1-8 until $M \geq N$, with N being the number of desired zones

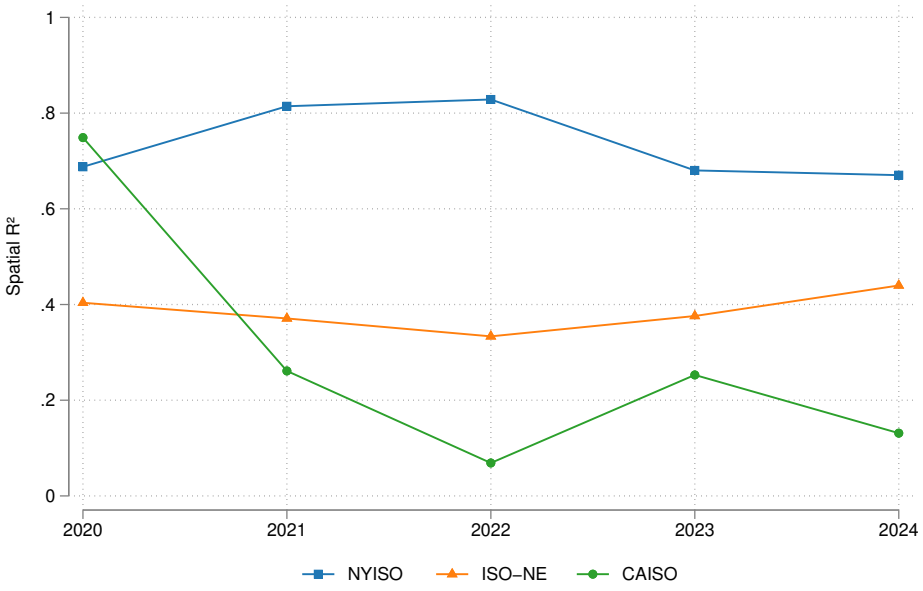
B Additional results

Table B1: Market efficiency of current price zones in 2020-2024 (all nodes)

	NYISO			ISO-NE			CAISO		
	LMP	Congestion	Losses	LMP	Congestion	Losses	LMP	Congestion	Losses
R^2_{spatial}	0.786	0.767	0.880	0.373	0.287	0.566	0.309	0.309	0.154
R^2	0.959	0.876	0.917	0.996	0.298	0.585	0.930	0.411	0.203
R^2_{temporal}	0.809	0.470	0.309	0.994	0.016	0.043	0.898	0.147	0.058
Nodes		548			1005			2177	
Obs.		24,025,964			44,062,215			95,393,963	

Note: The sample includes all locations with price, zone, and coordinate data from 2020 to 2024.

Figure B1: Spatial market efficiency of current price zones by year



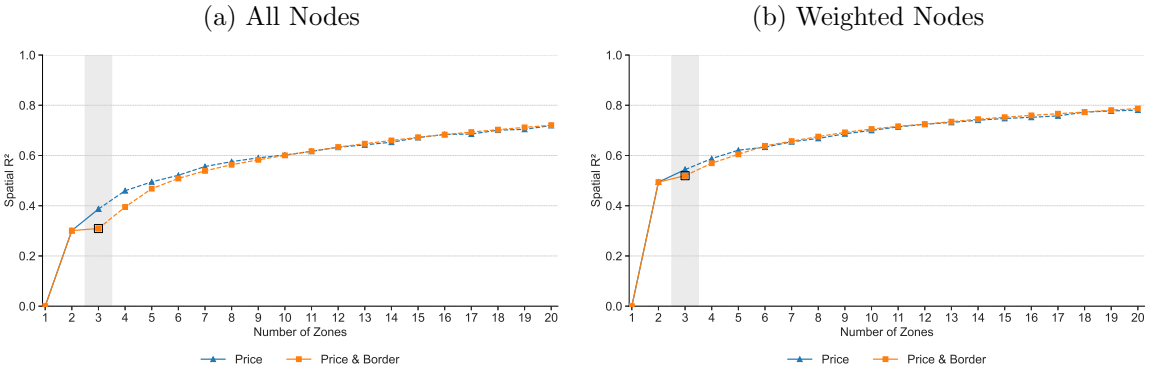
Note: The sample includes all locations with price, zone, and coordinate data from 2020 to 2024.

Table B2: Market efficiency of current price zones in 2020-2024 (CAISO, weighted)

	CAISO		
	LMP	Congestion	Losses
R^2_{spatial}	0.526	0.522	0.432
R^2	0.968	0.541	0.462
R^2_{temporal}	0.933	0.039	0.054
Nodes		1371	
Obs.		60,075,849	

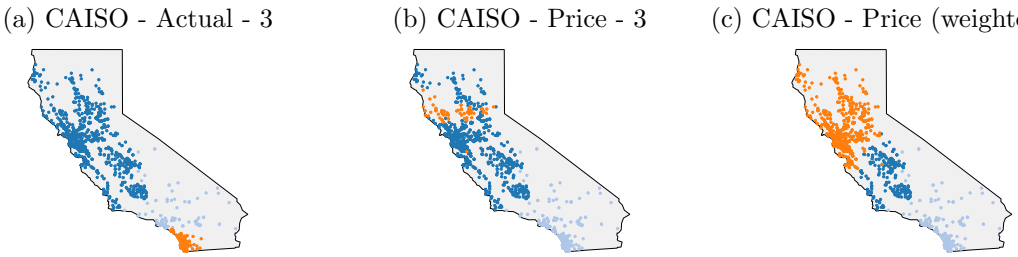
Note: The sample includes all locations with price, zone, coordinate, and load factor data from 2020 to 2024.

Figure B2: Spatial market efficiency for counterfactual zones (CAISO)



Note: Panel (a): The sample includes all locations with price and zone data from 2020 to 2024. Panel (b): The sample includes all locations with price, zone, coordinate, and load factor data from 2020 to 2024

Figure B3: Comparison of CAISO's counterfactual zones - unweighted vs. weighted



Note: The sample includes all locations with price, zone, coordinate, and load factor data from 2020 to 2024.