Designing Environmental Markets to Address Community Impacts

Danae Hernandez-Cortes* Bryan Leonard† Sara Sutherland‡

November 2026

PRELIMINARY DRAFT. PLEASE DO NOT SHARE WITHOUT PERMISSION.

In memory of Thomas D. Crocker (1936–2026)[1]

Abstract

Environmental markets minimize the aggregate costs of meeting environmental targets, but often change the spatial pattern of economic activity and environmental quality. This spatial incidence generates community opposition to markets and may affect efficiency when regulators cannot price impacts without knowing local damages and preferences. We propose a stacked permit system in which firms covered by an aggregate cap must also hold community-specific permits in proportion to their local impacts. Communities receive slack endowments of these permits and sell them to firms at posted prices, thereby revealing local damages. Under price-taking behavior, the equilibrium attains the constrained first best without the regulator knowing community damages. Because each community is the sole seller of its own permits and firms must assemble permits from multiple communities, however, strategic pricing generates monopoly markups and an input assembly problem across communities. We derive closed-form welfare comparisons to spatially differentiated and uniform taxes and trading systems. The ranking depends on the magnitude of regulator damage uncertainty relative to community market power. Dividing each community's selling authority among independent sub-sellers introduces within-community free-riding that counteracts between-community markups. In a simulation on California's carbon market, competitive community pricing improves on the carbon-only status quo by $66 million per year. Unitary community pricing more than erases that gain. Dividing each county's selling authority among ten sub-sellers recovers three-quarters of the gain with no damage information. Setting the subdivision

*Arizona State University Danae.Hernandez-Cortes@asu.edu

†University of Wyoming bryan.leonard@uwyo.edu

‡U.C. Davis sasutherland@ucdavis.edu

[1]Tom Crocker proposed the first market in pollution rights (Crocker, 1966) and, five decades later, documented the joint abatement technology linking carbon to its co-pollutants (Agee et al., 2014), two ideas central to this paper. He joined the University of Wyoming economics faculty in 1975 and helped transform it into an internationally recognized center for environmental and resource economics, serving as founding director of the Haub School of Environment and Natural Resources. He believed, in the words of his colleague Jason Shogren, in "exploring how alternative institutions allocate resources efficiently and with equal opportunity." We hope this paper is in that spirit.

from the regulator's noisy damage estimates recovers more than four-fifths of the gain and outperforms a differentiated tax built from the same estimates under plausible estimation error.

1 Introduction

Environmental markets use incentives and property rights to reduce the costs of meeting environmental objectives. In recent decades, regulators have applied them to fisheries, surface water, groundwater, habitat, wetlands, air pollution, and climate change mitigation. These markets rely on trade among firms to re-allocate production in a way that minimizes the costs of complying with an aggregate objective, without the regulator needing to know individual firms' cost functions (Montgomery, 1972). While market-based tools have intuitive theoretical appeal and have been favored by many regulators, they often face political opposition and concerns from affected communities (Fowlie et al., 2020). One reason is that regulators are typically focused on an aggregate environmental objective, while communities care about the spatial distribution of the associated costs and benefits and about whether they have a voice in determining those outcomes (Banzhaf et al., 2019).

When damages from the regulated activity have a spatial incidence, unrestricted trading among firms in an environmental market can have both efficiency and distributional implications (Muller and Mendelsohn, 2009; Fowlie et al., 2012). For instance, tradable fishing quota that reduce over-harvest of a fish stock can change where fishers offload their catch (Kokorsch and Benediktsson, 2018), and stabilizing aquifer or river levels through water markets can generate return flow and other hydrologic externalities (Chong and Sunding, 2006; Leonard et al., 2019). Offset markets that achieve "no net loss" of wetland acreage can still relocate flood protection away from certain communities (Aronoff and Rafey, 2026), and reducing CO_2 emissions via cap-and-trade can change communities' exposure to local co-pollutants (Hernandez-Cortes and Meng, 2023). Environmental markets can address some of these concerns with spatially differentiated prices, caps, or trading ratios (Tietenberg, 1995; Hung and Shaw, 2005). However, these approaches require the regulator to know community-level damages to function efficiently (Holland and Yates, 2015). That information may be difficult to observe, particularly when damages depend on locally held knowledge or community preferences not captured by standard estimates (Hayek, 1945).[2]

This paper develops an alternative approach to regulating spatially heterogeneous damages that are unknown to the regulator. We propose a stacked permit system in which firms must hold community-specific permits covering the relevant impacts (measured in terms of economic activity or environmental quality) they create in each affected community, in addition to the permits required under an aggregate cap. Communities receive a permit endowment that exceeds business-as-usual exposure and sell permits to firms at posted prices. The equilibrium level and spatial distribution of local impacts are therefore determined by communities' pricing decisions rather than by the regulator's damage estimates. The regulator assigns property rights over local exposure to the parties who bear the damages in the spirit of Coase (1960). It also addresses pro-

[2] A related mechanism-design literature elicits polluters private information about their abatement costs through incentive-compatible schemes (Kwerel, 1977; Dasgupta et al., 1980; Montero, 2008; Duggan and Roberts, 2002). However, these mechanisms assume that marginal *damages* are known to the regulator.

cedural environmental justice concerns by giving communities direct control over the price and level of local exposure, rather than relying on externally determined compensation (Schlosberg, 2007).

Assigning exclusive permit rights to individual communities introduces market-structure distortions, however. Because each community is the sole seller of its own permits, a unitary community authority has an incentive to act as a monopolist and set its permit price above marginal damage. Moreover, firms must assemble permits from every community they affect. This complementarity across communities' permits creates an anticommons externality (Heller, 1998; Buchanan and Yoon, 2000). When one community raises its permit price, it reduces firms' local emissions and thereby lowers permit demand in other communities. Because each community ignores this fiscal externality on other communities' permit revenues, equilibrium prices exceed the sole-seller monopoly benchmark.

We derive the community permit mechanism's properties in general form before turning to the specific case of CO_2 abatement with local co-pollutants. Cap-and-trade markets for CO_2 emissions provide an especially useful context in which to study the mechanism. The benefits of reducing CO_2 are global regardless of where abatement occurs, but carbon abatement also changes emissions of local co-pollutants such as NO_x and SO_2. The welfare consequences therefore depend on the spatial pattern of abatement, and previous work suggests these consequences are significant (Burtraw et al., 2003; Agee et al., 2014; Fullerton and Karney, 2018). Moreover, firm behavior, existing carbon markets, and the atmospheric and engineering relationships governing co-pollutant exposure are well documented, making it possible to quantify the mechanism's performance (e.g., Tessum et al., 2017; Holland et al., 2016).

Previous proposals for differentiated carbon policy require estimates of community-specific damages from co-pollutants. Source-specific prices and trading ratios build damage estimates directly into the instrument (Muller and Mendelsohn, 2009; Hung and Shaw, 2005). Facility-specific caps layered onto the carbon market ration quantities without prices, and require the regulator to set caps for specific facilities that must first be identified (Burtraw and Roy, 2023). Zonal designs and offset systems coarsen the information requirement (at the cost of optimal targeting) but still require damage estimates (Atkinson and Tietenberg, 1982; Krupnick et al., 1983). When the regulator's damage estimates include error, these spatially differentiated approaches can deliver lower expected welfare than uniform regulation, because the gains from differentiation are bounded by the quality of the underlying damage information (Fowlie and Muller, 2019). In contrast, assigning the permits to the communities themselves creates a mechanism through which markets can reveal community damages without relying on the regulator's damage estimates.

In our two-pollutant model, firms face an economy-wide cap on CO_2 emissions alongside a requirement to hold community-specific permits covering their contributions to local ambient pollution, as determined by a known source-receptor transport matrix. Under price-taking behavior, communities price permits at marginal damage and the equilibrium attains the constrained first best, given the carbon market cap. We generalize the linear-quadratic framework of Fowlie

and Muller (2019) to incorporate joint abatement costs between aggregate and local pollutants, allowing us to derive closed-form welfare comparisons between community permits and both undifferentiated and spatially differentiated taxes and cap-and-trade.

To address the market power that the mechanism generates, we develop a governance modification that divides each community's selling authority among multiple independent sub-sellers, such as neighborhood councils or local districts. Across communities, permits are complements because each firm must cover its exposure in every downwind jurisdiction, so each community quotes its own posted price. Within a community, however, all permits cover the same receptor and are perfect substitutes for regulated firms. Each sub-seller therefore consigns an authorized quantity of permits from the community's common pool into a single community clearinghouse. Much like a joint cap-and-trade auction into which participating jurisdictions consign allowances, the clearinghouse clears firms' compliance demand at one uniform community price. Because each sub-seller collects the revenue on the permits it consigns while bearing only a fraction of community-wide damage, this structure creates a within-community free-riding incentive that pushes prices down relative to the between-community monopoly markup. The design is also analogous to special business districts in the land-use context, which divide a jurisdiction's monopoly power over commercial siting among multiple local boards.

As a theoretical benchmark, we show that there exists a community-specific number of sub-sellers that exactly balances between-community market power and within-community free-riding and restores the constrained first best. Because computing that threshold requires knowing local damages, we focus on two practical implementations. First, if the regulator uses noisy damage estimates to choose the number of sub-sellers in each community, communities' own pricing decisions partially correct the regulator's error, so the mechanism can outperform differentiated taxes and cap-and-trade built from the same estimates. Second, if the regulator assigns the same number of sub-sellers to every community, a uniform subdivision that requires no damage estimates at all, the mechanism removes most of the market-power loss while avoiding damage-estimation error. We also show that the mechanism needs one further design element when some emissions cannot be abated at any price, which is the case for every facility in our data. A community whose receptor field contains such a facility can hold firms up over those tons, its payoff is unbounded in its own price, and no equilibrium exists. Exempting the unabatable tons from the permit requirement restores existence of the equilibrium.

We simulate the model on California's cap-and-trade program using the 279 large stationary sources that account for 99.9 percent of covered emitter-reported emissions in 2019. We map facility-level NO_x emissions into tract-level ambient $PM_{2.5}$ concentrations using the InMAP atmospheric transport model (Tessum et al., 2017). Each facility's abatement options come from two engineering menus. End-of-pipe NO_x controls come from EPA's Control Strategy Tool and South Coast staff reports, and CO_2 efficiency, fuel, and dispatch measures come from sector studies. We value local $PM_{2.5}$ exposure using the AP3 and AP4 mortality convention, the adult-mortality slope of Krewski et al. (2009), and a value of a statistical life of $10.3 million in 2019 dollars. Summed

across communities, the 2019 NO_x emissions of the 279 facilities generate $465 million per year in local mortality damages (45 deaths), or $431 million under the carbon-only market.

The quantitative stakes are large. The constrained first best, which price-taking community permits attain, gains $66 million per year over the carbon market alone and removes three times as much NO_x. We report every other regime by its welfare loss relative to that first best. A uniform tax or cap-and-trade loses $28 million by ignoring where damages occur. When each county prices as a single seller, the outcome closest to an equilibrium loses $113 million. Counties withhold permits and price at a median of four times marginal damage, erasing the first-best gain. Dividing each county's selling authority among ten sub-sellers cuts the loss to $15 million with no damage information. Setting each county's subdivision from the regulator's damage estimates cuts the loss to $11 million under full information and $8 million at our baseline estimation error. That rule outperforms a differentiated tax built from the same estimates once the error coefficient of variation exceeds 0.6 to 1.0, because communities price on their actual damages. On the distributional margin, the 29 counties containing disadvantaged communities receive nearly all of the $122 million in avoided damages and the $75 million in permit revenue.

Our paper builds on the foundational literature on environmental markets. Crocker (1966) and Dales (1968) first proposed markets in pollution rights, and Montgomery (1972) formalized the distinction between source-level *emission* licenses and receptor-level *ambient* licenses. Ambient permit systems have received limited practical attention because they require firms to assemble portfolios of receptor permits and require the regulator to choose receptor caps (Tietenberg, 1985, 1995). By combining an emission-permit layer for the aggregate pollutant with a community-owned ambient-permit layer for the local pollutant, our mechanism uses communities' supply decisions to determine local receptor caps endogenously. The closest precedent to our approach is Yang and Kaffine (2016), who study a single harmed community endowed with tradable permits for a local pollutant. We generalize that insight to multiple communities connected by an atmospheric transport matrix and embed the local market within an aggregate cap-and-trade system linked through joint abatement costs (Agee et al., 2014). We then characterize the resulting monopoly and anticommons distortions and show how within-community governance subdivisions can resolve them.

The remainder of the paper proceeds as follows. Section 2 presents the general model of a stacked permit market and characterizes its price-taking and strategic equilibria. Section 3 adapts the linear-quadratic framework of Fowlie and Muller (2019) to our two-pollutant setting and derives closed-form welfare comparisons across instruments under perfect information and damage uncertainty. Section 4 analyzes within-community subdivisions under known vs. estimated damages. Section 5 presents the empirical calibration to California's cap-and-trade program, and Section 6 concludes.

2 A Stacked Permit Market

This section presents the model environment, the efficient level and spatial allocation of abatement, and our proposed mechanism. We also derive welfare implications under differing assumptions about the competitiveness of community permit markets. Our exposition focuses on CO_2 abatement with a local co-pollutant for convenience. The same structure applies to any environmental market in which the regulator targets an aggregate level of activity while communities care about its spatial incidence. Communities may care directly about changes in economic activity (e.g., where fish are landed after catch shares are implemented) or about the associated changes in environmental quality (e.g., local pollution concentrations or wetland protections).

2.1 Model Setup

Let there be I regulated sources indexed by i and J affected communities indexed by j. Each source generates emissions e_i^A that affect the aggregate externality as well as local emissions e_i^L that affect specific communities. Abatement of each type of emissions by each firm is defined as $a_i^A \equiv \bar{e}_i^A - e_i^A$ and $a_i^L \equiv \bar{e}_i^L - e_i^L$, relative to business-as-usual levels $(\bar{e}_i^A, \bar{e}_i^L)$. The transport matrix $\mathbf{H} = [h_{ji}] \geq 0$, known to the regulator, specifies how the emissions of pollutants by each firm i affect each community j:

$$c_j = \sum_i h_{ji}\, e_i^L, \qquad c_j^0 \equiv \sum_i h_{ji}\, \bar{e}_i^L, \tag{1}$$

where c_j is community j's ambient concentration of local pollutants and c_j^0 is the business-as-usual concentration in the absence of any policy.

Aggregate and local emissions are linked because a firm's abatement decisions affect both at once, and we build this link into the abatement cost function in Assumption 1.

Assumption 1 (Joint abatement costs)**.** *Firm i's abatement cost $A_i(a_i^A, a_i^L)$ is twice continuously differentiable and strictly convex, with $A_i(0,0) = 0$ and $\nabla A_i(0,0) = \mathbf{0}$, and the two externalities are cost complements:*

$$\gamma_i\big(a_i^A, a_i^L\big) \;\equiv\; \frac{\partial^2 A_i}{\partial a_i^A\, \partial a_i^L} \;\leq\; 0. \tag{2}$$

Under Equation (2), increasing abatement of one pollutant weakly lowers the marginal cost of abating the other. Appendix B derives this cost complementarity ($\gamma_i \leq 0$) from a technology with three abatement margins. Firms can combine CO_2-only capture and NO_x-only controls (such as selective catalytic reduction) with joint measures that cut both pollutants at once (such as reducing output, switching fuels, or improving thermal efficiency). Provided the two pollutant-specific technologies have independent costs, the shared joint margin makes aggregate and local abatement weak cost complements. Empirical studies of fossil-fuel plants consistently find that CO_2 and local co-pollutant abatement are jointly produced in this manner (Burtraw et al., 2003; Considine and Larson, 2006; Agee et al., 2014).[3]

[3] Empirical modeling of jointly produced pollutants dates to Crocker (1981) and Färe et al. (2005). Bonilla et al. (2018)

Each community j has a damage function $D_j(c_j)$ that depends only on its own ambient concentration of local pollutants (Muller and Mendelsohn, 2009; Hsiang et al., 2019). We assume that $D_j(c_j)$ is increasing, weakly convex, and known only to community j. Throughout, we assume that damages from the global pollutant (CO_2) are separable from local damages and independent of where emissions occur (Holland et al., 2016). The regulator knows the constant social cost of carbon (SCC) (Rennert et al., 2022).

2.2 Planner's Problem and First Best

Next, we state the planner's problem to define the efficient level and spatial allocation of abatement across firms. We solve the planner's problem under the assumption of perfect information, so that it serves as a benchmark for our proposed mechanism. The planner chooses abatement of global and local emissions for each firm to minimize total abatement costs plus community damages, subject to the constraint that total aggregate abatement meets the pre-existing cap (Baumol and Oates, 1971; Montgomery, 1972):

$$\min_{\{a_i^A,\, a_i^L\}} \; \sum_i A_i\big(a_i^A, a_i^L\big) \; + \; \sum_j D_j\Big(\sum_i h_{ji}\big(\bar{e}_i^L - a_i^L\big)\Big) \qquad \text{s.t.} \qquad \sum_i a_i^A = R. \tag{3}$$

We assume there is a fixed emissions cap $\bar{E}^A$ that binds throughout our analysis, so required carbon abatement is determined by business-as-usual emissions in excess of the cap, $R \equiv \sum_i \bar{e}_i^A - \bar{E}^A$. We condition the planner's problem on the aggregate cap for three reasons. First, the mechanism is designed as a second layer on an existing carbon market, so holding it fixed isolates the local-margin comparison of interest. Second, with the cap binding, total carbon emissions (and carbon damages) are identical across all regimes, so the social cost of carbon drops out of every welfare comparison. Finally, the multiplier λ is the cap's implicit carbon price, so the cap-conditioned planner's problem and the full problem with carbon damages priced at SCC coincide when the cap is set optimally ($\lambda = \text{SCC}$). Corollary 1 in Section 3 establishes this point formally.

The planner's problem has a unique optimum because the objective is strictly convex and the constraint is an affine transformation of the choice variables.[4] Letting λ denote the multiplier on the carbon constraint, the first-order conditions are

$$\nabla A_i\big(a_i^A, a_i^L\big) = \begin{pmatrix} \lambda \\ \sum_j D_j'(c_j^*)\, h_{ji} \end{pmatrix} \qquad \text{for all } i. \tag{4}$$

Every firm's marginal cost of carbon abatement equals the common shadow price, and its marginal cost of local abatement equals the true marginal damage its emissions cause, summed

find that CO_2 and NO_x abatement can be cost substitutes in some settings. If $\gamma_i > 0$ for all i, the sign in Proposition 3 below reverses, so adding community permits raises rather than lowers the equilibrium carbon price.

[4] Each A_i is strictly convex in firm i's own abatement pair, so the separable sum is strictly convex in the full abatement vector. Damages add a weakly convex term, making the total objective strictly convex. The proof of Proposition 1 (Appendix A.1) establishes existence, uniqueness of the optimum and of the multiplier λ, and sufficiency of the first-order conditions.

over the communities its plume reaches. These composite marginal damages are the source-specific prices that damage-based regulation seeks to implement directly (Muller and Mendelsohn, 2009). We denote this composite marginal damage at source i by $\mathrm{MD}_i^* \equiv \sum_j D_j'(c_j^*)\, h_{ji}$.

2.3 The Mechanism and its Properties

Our mechanism stacks on top of an existing cap-and-trade scheme that caps total emissions at $\bar{E}^A$, with an equilibrium permit price of τ^A. In addition to holding a permit for every unit of aggregate emissions e_i^A, each firm must hold permits covering its contribution to ambient concentration c_j in each community j, as determined by the transport matrix $\mathbf{H}$. Each firm i must therefore hold $h_{ji}e_i^L$ permits from every j. Writing σ_j for the price of community j's permits, the effective price for local emissions facing firm i is

$$\tau_i^L = \sum_j \sigma_j h_{ji}. \tag{5}$$

The functioning of the mechanism depends on the initial allocation of ambient permits, under which each community j is given a slack allocation $\bar{L}_j \geq c_j^0$ of concentration-denominated permits that it can sell to firms as it chooses.[5] Because the initial allocation of permits exceeds baseline exposure, the equilibrium level of abatement and emissions for the local pollutant is completely determined by the communities' supply decisions, rather than by the regulator. We assume that communities sell at posted, take-it-or-leave-it prices that apply uniformly to every firm that must buy their permits. Hence, the mechanism operates through arm's-length markets rather than bilateral negotiation between firms and communities. Posted prices also prevent bilateral hold-up, because no community can tailor its terms to an individual firm's willingness to pay.

Next, we describe the equilibrium of the stacked permit system under perfect vs. imperfect competition, as well as the associated welfare implications. In each scenario, firms take all prices as given, and each firm i minimizes abatement cost plus permit expenditure,

$$\min_{a_i^A,\, a_i^L} \; A_i\big(a_i^A, a_i^L\big) + \tau^A\big(\bar{e}_i^A - a_i^A\big) + \tau_i^L\big(\bar{e}_i^L - a_i^L\big). \tag{6}$$

First, we characterize a perfectly competitive set of permit markets where communities also take prices as given. In this case, community j chooses how many permits to sell, taking σ_j as given and trading off permit revenue and the damage from the exposure its sales admit. Under the assumption of quasi-linear payoffs,[6] community j solves

$$\max_{s_j \in [0, \bar{L}_j]} \; \sigma_j s_j - D_j(s_j). \tag{7}$$

[5] In the calibration we set $\bar{L}_j = 1.5\, c_j^0$. The multiple is a modeling choice; any endowment above baseline exposure gives the same equilibrium, because firms never demand more than c_j^0.

[6] Permit revenue enters linearly, so damages are denominated in dollars, there are no income effects, and revenue received does not alter the community's marginal damage. Quasi-linearity is the standard assumption underlying willingness-to-pay measures of pollution damages (Freeman et al., 2014).

We formally define a competitive equilibrium in the context of the stacked permit system in Definition 1 and then address the existence and properties of such an equilibrium in Proposition 1.

Definition 1 (Competitive equilibrium)**.** *A competitive equilibrium of the stacked permit system is a price system* $\left(\tau^A, \{\sigma_j\}_{j=1}^J\right)$ *and an allocation* $\left(\{a_i^A, a_i^L\}_{i=1}^I, \{s_j\}_{j=1}^J\right)$ *such that (i) each firm's abatement solves Equation (6), taking* $\left(\tau^A, \{\sigma_j\}\right)$ *as given; (ii) each community's permit supply solves Equation (7), taking* σ_j *as given; (iii) the aggregate permit market clears,* $\sum_i a_i^A = R$*; and (iv) each community's permit market clears, in the sense that* $\sigma_j \geq 0$*,* $s_j \geq \sum_i h_{ji}\left(\bar{e}_i^L - a_i^L\right) = c_j$*, and* $\sigma_j\left(s_j - c_j\right) = 0$*, so that at any positive price the permits sold equal realized exposure and supply offered beyond exposure goes unsold only at a zero price.*

Proposition 1 (Efficiency and damage robustness of the competitive equilibrium)**.** *A competitive equilibrium exists. In any competitive equilibrium: (a) each community prices exposure at its own marginal damage,* $\sigma_j^* = D_j'(c_j^*)$*. (b) The allocation is the unique solution to the planner's problem in Equation (3), and the aggregate permit price equals the shadow price of the cap,* $\tau^A = \lambda$*. (c) No equilibrium object depends on the regulator's beliefs about community damages, conditional on the realized cap.*

Proof: Appendix A.1.

The logic of Proposition 1 rests on the fact that the planner's problem has a unique optimum, characterized by the first-order conditions in Equation (4). In a competitive equilibrium, communities price exposure at marginal damage and firms equate marginal abatement cost to the prices they face, so the conditions of Definition 1 reproduce the planner's first-order conditions exactly. Any competitive equilibrium must therefore coincide with the planner's unique optimum. Community damages enter the system only through communities' own supply decisions. This equilibrium is *conditional on the realized cap*, and Section 3 examines this issue further.

The competitive equilibrium illustrates how community permit endowments and posted prices can decentralize the efficient spatial allocation of abatement. In practice, however, endowing each community with exclusive control over its own exposure permits grants communities local monopoly power (Fullerton and Metcalf, 2002) while requiring firms to assemble permits across multiple jurisdictions. Together, these two features make price-taking behavior unlikely when each community acts as a unitary seller. We therefore model each community j as a unitary seller that sets its posted permit price $\sigma_j \geq 0$, taking other communities' prices $\boldsymbol{\sigma}_{-j}$ and the carbon price τ^A as given.

$$\max_{\sigma_j \geq 0} \; \sigma_j\, s_j\left(\sigma_j;\, \boldsymbol{\sigma}_{-j}, \tau^A\right) - D_j\Big(s_j\left(\sigma_j;\, \boldsymbol{\sigma}_{-j}, \tau^A\right)\Big), \qquad s_j\left(\sigma_j;\, \boldsymbol{\sigma}_{-j}, \tau^A\right) = c_j^0 - \sum_i h_{ji}\, a_i^L\Big(\tau^A, \textstyle\sum_k \sigma_k h_{ki}\Big). \tag{8}$$

Here $s_j(\sigma_j;\, \boldsymbol{\sigma}_{-j}, \tau^A)$ is the residual demand for community j's permits (and thus its realized pollution exposure) as a function of its posted price σ_j, other communities' permit prices, and the aggregate carbon price. Because each firm's local abatement a_i^L is strictly increasing in its effective local price τ_i^L, residual demand s_j is strictly decreasing in σ_j, with well-defined inverse $\sigma_j(s_j)$ on the interior range $s_j \in (0, c_j^0]$. Choosing a posted price σ_j is therefore equivalent to selecting a

point along community j's residual demand curve.[7] Unlike standard differentiated-product markets, price competition between communities does not bid markups down because communities' permits are complements rather than substitutes for regulated firms. A firm must hold permits from every community in its plume, so no community can win another community's sales by undercutting its price. Proposition 2 describes the resulting Nash equilibrium in community prices and its welfare consequences.

Proposition 2 (Strategic pricing equilibrium: existence, distortions, welfare loss)**.** *(a) A pure-strategy Nash equilibrium of the cross-community pricing game exists.*[8] *(b) Relative to the price-taking benchmark, the equilibrium embeds two distinct distortions. Each community charges a within-community monopoly markup, $\sigma_j > D'_j(c_j)$ at any interior supply, which is present even with a single community. In addition, when $J \geq 2$, each community ignores the effect of its price on other communities' permit demands. Whenever these net cross-demand externalities are positive and outweigh the community's own carbon-price feedback, markups stack above the level that a coordinated seller of all communities' permits would choose. (c) Whenever equilibrium supplies are interior, the welfare loss relative to the constrained first best is strictly positive.*

Proof: Appendix A.2.

Decentralizing permit ownership creates two compounding sources of inefficiency, the first of which is standard market power. Because each community is a monopolist over its own permits, it can strategically withhold supply to raise its price, making both prices and local abatement too high. The second inefficiency comes from the fact that firms need to assemble permits from *each* community that they are connected to via $\mathbf{H}$. This creates an "anticommons" problem (Buchanan and Yoon, 2000; Heller, 1998), whereby each community's price increase reduces the residual demand for every other community's permits. Unless communities collude to maximize joint profits, this externality pushes prices even higher and induces further over-abatement. Both inefficiencies depend on market structure and true damages, but they are unaffected by the regulator's beliefs about damages, because no damage estimate enters the mechanism anywhere.

Both results take for granted that a firm can remove its entire local stream at some price. Engineering abatement menus of the kind we use in Section 5 do not have that property. Every facility has some emissions that no listed control technology can fully abate, and a community whose receptor field contains such a facility faces a residual demand that never reaches zero. The community's revenue then rises without bound in its own price, so no best response exists. The same is true for any finite number of sub-sellers. Proposition 9 in Appendix C states this hold-up result and shows that it can be addressed by exempting the unabatable emissions from the community permit requirement. We adopt that exemption in the empirical application, and California's existing best-available-control rules provide a regulatory definition of the unabatable residual.

[7] Solving the first-order condition of Equation (8) assumes that community j knows the local slope of its residual demand curve $\partial s_j / \partial \sigma_j$.

[8] Existence requires residual demand to be continuous and reach zero at a finite price. When some tons cannot be abated at any price, demand never reaches zero and no best response exists. Proposition 9 in Appendix C states that result and an exemption on some abatement that restores existence.

Whether the market-structure inefficiencies outweigh the mechanism's informational advantages cannot be resolved at this level of generality, and quantifying the tradeoff requires the parametric structure of Section 3. First, however, we state one more general result that holds whether or not the permit market is competitive.

The final general result concerns the feedback from the community permit market to the cap-and-trade market. When abatement costs are not separable ($\gamma_i < 0$), the local abatement that community prices induce also reduces carbon emissions, so the carbon market clears at a lower price, as Proposition 3 states formally.

Proposition 3 (Layer interaction)**.** *Hold the cap R fixed. Adding community permit markets with prices $\boldsymbol{\sigma} \succeq \mathbf{0}$ lowers the market-clearing carbon price, $\tau^A(\boldsymbol{\sigma}) \leq \tau^A(\mathbf{0})$. The fall is strict whenever some firm with $\gamma_i < 0$ lies in the receptor field of a community whose price is positive. When costs are separable, so that $\gamma_i = 0$ for every firm, the carbon price does not change.*

Proof: Appendix A.3.

Proposition 3 has an important implication for our mechanism. The mechanism achieves the efficient spatial distribution of local abatement without the regulator knowing community damages. Because local abatement changes each firm's marginal cost of *aggregate* abatement, however, setting the aggregate cap optimally requires the regulator to estimate local damages. The optimal cap equates the marginal cost of aggregate abatement with the SCC. That marginal cost depends on local abatement and therefore on local damages. A simple alternative is to make the aggregate instrument a *tax* rather than a *cap*. In that case, the regulator needs to know only the SCC, which is unaffected by local damages. Corollary 1 in Section 3 states this result formally.

3 Comparison to Other Policy Instruments

Thus far, we have shown that a stacked system of community exposure permits largely avoids the informational requirements of other approaches for spatially targeted carbon policy, but at the cost of inefficiencies related to market structure. In this section, we parameterize the abatement cost function and the community damage function to facilitate direct welfare comparisons between our proposed mechanism and four leading alternatives. We consider carbon taxes and traditional cap-and-trade, both undifferentiated and differentiated by community. This parameterization enables the closed forms that enable the comparison. The empirical illustration in Section 5 replaces the quadratic costs with facility-level engineering abatement menus and solves every regime, and we note where the two economies differ.

3.1 Parameterization

We base our analysis on a generalization of the functional forms in Fowlie and Muller (2019) (FM), so that our characterizations of each policy regime are directly comparable to theirs. We state this formally in Assumption 2:

Assumption 2 (FM parameterization). $A_i = \frac{\alpha_i^A}{2}(a_i^A)^2 + \gamma_i a_i^A a_i^L + \frac{\alpha_i^L}{2}(a_i^L)^2$ *with constants satisfying Assumption 1;* $D_j(c_j) = d_j c_j$ *with* $d_j > 0$.

FM study a single non-uniformly mixed pollutant with quadratic abatement costs and a constant source-specific marginal damage that bundles all spatial variation into one reduced-form parameter. We make two adjustments to facilitate analysis of the tiered system we propose. First, rather than a single pollutant we model a uniformly mixed aggregate pollutant and a non-uniformly mixed local pollutant. Second, the two pollutants are linked on the cost side by $\gamma_i \leq 0$, a channel that does not exist in the FM single-pollutant version. Appendix E derives the equivalence between the local layer at $\gamma_i = 0$ and their model, along with each of their analytical results.[9]

For notational convenience, we express firms' price responses using three constants derived from the Hessian of the cost function:

$$\psi_i = \frac{\alpha_i^L}{\alpha_i^A \alpha_i^L - \gamma_i^2}, \qquad p_i = \frac{-\gamma_i}{\alpha_i^A \alpha_i^L - \gamma_i^2} \geq 0, \qquad \tilde{\beta}_i = \frac{\alpha_i^A}{\alpha_i^A \alpha_i^L - \gamma_i^2}. \tag{9}$$

These parameters give the slopes of firm i's linear abatement schedules with respect to emission prices. A one-dollar increase in the carbon price raises carbon abatement by ψ_i, and a one-dollar increase in the effective local price raises local abatement by $\tilde{\beta}_i$. The cross-price parameter p_i measures the co-benefit response by which a higher price on either pollutant induces additional abatement of the other. On the damage side, Assumption 2 implies a constant marginal damage d_j per unit of exposure in community j. Under this linear-quadratic structure, firm abatement, the market-clearing carbon price, community residual demands and the strategic pricing equilibrium all admit closed-form linear characterizations (Appendix A). In Lemma 1 (Appendix A.4), we establish a sufficient diagonal-dominance condition for uniqueness of the strategic pricing equilibrium. Uniqueness is guaranteed analytically whenever each community's own-price demand slope exceeds the combined cross-price feedback from other communities and from the aggregate carbon market, as occurs when the transport matrix $\mathbf{H}$ is sufficiently sparse.[10]

3.2 Welfare Comparison

Next, we compare the welfare losses of the policy regimes analyzed by FM to those associated with both the competitive and non-competitive versions of our proposed mechanism. Under Assumption 2, the efficient local price at source i is its constant marginal damage, $\text{MD}_i \equiv \sum_j d_j h_{ji}$. The efficient carbon price is the shadow price of the cap, λ. Every regime we consider confronts

[9] The cost linkage adds three channels that are absent from their analysis. First, local pricing affects the aggregate permit price (Proposition 3). Second, every regime's welfare loss includes a partial offset because the carbon market re-clears in response to local mispricing. Third, the optimal level of the aggregate cap depends on local damages (Corollary 1). All three channels vanish as $\gamma_i \to 0$. An alternative response to spatially varying co-pollutant damages is to differentiate the *carbon* instrument itself by source (Muller, 2012). We instead keep the carbon layer uniform and price the local margin directly so that community prices reflect local damages.

[10] Because atmospheric $PM_{2.5}$ transport links every source to every receptor in our empirical application, the condition fails there. On the engineering abatement menus of Section 5 the pricing game has several equilibria at every subdivision we analyze. We therefore report equilibria as sets in Section 5 for the calibrated economy.

firms with some pair of prices. Proposition 4 establishes that each regime's welfare loss can be written as a function of the *wedges* between those prices and the efficient prices from the planner's problem.

Proposition 4 (Welfare-loss identity)**.** *For any regime with wedges* $\Delta^A = \tau^A - \lambda$ *and* $\Delta_i^L = \tau_i^L - \mathrm{MD}_i$*, market clearing implies* $\Delta^A = -\sum_i p_i \Delta_i^L / \sum_i \psi_i$ *and*

$$\mathcal{L} \;=\; \frac{1}{2}\sum_i \tilde{\beta}_i (\Delta_i^L)^2 \;-\; \frac{\left(\sum_i p_i \Delta_i^L\right)^2}{2\sum_i \psi_i} \;\geq\; 0. \tag{10}$$

Proof: Appendix A.5.

Equation (10) has two important features that are central to the rest of our welfare comparisons. First, only local mispricing matters, because market clearing implies that the carbon wedge is mechanically determined by the local wedges. Second, the negative correction term is essentially a carbon-market rebate. Local prices affect how much carbon firms co-abate, so when those prices are wrong the carbon price re-clears in response to the error. This re-clearing undoes the common component of the mispricing, so only its dispersion across sources appears in net welfare loss. For this reason, the entries in Table 1 depend on the dispersion of pricing errors across sources rather than their level, and estimation errors that shift all sources in the same direction are less costly.

We compare five regimes implemented by a regulator who holds beliefs $\hat{d}_j$ about community damages, with perfect information as the special case $\hat{d}_j = d_j$. Regimes I and II are undifferentiated and differentiated cap-and-trade systems, respectively, while Regimes III and IV are undifferentiated and differentiated emissions taxes. Firms face uniform local prices under Regimes I and III and source-specific prices determined by the regulator's damage estimates under Regimes II and IV.[11] Regime V is our proposed mechanism, in its price-taking and strategic variants. For each regime, Table 1 reports the local price facing source i, the welfare loss from evaluating Equation (10) at the implied wedges, and whether that loss depends on the regulator's beliefs. Proposition 5 states the welfare comparison across instruments for a given estimate of community damages $\hat{\mathbf{d}}$.

Proposition 5 (Damage robustness and regime ranking)**.** *Fix any beliefs* $\hat{\mathbf{d}}$ *and condition on the realized aggregate cap. (a) The price-taking equilibrium of Regime V attains zero welfare loss for every* $\hat{\mathbf{d}}$*, and no other regime in Table 1 has this property. (b) The strategic-pricing loss of Regime V does not depend on* $\hat{\mathbf{d}}$*. (c) In Regimes I–IV the local wedge inherits the belief error (*$\Delta_i^L = b_i$ *under differentiation and* $\Delta_i^L = t^*(\hat{\mathbf{d}}) - \mathrm{MD}_i$ *under uniformity), and when* $\gamma_i < 0$ *the error also displaces the carbon price through the linkage equation in Proposition 4. (d) Neither Regime V under strategic pricing nor Regime IV dominates the other. At* $\hat{\mathbf{d}} = \mathbf{d}$*, Regime IV attains zero loss while the strategic-pricing loss is positive at interior supplies. As the belief error grows without bound in any fixed direction with positive loss, Regime IV's loss*

[11] In practice, differentiated caps are implemented with trading ratios, and even optimally chosen ratios are generally suboptimal under asymmetric information (Holland and Yates, 2015), and integrated multipollutant ratio systems embed the same damage estimates (Reeling et al., 2018).

Table 1. Regime Comparison Under Damage beliefs $\hat{d}_j$

Regime	Local price at source i	Welfare loss $\mathcal{L}$ (eliminated form, Equation (10))	Robust to $\hat{d}_j$?
I. Undifferentiated cap-and-trade	$t^*(\hat{\mathbf{d}})$	$\frac{1}{2}\sum_i \tilde{\beta}_i\big(t^* - \mathrm{MD}_i\big)^2 - \dfrac{\big[\sum_i p_i\big(t^* - \mathrm{MD}_i\big)\big]^2}{2\sum_i \psi_i}$	No
II. Differentiated cap-and-trade	$\widehat{\mathrm{MD}}_i$	$\frac{1}{2}\sum_i \tilde{\beta}_i\, b_i^2 - \dfrac{\big(\sum_i p_i\, b_i\big)^2}{2\sum_i \psi_i}$	No
III. Undifferentiated tax	$t^*(\hat{\mathbf{d}})$	$\frac{1}{2}\sum_i \tilde{\beta}_i\big(t^* - \mathrm{MD}_i\big)^2 - \dfrac{\big[\sum_i p_i\big(t^* - \mathrm{MD}_i\big)\big]^2}{2\sum_i \psi_i}$	No
IV. Differentiated tax	$\widehat{\mathrm{MD}}_i$	$\frac{1}{2}\sum_i \tilde{\beta}_i\, b_i^2 - \dfrac{\big(\sum_i p_i\, b_i\big)^2}{2\sum_i \psi_i}$	No
V. Community (price-taking)	MD_i	0	Yes
V. Community (strategic)	$\mathrm{MD}_i + \Delta_i$	$\frac{1}{2}\sum_i \tilde{\beta}_i \Delta_i^2 - \dfrac{\big(\sum_i p_i \Delta_i\big)^2}{2\sum_i \psi_i}$	Yes

Notes: Each row reports the local price facing source i, the welfare loss relative to the constrained first best from evaluating Equation (10) at the implied wedges, and whether the equilibrium outcome depends on the regulator's damage beliefs $\hat{d}_j$. $b_i = \sum_j (\hat{d}_j - d_j) h_{ji}$ is the belief bias at source i; $\Delta_i = \sum_j \mu_j h_{ji}$, where $\mu_j \geq 0$ is community j's equilibrium markup over marginal damage (Lemma 1); t^* is the optimal uniform local price (Appendix D). Under known costs the cap-and-trade and tax versions coincide (I = III, II = IV); they differ under cost uncertainty (Appendix C).

diverges while the strategic-pricing loss is unchanged.

Proof: Appendix A.6.

Part (d) establishes that the comparison between the market-power losses and the informational advantages of the mechanism has an ambiguous sign in general. We can, however, provide some intuition about when each wedge is likely to dominate. First, the information wedge scales with the regulator's error, while the market-power wedge does not depend on the error. Hence, the mechanism is more likely to be favored when community damages are harder to estimate. Second, by the carbon-market rebate, only the *dispersion* of the error across sources survives in $\mathcal{L}(\{b_i\})$, so the mechanism's advantage is largest when estimation errors are idiosyncratic across space rather than a common mis-scaling of all damages. Third, by Lemma 1 (Appendix A.4), community j's markup equals one-half the gap between the choke price of its residual demand and its own marginal damage. Market power is therefore most severe for communities facing large permit demand relative to their marginal damage, and mildest where marginal damages are high. Finally, the stacking externality grows with the number of communities that a firm must obtain permits from, so the market-power wedge rises as communities are defined more finely.

Proposition 5 applies at a fixed belief error and therefore does not establish a general ranking of the regimes, because at any fixed $\hat{\mathbf{d}}$ the comparison between uniform and differentiated instruments depends on the particular error realized. Proposition 6 instead treats the error as a random variable and states the comparison in expectation. To do so, we assume that the regulator must implement their instrument before damages are known, basing it on an estimate that equals

the truth plus an estimation error with mean zero and known covariance. We then compare the regimes by their expected welfare losses, averaged over the possible errors. Throughout, communities act on the damages they actually experience.[12] To streamline notation, we write $\mathcal{L}^{\text{diff}}$ for the welfare loss of the differentiated instruments (Regimes II and IV, which coincide) and $\mathcal{L}^{\text{undiff}}$ for that of the undifferentiated instruments (Regimes I and III).

Proposition 6 (Expected losses under damage uncertainty)**.** *Suppose the regulator implements a regime using beliefs $\hat{\mathbf{d}} = \mathbf{d} + \boldsymbol{\varepsilon}$, where the estimation error $\boldsymbol{\varepsilon}$ has mean zero and covariance Σ, and let $b_i = \sum_j \varepsilon_j h_{ji}$ be the induced bias at source i. (a) The expected loss of the differentiated instruments is the loss formula of Table 1 with variances in place of squared wedges,*

$$\mathbb{E}[\mathcal{L}^{\text{diff}}] \;=\; \frac{1}{2}\sum_i \tilde{\beta}_i \operatorname{Var}(b_i) \;-\; \frac{\operatorname{Var}\big(\sum_i p_i b_i\big)}{2\sum_i \psi_i}. \tag{11}$$

(b) Under the uniform instruments, the estimation error affects the loss only through the uniform price $t^*(\hat{\mathbf{d}})$, so

$$\mathbb{E}[\mathcal{L}^{\text{undiff}}] \;=\; \mathcal{L}_0^{\text{undiff}} \;+\; \frac{\Lambda}{2}\operatorname{Var}\big(t^*(\hat{\mathbf{d}})\big), \tag{12}$$

where $\mathcal{L}_0^{\text{undiff}}$ is undifferentiated instruments' loss under perfect information and $\Lambda = \sum_i \tilde{\beta}_i - (\sum_i p_i)^2 / \sum_i \psi_i$ is the curvature of the welfare loss with respect to a uniform price. For any error distribution, the error term in (b) is no larger than the expected loss in (a), with equality only if the error shifts all sources uniformly. Under the scaling $\Sigma = \mathrm{CV}^2\,\Sigma_0$, the expected gain from differentiation is strictly decreasing in CV and, provided $\mathcal{L}_0^{\text{undiff}} > 0$, changes sign at a unique threshold $\mathrm{CV}^* > 0$ whenever that inequality is strict. (c) By Proposition 5, the community mechanism is unaffected because the price-taking equilibrium has zero loss at every realization, and the monopoly loss (one seller per community) is constant in $\boldsymbol{\varepsilon}$.

Proof: Appendix A.7.

Remark 1 (Finite abatement menus)**.** *The threshold CV^* in Proposition 6(b) is a property of the quadratic cost form of Assumption 2, under which a source's loss from a mispriced tax grows with the square of the pricing error without bound. When a source instead chooses from a finite menu of control options, as in the engineering data of Section 5, its response to a price is a step function. In that setting, a pricing error only affects welfare if it crosses one of these steps. When this occurs, the cost is bounded by the net cost of the step, so the differentiated loss is bounded by the menu rather than growing in CV^2. The uniform price loses its averaging advantage for the same reason. Because the welfare-maximizing uniform price lies at a breakpoint of the aggregate menu, a shift in beliefs can move it to a different breakpoint. Both departures weaken the case for the reversal, and in the simulations of Section 5 it does not occur within any error level we examine.*

[12] The same calculation applies when the damages themselves are the random object and the regulator's estimate is their mean. The comparison of the subdivision with the differentiated instruments under matched information in Section 4 uses this second framing. Nothing in the expected losses depends on which object is held fixed, because every loss below is quadratic in the error and depends only on its covariance.

Proposition 6 shows that the expected-loss ranking between uniform and differentiated regulation depends on how each instrument transmits estimation error into source prices. Because an optimal uniform price depends only on a weighted average of estimated marginal damages across sources, independent community-level estimation errors largely cancel out in the aggregate. Under differentiated taxes or trading ratios, by contrast, each community's estimation error passes directly into the prices faced by upwind sources without averaging out. This tradeoff is particularly relevant in light of recent general-equilibrium estimates showing substantial potential welfare gains from spatially differentiating U.S. air pollution policy (Hollingsworth et al., 2026). Those potential gains are driven by wide spatial dispersion in marginal damages, which is also the condition under which noisy damage estimates can generate large pricing errors.

Finally, we return to a point raised in Section 2. Damage-estimation error also affects the level of the aggregate cap, so it can cause welfare losses even when the community permit market is perfectly competitive. Here, we use the FM parameterization to formally compare the performance of cap-and-trade and taxes for the *aggregate* layer when the abatement cost feedbacks from the stacked community layer are accounted for.

Corollary 1 (Cap versus tax for the aggregate instrument)**.** *Suppose the regulator also chooses the level of the aggregate instrument, and assume the mechanism attains the* full *first best when, in addition to the conditions of Proposition 1, the shadow price of carbon equals the SCC. (a) The optimal cap must anticipate the carbon co-abatement induced by community pricing, with required abatement*

$$R^* \;=\; \Big(\sum_i \psi_i\Big)\mathrm{SCC} \;+\; \sum_i p_i \sum_j d_j h_{ji}, \tag{13}$$

and therefore requires an estimate of community damages. (b) A cap set under beliefs $\hat{\mathbf{d}}$ *misprices carbon by* $\sum_i p_i b_i / \sum_i \psi_i$*, where* $b_i \equiv \sum_j (\hat{d}_j - d_j) h_{ji}$ *is the regulator's damage bias at source* i*, generating a carbon-side welfare loss of*

$$\mathcal{L}^{\mathrm{cap}} \;=\; \frac{\big(\sum_i p_i b_i\big)^2}{2\sum_i \psi_i}. \tag{14}$$

(c) An aggregate tax set at the SCC requires no damage estimate, and the pair (aggregate tax, competitive community permits) attains the full first best for any regulator beliefs.

Proof: Appendix A.8.

This comparison could not arise under separable abatement costs, where the cap could be set without reference to the local layer. When technology links the two pollutants, quantity-based aggregate regulation allows damage forecast errors to enter indirectly through the cap level, and the tax-based variant of our mechanism avoids this problem. The cap–tax asymmetry is in line with the prices-versus-quantities tradition (Weitzman, 1974), except that here it arises from information about a *different* pollutant.

4 Resolving Market Power Distortions

In this section, we show how dividing permit-selling authority within each community can counteract the market-power distortions introduced by our proposed mechanism. We first characterize the theoretical benchmark in which a full-information subdivision exactly restores the constrained first best, and we then analyze two practical implementations that either use noisy damage estimates or avoid damage information altogether.

4.1 Within-Community Subdivisions

Rather than vesting selling authority in a single community-wide body, the regulator can divide selling authority within community j among n_j independent local sub-sellers, such as neighborhood councils, municipal wards, or special-purpose local districts. These bodies are analogous to business improvement districts or community facilities districts that exercise localized fiscal authority within a city or county. The market institution differs across the two margins because of the physical structure of pollution exposure. Across communities, exposure permits are complements because each firm must assemble h_{ji} permits from every downwind community j, giving each community a distinct compliance product and its own price σ_j. Within community j, by contrast, all permits cover the same receptor and are homogeneous substitutes for regulated firms. Under a within-community subdivision, the community therefore operates a single clearinghouse (or uniform-price consignment auction) that quotes one market-clearing price σ_j to firms. Each sub-seller $\ell \in \{1, \dots, n_j\}$ independently consigns an authorized permit quantity $s_{j\ell} \geq 0$ from the community's slack pool $\bar{L}_j$, much as participating jurisdictions or utilities consign allowance budgets into the joint Western Climate Initiative or Regional Greenhouse Gas Initiative auctions. The clearinghouse then sets the uniform price $\sigma_j(s_j)$ at which firms' residual demand absorbs total consigned supply $s_j = \sum_\ell s_{j\ell}$. Firms continue to face a single posted price σ_j in each community and solve the same compliance problem as in Section 2.[13]

Assuming the n_j sub-sellers represent equal population shares and therefore bear equal shares $1/n_j$ of community j's pollution damages, each sub-seller solves:

$$\max_{s_{j\ell} \geq 0} \; \sigma_j\big(s_{j\ell} + S_{j,-\ell}\big)\, s_{j\ell} \; - \; \frac{1}{n_j} D_j\big(s_{j\ell} + S_{j,-\ell}\big), \qquad s_j = \sum_\ell s_{j\ell}. \tag{15}$$

In the static equilibrium of Equation (15), each sub-seller knows its constituents' damage share d_j/n_j, the clearinghouse price σ_j, and the local slope $\sigma_j'(s_j)$ of the community's inverse residual demand curve. The clearinghouse can report that slope directly if firms submit demand schedules.[14] Subdividing selling authority in this way introduces two incentives that lead sub-sellers to

[13] If sub-sellers instead posted competing prices on an unallocated common pool of identical community-j permits, firms would buy exclusively from the lowest-priced seller, and Bertrand competition would push the community price to zero even with two sellers. Operating a single clearinghouse into which local boards consign quantities preserves a uniform community price while allowing sub-sellers' supply decisions to discipline local markups.

[14] Administratively, this requires a posted community price and a record of each body's authorizations and receipts.

authorize more permits in aggregate than a unitary authority would. First, each sub-seller collects the full market price σ_j on every permit it authorizes, but its constituency bears only $1/n_j$ of the resulting pollution damage. This common-pool externality weakens each sub-seller's incentive to restrict exposure relative to a unitary community seller. Second, when expanding its own authorization $s_{j\ell}$, each sub-seller internalizes the price reduction $-\sigma_j' s_{j\ell}$ on its own sales rather than on the community's total sales s_j, which weakens its incentive to withhold supply.[15]

Figure 1 illustrates how dividing a community's selling authority among sub-sellers erodes community-level markups. At an interior equilibrium, every sub-seller's first-order condition for Equation (15) gives the same offer, so the within-community equilibrium is symmetric. Summing the n_j first-order conditions gives a single community-level pricing rule, which at $n_j = 1$ is the monopoly equilibrium of Section 2. The community as a whole behaves like an oligopoly of n_j sellers sharing one market. Its equilibrium price declines in n_j from the monopoly price at $n_j = 1$ toward zero as n_j grows large and each seller's damage share vanishes. The free-riding that the subdivision creates *within* communities thus counteracts the monopoly withholding and anticommons stacking that operate *between* communities. This approach also maps onto familiar governance structures, such as multi-member commissions or resident voting bodies whose members authorize supply independently.

Devolving price-setting authority within communities also helps address a potential political economy and environmental justice concern with our mechanism. Some decision-making body must represent "the community" when setting prices for pollution exposure. Depending on the scale at which a community is defined (a topic we return to in Section 5), this setup may itself induce various political tensions and questions of procedural justice within communities. Dividing each community's selling authority among sub-sellers helps address these concerns by making decision-making even more local (e.g., moving from a board of county commissioners to neighborhood associations, school districts, or other pre-existing bodies).[16]

Proposition 7 states the equilibrium of the within-community pricing scheme and the number of sub-sellers $n_j^\star$ at which the two forces exactly offset. At $n_j^\star$ the community-level price is efficient even in the absence of perfect competition between communities.

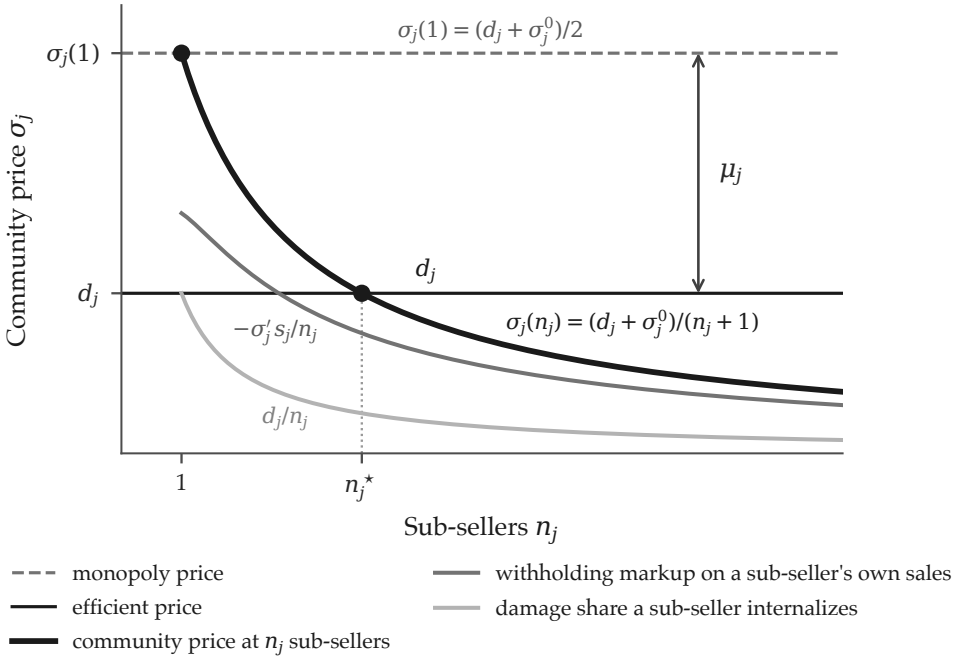
Proposition 7 (Optimal subdivision)**.** *Fix a committed subdivision* $\mathbf{n} = (n_1, \ldots, n_J)$*. (a) Given rivals' prices* $\boldsymbol{\sigma}_{-j}$ *and the carbon price* τ^A*, community* j*'s equilibrium price is*

$$\sigma_j(n_j) \;=\; \frac{1}{n_j + 1}\left(d_j \;+\; \frac{c_j^0 - \sum_{k \neq j} \left(\sum_i \tilde{\beta}_i h_{ji} h_{ki} \right) \sigma_k - \left(\sum_i p_i h_{ji} \right) \tau^A}{\sum_i \tilde{\beta}_i h_{ji}^2} \right), \tag{16}$$

[15] Note that the optimization problem in Equation (15) requires no cap on an individual sub-seller's authorization. Because the community's total endowment $\bar{L}_j$ exceeds baseline exposure c_j^0, firms never demand more than c_j^0 permits in total. Any excess supply authorized beyond c_j^0 pushes the clearing price to zero, generating zero revenue while still exposing the sub-seller's constituents to baseline damages.

[16] Both n_j and the geographic definition of a community are important parts of the governance design, because community size trades off local heterogeneity and cross-boundary spillovers (Costello and Kaffine, 2018).

Figure 1. Dividing a Community's Selling Authority Among Sub-sellers



Notes: A single stylized community with linear residual inverse demand and constant marginal damage d_j, where σ_j^0 denotes the choke price of residual demand (the bracketed ratio in Equation (16)). The community price at n_j sub-sellers, $\sigma_j(n_j)$, is given by Proposition 7(a). It starts at the monopoly price, a markup of μ_j over the efficient price d_j, and reaches d_j at $n_j^\star$. The two lower curves plot the two terms of a sub-seller's first-order condition, namely the damage share d_j/n_j it internalizes and the markup $-\sigma_j' s_j/n_j$ it earns on its own sales. Their sum equals the community price, and both fall as sub-sellers are added.

which is strictly decreasing in n_j and coincides with the unitary monopoly price at $n_j = 1$. (b) Let

$$n_j^\star \;=\; \frac{c_j^0 - \sum_{k\neq j}\left(\sum_i \tilde{\beta}_i h_{ji} h_{ki}\right) d_k - \left(\sum_i p_i h_{ji}\right)\lambda}{d_j \sum_i \tilde{\beta}_i h_{ji}^2}, \tag{17}$$

the same expression with rivals priced at their marginal damages and the carbon market at its first-best price λ. At $\mathbf{n} = \mathbf{n}^\star$, the profile $\sigma_j = d_j$ for every j is a joint equilibrium, and it is the unique one under a diagonal-dominance condition stated in Appendix A.4. The within-community markup is competed away and the allocation attains the constrained first best.

Proof: Appendix A.9.

The intuition behind Proposition 7 follows directly from the sub-sellers' symmetric first-order conditions. Each sub-seller equates its perceived marginal revenue, $\sigma_j + \sigma_j'(s_j)s_{j\ell}$, to its share of marginal damage, d_j/n_j. Summing across the n_j sub-sellers shows that the equilibrium community price equals the sum of the residual-demand choke price σ_j^0 and the community's marginal damage d_j, divided by n_j+1 as depicted in Figure 1. Setting this equilibrium price equal to d_j gives the first-best number of sub-sellers, $n_j^\star = \sigma_j^0/d_j$, where the choke price σ_j^0 is evaluated at rivals' marginal damages and $\tau^A = \lambda$. With fewer than $n_j^\star$ sub-sellers, monopoly withholding dominates and $\sigma_j > d_j$; with more than $n_j^\star$ sub-sellers, within-community free-riding dominates and $\sigma_j < d_j$. When every community is divided into $n_j^\star$ sub-sellers simultaneously, marginal-damage pricing in every community is a mutually consistent fixed point of the cross-community game, eliminating both the local monopoly markup and the cross-community anticommons externality.

Proposition 7 serves as a theoretical benchmark demonstrating that within-community institutional design can, in principle, completely neutralize between-community market power. Equation (17) also clarifies what determines the ideal degree of subdivision across places. The benchmark $n_j^\star$ is large where firms' baseline permit demand is high relative to local marginal damage (where a unitary seller would exercise the most market power), and small where marginal damages are high relative to permit demand.[17] Because $n_j^\star$ depends on true marginal damages $\mathbf{d}$, however, a regulator who does not know community damages cannot implement $\mathbf{n}^\star$ directly. We next examine two operational policies available to an imperfectly informed regulator, one that uses noisy damage estimates to set $\hat{n}_j$ and one that sets a single uniform subdivision $\bar{n}$ across all communities.

4.2 Estimating the Subdivision

In practice, $n_j^\star$ cannot be implemented directly because it depends on community damages, and so it raises the same information problem that the community permit mechanism is meant to avoid. Because the subdivision and the differentiated instruments require the same damage estimates,

[17] Communities whose residual-demand choke price lies below d_j have $n_j^\star \leq 1$ and would not be subdivided. Although Proposition 7 uses the linear-quadratic parameterization of Assumption 2, Proposition 10 in Appendix C shows that an exact first-best subdivision $n_j^\star = 1 + (-\sigma_j'(s_j^\star))s_j^\star/D_j'(c_j^\star)$ exists under the general convex cost and damage assumptions of Section 2.

the question is whether one approach does better than the other in the presence of incomplete information. To make this comparison, we assume that the regulator uses its estimate $\hat{\mathbf{d}} = \mathbf{d} + \boldsymbol{\varepsilon}$, with $\boldsymbol{\varepsilon}$ mean zero and covariance Σ, while communities price at the damages they actually experience (as in Proposition 6). The regulator must select a subdivision $\hat{n}_j$ from its estimates:

$$\hat{n}_j \;=\; \frac{c_j^0 - \sum_{k\neq j}\left(\sum_i \tilde{\beta}_i h_{ji} h_{ki}\right)\hat{d}_k - \left(\sum_i p_i h_{ji}\right)\tau^A_{\hat{\mathbf{d}}}}{\hat{d}_j \sum_i \tilde{\beta}_i h_{ji}^2}, \tag{18}$$

where $\tau^A_{\hat{\mathbf{d}}}$ is the carbon price that clears the aggregate market when damages equal their estimates. Equation (18) is the limited-information analog of Equation (17), with estimated damages in place of known damages. This substitution introduces potential welfare losses when damages are estimated with error, as with differentiated taxes or cap-and-trade. Proposition 8 states the welfare comparison across regimes when damages are unknown and must be estimated.

Proposition 8 (Information aggregation under damage uncertainty)**.** *Suppose the regulator implements the subdivision at* $\hat{\mathbf{n}}$ *from Equation* (18) *and communities then price knowing their realized damages. Let* $\varepsilon_j = \hat{d}_j - d_j$ *denote the estimation error, as in Proposition 6. (a) Under the diagonal-dominance condition of Appendix A.4, equilibrium prices are affine in realized damages. When cross-community spillovers are negligible, each community's price is the convex combination*

$$\sigma_j = (1-\theta_j)\,\hat{d}_j + \theta_j\, d_j, \qquad \theta_j = \frac{1}{\hat{n}_j + 1}, \tag{19}$$

blending the regulator's estimate with the community's real damages. (b) Define the damped bias $b_i^\theta \equiv \sum_j (1-\theta_j)\,\varepsilon_j h_{ji}$*. In the same negligible-spillover limit as part (a), the mechanism's expected loss is the differentiated instruments' formula in Equation* (11) *evaluated at the damped bias,*

$$\mathbb{E}[\mathcal{L}^{\text{mech}}] \;=\; \frac{1}{2}\sum_i \tilde{\beta}_i \,\mathrm{Var}(b_i^\theta) \;-\; \frac{\mathrm{Var}\big(\sum_i p_i b_i^\theta\big)}{2\sum_i \psi_i}. \tag{20}$$

(c) In that limit, if estimation errors are independent across communities with positive variances, then $\mathbb{E}[\mathcal{L}^{\text{mech}}] < \mathbb{E}[\mathcal{L}^{\text{diff}}]$*. If* $\theta_j = \bar{\theta}$ *for all* j*, then* $\mathbb{E}[\mathcal{L}^{\text{mech}}] = (1-\bar{\theta})^2\,\mathbb{E}[\mathcal{L}^{\text{diff}}]$ *for any error distribution.*

Proof: Appendix A.11.

Proposition 8 implies that, under uncertainty, the community mechanism is likely to produce smaller welfare losses than the differentiated instruments. Under the differentiated instruments, the regulator's estimate and any associated error pass into prices one-for-one. Under our mechanism, the regulator's damage estimate sets only the subdivision, and communities then price on their *actual* damages.[18] Section 5 evaluates this error-damping channel on the calibrated abatement menus.

[18]The within-community subdivision can fail to dominate when highly heterogeneous θ_j combine with strongly negatively correlated errors. Uneven damping then distorts the spatial pattern of the surviving error across communities. Whether that case is empirically relevant depends on the joint distribution of damages.

4.3 A Uniform Subdivision

Although the subdivision $\hat{n}_j$ dampens the regulator's estimation error, it still requires J damage estimates. Here, we consider a simpler alternative. Instead of trying to calibrate a unique subdivision for each community, the regulator can give every community the same number of sub-sellers, $\bar{n}$, which requires no damage information. Corollary 2 characterizes the equilibrium under a uniform subdivision and establishes that $\bar{n}$ can be chosen without damage estimates.

Corollary 2 (Uniform subdivisions)**.** *Fix a uniform subdivision $n_j = \bar{n}$ for every j, and let $n_j^\star(\bar{n})$ denote the expression in Equation (17) evaluated at the carbon price that clears the aggregate market under that subdivision rather than at λ; the two coincide when $p_i = 0$ for every i. In the negligible-spillover limit of Proposition 8: (a) community j's equilibrium price is*

$$\sigma_j(\bar{n}) \;=\; d_j \,\frac{n_j^\star(\bar{n})+1}{\bar{n}+1}, \tag{21}$$

so its markup over marginal damage, $\mu_j(\bar{n}) = d_j\,(n_j^\star(\bar{n})-\bar{n})/(\bar{n}+1)$, falls in magnitude at rate $1/(\bar{n}+1)$, changes sign at $\bar{n} = n_j^\star(\bar{n})$, and makes the aggregate welfare loss U-shaped in $\bar{n}$; (b) for every community with $n_j^\star(\bar{n}) \geq \bar{n}$, the markup is at most $2/(\bar{n}+1)$ times the unitary markup at the same residual demand, $d_j\,(n_j^\star(\bar{n})-1)/2$, so that community's contribution to the welfare loss falls at least at the rate $4/(\bar{n}+1)^2$; (c) no element of the equilibrium depends on the regulator's damage beliefs, so the welfare loss under a uniform subdivision is deterministic rather than expected; (d) $\bar{n} = 1$ is the monopoly equilibrium and, as $\bar{n} \to \infty$, every community's price falls to zero and the equilibrium converges to the carbon-only status quo, so that the overshoot wedge is bounded by $|\mu_j(\bar{n})| < d_j$ whenever $\bar{n} \geq n_j^\star(\bar{n})$, whereas the undershoot wedge for $\bar{n} < n_j^\star(\bar{n})$ is bounded only by the unitary markup.

Proof: Appendix A.10.

By Corollary 2, the markup in every community shrinks at rate $1/(\bar{n}+1)$ and its contribution to the welfare loss shrinks at rate $1/(\bar{n}+1)^2$ as $\bar{n}$ grows large.[19] However, the markup changes sign at the community's own $n_j^\star$, so a single $\bar{n}$ overshoots communities with high damages relative to demand while it still undershoots communities with low damages relative to demand. The two errors are not symmetric, because adding sub-sellers beyond $n_j^\star$ pushes prices toward zero and the equilibrium toward the carbon-only status quo. The cost of too many sub-sellers is therefore bounded by the status-quo welfare loss of having no local instrument at all. Too few sub-sellers, by contrast, leave a markup that is bounded only by the monopoly loss.

Because the free-riding discount takes the same form in every community, we can bound the undershoot loss under linear residual demand without knowing community damages. In any

[19] Corollary 2 is stated for the linear parameterization. Its explicit rates (the $1/(\bar{n}+1)$ decline in the markup and the $4/(\bar{n}+1)^2$ decline in each community's loss) rely on that structure, but its qualitative content does not. Under the maintained assumptions of Section 2 alone, Proposition 10(b) establishes that the within-community price is strictly decreasing in the number of sub-sellers and crosses marginal damage exactly once, at $n_j^\star$. A uniform subdivision therefore still overshoots the communities with $n_j^\star < \bar{n}$, undershoots the rest, and depends on no damage estimate. Its outcome is still bracketed by the monopoly outcome at one sub-seller and by the carbon-only status quo, which it approaches as each sub-seller's damage share and price impact vanish with the number of sub-sellers.

community with $n_j^\star \geq \bar{n}$, at most $4/(\bar{n}+1)^2$ of the unitary market-power loss survives (3.3 percent at $\bar{n} = 10$ and 0.9 percent at $\bar{n} = 20$). On the discrete engineering menus of Section 5, the aggregate loss is U-shaped in $\bar{n}$ with a minimum near ten sub-sellers. At $\bar{n} = 3$ to 10, low-damage rural counties still price above d_j while high-damage urban counties already price near or below d_j. Beyond twenty sub-sellers, urban sub-sellers with small damage shares give permits away across whole steps of demand, pushing the aggregate loss back toward the carbon-only status quo.

5 Empirical Illustration: California's Carbon Market

5.1 Calibration

We calibrate the model to California's cap-and-trade program in 2019, covering $I = 279$ large stationary sources including power plants, petroleum refineries, cement kilns, hydrogen plants, and other industrial facilities.[20] To construct the source-receptor transport matrix $\mathbf{H}$, we use facility-specific runs of the InMAP reduced-complexity atmospheric transport model (Tessum et al., 2017), which maps each facility's NO_x emissions into ambient $PM_{2.5}$ concentrations across 9,109 California census tracts.[21] To parameterize marginal damages d_j, we follow the mortality convention of the AP3 and AP4 integrated assessment models (Muller and Mendelsohn, 2009). We multiply each tract's 2020 Census population at risk (adults aged 30 and over and infants) by county age-group death rates and the adult-mortality slope of Krewski et al. (2009), $\ln(1.06)/10$ per $\mu g/m^3$. We value mortality at a value of a statistical life of $10.29 million in 2019 dollars, which carries EPA's $7.4 million (2006 dollars) forward for inflation and real income growth.[22] Evaluating total local damages $\sum_j d_j c_j$ at the 2019 NO_x emissions of the 279 facilities gives $465 million per year (45 deaths), or $431 million under the carbon-only status quo. The median facility's marginal damage is $14,093 per short ton of NO_x.

We focus on NO_x as the local co-pollutant of CO_2 regulation for three reasons. First, NO_x is a primary criteria pollutant emitted by stationary fossil-fuel sources in California. Second, atmospheric NO_x reacts to form secondary particulate matter ($PM_{2.5}$), the dominant driver of air-pollution mortality damages. Third, facility-level changes in greenhouse gas emissions are strongly correlated with changes in NO_x emissions (Burtraw et al., 2003), as we confirm in our

[20] The sample is the 287 facilities for which we have atmospheric transport runs, minus eight that are absent from 2019 reporting, report no covered emissions, or are reporting artifacts. The 279 account for 99.9 percent of the covered emissions that emitters reported in 2019 and for 30 percent of the program's covered emissions once fuel suppliers are included. We take NO_x from CARB's criteria-pollutant inventory and greenhouse gases from its mandatory reporting data, and we fill six facilities' missing 2019 NO_x from their 2018 emission ratios (2.05 percent of the sample total).

[21] Each run releases 1,000 short tons of NO_x at ground level and records the change in $PM_{2.5}$ at every tract. We count exposure at California receptors only, which omits the 14 percent of the median facility's population exposure that falls outside the state.

[22] Section 5.5 reports the 2.5th and 97.5th percentiles of the AP4 VSL distribution ($0.94 million and $25.6 million) and the steeper slope of Di et al. (2017) for the population aged 65 and over. We count $PM_{2.5}$ mortality only and exclude ozone and morbidity. In EPA's 2024 benefit-per-ton technical support document (Table 8, 2 percent discount rate), ozone mortality damages per ton of NO_x are $49,400 to $98,900 across our sectors in 2019 dollars. Those figures are 2.7 to 12.8 times the corresponding $PM_{2.5}$ mortality estimates.

2018–2019 California facility data in Figure G1 (Appendix G). Because we abstract from other co-pollutants such as SO_2 and direct primary $PM_{2.5}$, our \$465 million baseline is a lower bound on total co-pollutant damages from covered facilities.

There is no quasi-experimental price variation available to identify abatement costs across facilities. Instead, we take the three abatement channels of Appendix B from engineering data and keep them as discrete menus rather than fitting the quadratic forms of Assumption 2. First, we measure each facility's co-emission ratio ϕ_i as its 2019 ratio of NO_x to covered CO_2 emissions in the California Air Resources Board (CARB) mandatory reporting data. Second, for NO_x-only abatement we build a frontier of end-of-pipe controls (selective catalytic reduction, low-NO_x burners, and the other measures in EPA's Control Strategy Tool and the South Coast Air Quality Management District's BARCT assessments (U.S. Environmental Protection Agency, 2020; South Coast Air Quality Management District, 2021)). We assign each facility the controls listed for its equipment beyond those it already had in 2019. This corresponds to a single control for most facilities and a ladder of two to eight for refineries and the largest units (each priced over the incremental tons it would remove) in 2019 dollars. The resulting frontier is capable of removing at most 63 percent of the sample's NO_x, and 91 facilities have no option beyond their 2019 controls. Because engineering menus understate the costs that firms reveal in permit markets (Fowlie, 2010; Fowlie and Muller, 2019), we report every result under the engineering costs and under the two multipliers, 1.85 and 2.39. These are drawn from Fowlie and Muller's estimates of the difference between observed vs. engineering costs in their setting.[23]

Third, for joint CO_2–NO_x abatement we build a per-facility menu of efficiency and fuel measures from engineering sources by sector. These include heat-rate data and benchmarks from the Energy Information Administration for electricity generators, CARB's measure documentation for refineries and hydrogen plants, and sector studies for cement, oil and gas, cogeneration and other combustion. Each option entails a cost per ton of CO_2 in 2019 dollars, the share of the facility's CO_2 it can remove, and the NO_x it removes alongside. Hence, the co-emission ratio of Appendix B becomes an option-specific factor. Electricity generators also abate through dispatch, which we represent by the demand response estimated by Borenstein et al. (2019) at an elasticity of -0.2, scaled to the share of marginal load that in-state plants serve. We found no sourced cost for output reductions at refineries, oil and gas, hydrogen, cogeneration, or other combustion facilities. Those sectors are therefore bounded by their efficiency options, and Section 5.5 reports the cement output response as a sensitivity. Facilities already faced an allowance price of \$16.84 in 2019, so the cheapest measures on the menu were already in place. We therefore add \$16.84 to every option's cost, which sets the marginal abatement cost at 2019 emissions equal to the 2019 allowance price.[24] Finally, consistent with 2019 CARB compliance records, we treat CO_2-only

[23] Appendix F describes the frontier, its range across the cost cells reported for cement kilns, and the RECLAIM trades that bracket the multipliers. We do not scale the menu to RECLAIM credit prices, for the reasons given there.

[24] Allowances settled at a mean of \$16.84 per ton across the four 2019 auctions (\$15.73, \$17.45, \$17.16 and \$17.00; \$16.78 weighted by volume), with a reserve price of \$15.62 (California Air Resources Board, 2019). All prices in this section are in 2019 dollars.

post-combustion capture (c_i^g) as unavailable in the 2019 baseline and evaluate it at engineering retrofit costs in Appendix G.

The calibration requires several other design choices. First, the sourced abatement cost menu implies that the program's statutory reductions are not feasible. At any price up to $300 per ton, the engineering cost menu implies removal of about 12 percent of the facilities' 2019 CO_2, whereas the 2030 budget would require 42 percent. We therefore set the carbon side of the model by setting the price rather than specifying a specific cap. Our headline regimes hold the allowance price at the social cost of carbon, $44.28 per ton in 2019 dollars (the $41 of Holland et al. (2016) in 2014 dollars), which is the exogenous-price case of Corollary 1. At that price the menu implies the removal of 6.77 million tons, 7.3 percent of covered emissions, and we report the cap that clears at that price as a robustness check. Second, every facility has NO_x that no listed control can remove, so we apply the exemption of Proposition 9. The exempt residual is the share of each facility's 2019 NO_x that remains after the control that the South Coast district's BARCT rules require for its equipment, which we apply statewide. This represents 42 percent of the sample's NO_x. Because the exempt tons sit at the facilities nearest to population, this represents 73 percent of the median county's exposure. Third, we solve every regime by iterating prices on the discrete abatement cost menus, so that firms' responses are step functions. We assume that a firm indifferent at a step keeps its permits rather than selling them. The strategic regimes are solved as exact games on those steps. A community's best response is found by enumerating the breakpoints of its residual demand with the carbon market re-cleared at each. A profile of sub-sellers is an equilibrium only if no sub-seller gains from any unilateral deviation. Because the pricing game has several equilibria we report in the text the equilibrium reached from the price-taking profile and in Appendix G every equilibrium found from eleven starting profiles for the unitary game and five for each subdivided game. Fourth, our core results use counties to define communities. We compute the non-strategic regimes for census places and tracts as well, and we restrict the strategic regimes to counties because the exact game is not computable at the finer units. Table 2 summarizes the key model objects.

Using this calibrated economy, we compute the equilibrium of each policy regime in Table 1. These include the four traditional instruments (Regimes I–IV) and the community permit mechanism under price-taking and unitary strategic pricing ($n_j = 1$). To evaluate the within-community governance subdivisions from Section 4, we compute the estimated subdivision $\hat{n}_j$ set from noisy damage estimates and uniform subdivisions that assign the same $\bar{n}$ sub-sellers to every county ($\bar{n} \in [2, 1{,}000]$). We evaluate each regime's welfare loss relative to the constrained first best under both perfect information and damage-estimation error.

Table 2. Calibration Summary

	Tract	Place	County
Communities J	9,109	529	58
Facilities I	279	279	279
Median d_j (\$/μg/m^3)	1,533,368	18,116,791	79,725,011
Median MD_i (\$/ton NO_x)	14,093	14,094	14,093
Abatable share of 2019 NO_x	0.58	0.58	0.58
Exempt share (BARCT residual)	0.42	0.42	0.42
Allowance price τ^A (\$/tCO_2e), held	44.28	44.28	44.28
τ^A at price-taking under the fixed cap	–	–	42.82
Distinct verified equilibria, $\bar{n} = 3$	–	–	5

Notes: Discrete-menu calibration, 2019 dollars. $MD_i = \sum_j h_{ji} d_j$ is the same facility vector at every community definition up to aggregation, so its median does not change. The exempt share is the BARCT residual (the NO_x that remains after the required controls), which no community permit covers. The headline holds the allowance price at \$44.28; the robustness case fixes the cap that clears at \$44.28 in the status quo and lets the price adjust, and was solved at the county definition only. The last row counts the distinct verified equilibria of the sub-seller game at $\bar{n} = 3$ from 5 starts; the strategic layer is computed at the county definition only.

Notes: Median d_j is the median community damage per unit of ambient concentration; it rises with aggregation because larger communities combine more exposed population. Median MD_i, the abatable share of NO_x and the exempt share are source-level objects and do not vary with the community definition. The carbon price row reports the market-clearing τ^A under the price-taking mechanism at the cap that clears at \$44.28 without community permits. It lies below \$44.28 because community pricing induces carbon co-abatement (Proposition 3); under the exogenous-price framing of the headline results it equals \$44.28 by construction. The equilibria row reports the number of distinct verified equilibria of the sub-seller game at $\bar{n} = 3$, the smallest subdivision at which every starting profile reaches a verified equilibrium, computed for counties only.

5.2 Aggregate Welfare Results

Table 3 reports the annual welfare loss of each regime relative to the constrained first best, where zero loss corresponds to a \$66.5 million annual gain over the carbon-only status quo. The first best removes 6,522 tons of NO_x per year, whereas the carbon market alone removes 1,964 tons. That \$66.5 million gain equals one-seventh of the \$465 million in 2019 local damages because 42 percent of baseline tons are exempt post-BARCT emissions and the rest are abated only where controls cost less than avoided damages. Column 1 presents the full-information benchmark ($\hat{d}_j = d_j$). Under certainty, differentiated taxes and cap-and-trade (Regimes II and IV) and price-taking community permits all attain the constrained first best. The optimal uniform local price of \$14,528 per ton loses \$28.3 million per year by ignoring spatial variation in marginal damages. When each county prices as a unitary seller ($n_j = 1$), best-response iteration does not converge. Across eleven starting profiles, the end states after 200 rounds lose \$82 million to \$114 million per year relative to the first best. In the state closest to an equilibrium, where no county can gain more than \$1.2 million by deviating, the loss is \$113 million (\$46 million worse than having no local instrument). Counties in that state price at a median of four times marginal damage, and firms remove 15,100 tons of NO_x compared with 6,522 tons at the first best. Table G2 reports all eleven end states. The subdivided games admit multiple verified equilibria at each subdivision, and Table 3 reports the equilibrium reached from the price-taking start alongside the range across five starting profiles.

Even under full information, the estimated subdivision $\hat{n}_j$ loses \$11.3 million (\$8.6 million to \$13.2 million) because integer sub-seller counts and linear slope fits only approximate step-function demand (Figure G2). Under a uniform subdivision of $\bar{n} = 10$, the loss is \$15 million (\$15 million to \$22 million across equilibria), recovering three-quarters of the first-best gain without damage information. At $\bar{n} = 5$ the loss is \$15 million from the price-taking start (up to \$31 million across equilibria), at $\bar{n} = 3$ it is \$30 million (\$19 million to \$45 million), and at $\bar{n} \geq 50$ it exceeds the uniform-tax loss. At $\bar{n} = 3$ to 10, the median county prices at two to ten times d_j as low-damage rural counties withhold permits, whereas beyond $\bar{n} = 20$ high-damage urban counties price below d_j and give permits away across steps.

Table 3. Welfare Losses by regime

Welfare loss relative to the price-taking first best (\$M per year; larger = worse)			
Regime	Perfect information	Estimation error (lognormal, CV = 1)	Verified equilibria / starts
Uniform tax or cap-and-trade (I/III)	28.31	38.18	–
Differentiated tax or cap-and-trade (II/IV)	0	11.81	–
Community permits, price-taking (V)	0	0	–
Community permits, unitary, n_j=1 (V); approximate equilibrium, residual gain \$1.2M	112.76 (81.50–113.63)	112.76 (81.50–113.63)	0 / 11
Community permits, subdivision at $\hat{n}_j$ (V)	11.30 (8.58–13.21)	7.88 (6.61–11.73)	5 / 5
Community permits, uniform $\bar{n}$=10 (V)	15.24 (15.16–21.78)	15.24 (15.16–21.78)	5 / 5
Community permits, uniform $\bar{n}$=3 (V)	30.48 (18.97–45.36)	30.48 (18.97–45.36)	5 / 5
Community permits, uniform $\bar{n}$=5 (V)	15.33 (15.33–31.04)	15.33 (15.33–31.04)	5 / 5

Notes: Communities are counties ($J = 58$); 2019 dollars. Allowance price held at \$44.28/tCO$_2$e. Every entry is a loss relative to the price-taking community-permit equilibrium, whose gain over the carbon-only status quo is \$66.46M per year. For the subdivided regimes the entry is the loss at the equilibrium reached from the price-taking start, with the range across the distinct verified equilibria beneath, and the last column counts verified equilibria over the starts (the price-taking start plus random starts). No exact equilibrium of the unitary game was found from any start; the row reports the approximate equilibrium, the end state after 200 rounds with the smallest remaining best-response gain (the price-taking start), with the range across all end states beneath. The appendix table reports every end state. The estimation-error column draws $\hat{d}_j$ lognormal around d_j with sd = d_j from the rerun of 2026-11-02 at the held allowance price (200 lognormal draws); the two taxes are re-solved at the estimates, the $\hat{n}_j$ row re-solves the sub-seller game at $\hat{n}_j(\hat{d})$ (carried from the fixed-cap band run of 2026-10-29 through 11-01; the price-taking-start equilibrium, with the best and worst equilibria beneath, averaged over 13 draws with a verified equilibrium), and the regimes that use no estimate repeat the first column.

Column 2 introduces damage-estimation error. We assume that the regulator observes each county's damage times a lognormal error with mean one and a standard deviation equal to the true damage (CV = 1) and sets its instrument from those estimates. We evaluate the resulting welfare outcome at true damages. At this noise level the expected loss under differentiated taxes or cap-and-trade is \$11.8 million per year, vs. \$38.2 million under undifferentiated regulation. The reversal of Proposition 6 does not occur at any error level we compute, for the reason given in Remark 1. The price-taking and uniform-subdivision outcomes do not change across the two

columns because neither uses the regulator's estimates. When the regulator uses its noisy estimates $\hat{\mathbf{d}}$ to set the within-community subdivision $\hat{n}_j$, the expected loss at CV $= 1$ is \$8 million from the price-taking start (\$7 million to \$12 million across the equilibria we find). At every equilibrium it is below the differentiated instruments' \$11.8 million and well below the uniform tax.[25]

The uniform subdivisions in the lower rows of Table 3 use no damage information at all, so their losses are deterministic given the equilibrium that obtains. At $\bar{n} = 10$ the loss is \$15 million per year, which beats the uniform tax at every error level and the differentiated instruments beyond a coefficient of variation of 1.3 (1.3 to 2.0 across its equilibria). Section 5.5 reports the sensitivity of the best $\bar{n}$ to the cost of NO_x abatement.

5.3 Distributional Implications

Figure 2 shows the welfare incidence of each regime for individual counties, measured as avoided damages plus permit revenue according to the spatially resolved damage accounting of Holland et al. (2016, 2019). All gains are measured relative to the status quo, where there are no local co-pollutant adjustments.[26] We also plot the regulated facilities for reference and indicate disadvantaged communities based on California's SB 535 (California Office of Environmental Health Hazard Assessment, 2021) with bold outlines.

Under the traditional instruments (Regimes I–IV), community welfare gains consist solely of avoided pollution damages, which total \$122 million per year statewide at the first best. Under the community permit mechanism, communities also receive the revenue from selling exposure permits, amounting to an additional \$75 million per year under price-taking. Because permit payments are a transfer from regulated firms to communities, these community-side figures are larger than the net social welfare gains reported in Table 3. Across all panels of Figure 2, community gains are concentrated in the South Coast Air Basin, the San Francisco Bay Area, and the southern San Joaquin Valley. Both covered industrial facilities and downwind populations are dense in those three regions.

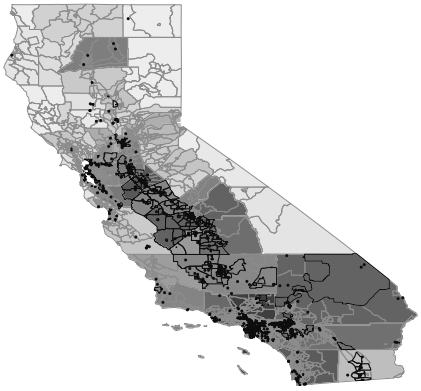
Figure 3 disaggregates community welfare gains between the 29 counties containing disadvantaged communities (DACs) and the 29 non-DAC counties, reporting mean gains per county and separating avoided damages from permit revenue. The scarcity rents that standard cap-and-trade or tax designs allocate to firms or the state accrue to exposed communities instead (Fullerton and Metcalf, 2001; Fullerton, 2011). At the first best, DAC counties receive \$120 million of the \$122 million in statewide avoided damages (\$4.14 million per county vs. \$0.08 million in non-DAC

[25] The rule sets $\hat{n}_j = 1 + c_j(\hat{d}_j)/(B_j\hat{d}_j)$, where B_j is the fitted slope of the county's residual demand over prices up to $2\hat{d}_j$. On linear demand $c_j(0) = c_j(\hat{d}_j) + B_j\hat{d}_j$, so this rule matches Equation (18). In the 21 counties where the first best abates more than 0.5 percent of residual demand, $\hat{n}_j$ at true damages ranges from 1 to 77 with a median of 6 (for example, 1 in Kern and 2 in Los Angeles). In the remaining rural and upwind counties, no control on the menu costs less than $2d_j$. Step demand is therefore flat near d_j and $\hat{n}_j$ is large, giving a statewide median of about 12,000 and 16 counties at the one-million ceiling. In those counties, first-best abatement is zero and a large $\hat{n}_j$ simply pushes the permit price to zero. Varying the ceiling between 100,000 and ten million changes statewide welfare by \$0.5 million.

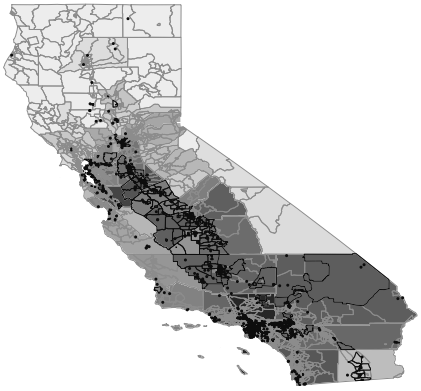
[26] Even the undifferentiated regimes produce gains relative to this baseline because they incorporate information about damages from local co-pollutants, which the status-quo regime does not.

counties). They also receive $72 million of the $75 million in permit revenue ($2.50 million per county vs. $0.10 million) because covered facilities and exposed populations sit in the same counties. Under unitary pricing, the approximate equilibrium lowers statewide permit revenue to $36 million because firms over-abate at prices four times marginal damage. Under the uniform subdivisions, DAC counties continue to account for nearly all avoided damages, but low-damage non-DAC counties still withhold permits. At $\bar{n} = 5$ (shown in Figure 3), non-DAC counties price at a median of 16 times marginal damage and collect $23 million in revenue ($0.78 million per county) on $4.5 million of avoided damages. DAC counties collect $135 million in revenue ($4.66 million per county) on $121 million of avoided damages ($4.16 million per county). At $\bar{n} = 3$, non-DAC counties price at 30 times marginal damage and collect $28 million in revenue on $5 million of avoided damages, while DAC counties collect $108 million in revenue on $143 million of avoided damages. Under the estimated subdivision $\hat{n}_j$, non-DAC counties' median price falls to half of marginal damage and their statewide revenue drops to $3.5 million ($0.12 million per county). DAC counties receive $123 million in avoided damages ($4.26 million per county) and $90 million in revenue ($3.10 million per county), close to the first best.

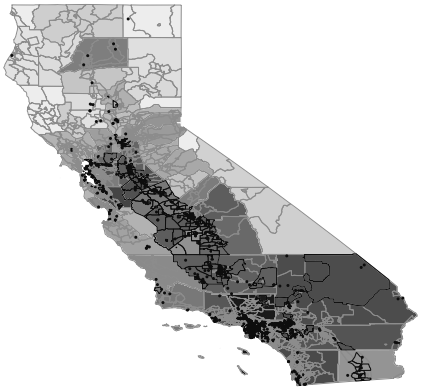
Figure 2. Community Welfare Incidence by County



(a) Uniform tax or cap-and-trade (I/III)



(b) Differentiated tax or cap-and-trade (II/IV)



(c) Community permits, price-taking (V)

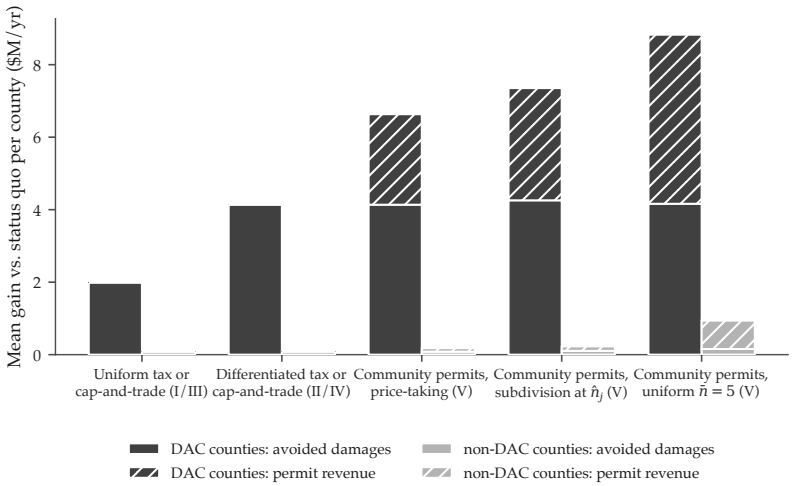


(d) Community permits, subdivision at $\hat{n}_j$ (V)



Notes: Each panel shows county-level welfare incidence relative to the status quo. Darker shading indicates larger gains, hatching indicates losses. Shading is log-scaled with a $10,000 floor. Dots mark the 279 regulated facilities. Thin lines are census tract boundaries, and bold outlines mark the tracts designated as disadvantaged communities (DACs) under SB 535, drawn on the CalEnviroScreen 4.0 tract geography.

Figure 3. Incidence in DAC vs. non-DAC Counties



Notes: Bars report the mean welfare gain per county relative to the carbon-only status quo, in $M per county per year, averaged separately over DAC and all other counties. Solid segments are reductions in local air-pollution damages; hatched segments are community permit revenue (V regimes only). A county is DAC if it contains at least one census tract designated as disadvantaged under SB 535 (CalEnviroScreen 4.0); 29 of 58 counties qualify.

5.4 Damage Uncertainty and Community Scale

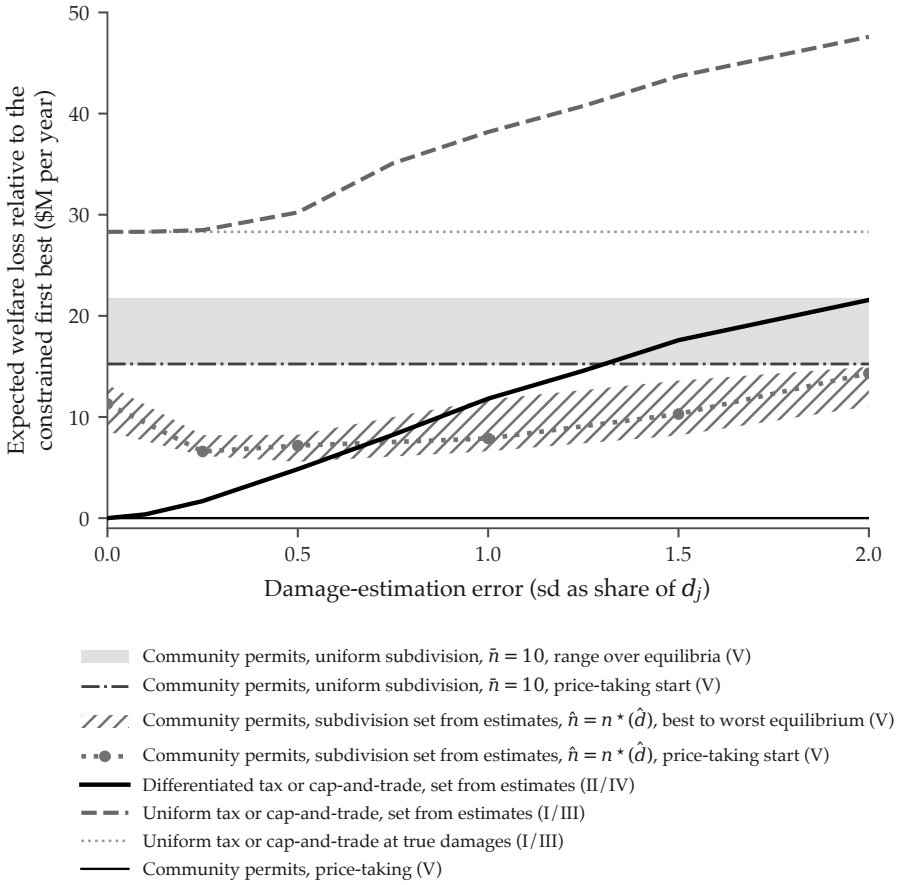
The comparisons in Table 3 hold damage-estimation error fixed at CV $= 1$ and define communities as counties. Because Propositions 6 and 8 show that relative instrument performance depends on both the magnitude of estimation error and the geographic scale of communities, we now vary each dimension in turn.

Figure 4 plots expected welfare losses at the county scale as the coefficient of variation (CV) of independent lognormal damage-estimation errors increases from 0 to 2, with 200 draws at each point. The regimes that do not rely on regulator damage estimates are flat. Price-taking community permits lose nothing at every error level, and the uniform subdivision at $\bar{n} = 10$ loses $15 million to $22 million across equilibria (shaded band).

Among the regimes that use the regulator's damage estimates, the differentiated taxes and cap-and-trade lose $4.8 million at CV $= 0.5$, $11.8 million at CV $= 1$ and $21.6 million at CV $= 2$. They also stay below the uniform tax at every error level we compute, which is the opposite of the ranking the closed form of Proposition 6 gives beyond CV*. The differentiated loss grows roughly linearly in the error rather than with its square. Two features of the abatement cost menus explain why (as Remark 1 anticipates). First, price errors have no impact until they trigger a facility's next control. Second, the 91 facilities with no control beyond their 2019 equipment cannot respond to any error. The uniform tax, by contrast, loses the protection of averaging because the optimal uniform price lies at a breakpoint of the aggregate menu. Shifts in beliefs move it across breakpoints, raising its expected loss from $28.3 million to $38.2 million at CV $= 1$ and $47.6 million at CV $= 2$. When the regulator instead uses $\hat{\mathbf{d}}$ to set the within-community subdivision $\hat{n}_j$, communities' equilibrium pricing dampens the estimation error (Proposition 8), and the mecha-

nism's expected loss stays between \$6 million and \$15 million across error levels and equilibria. The differentiated instruments therefore outperform the estimated subdivision at small errors but do worse once the coefficient of variation exceeds 0.6 to 1.0. The uniform subdivision at $\bar{n} = 10$ beats the differentiated instruments beyond a coefficient of variation of 1.3 to 2.0, depending on the equilibrium.

Figure 4. Instrument Performance as a Function of Estimation Error



Notes: Each curve plots the expected welfare loss relative to the constrained first best (\$M per year, county unit, allowance price \$44.28) when the regulator acts on estimates $\hat{d}_j = d_j \times \eta_j$, with independent lognormal errors of mean one whose standard deviation equals the horizontal-axis multiple of d_j, over 200 draws per point. The uniform and differentiated taxes are set from the estimates and evaluated at true damages. The $\hat{n}_j$ band spans the verified equilibria of the subdivided game at each error level, with the equilibrium reached from the price-taking start as a line.

Community scale matters for the taxes as well. The first best and the full-information uniform price do not depend on the community definition, because both are set from the same facility-level marginal damages, but the two taxes under estimation error do. Figure G6 in Appendix G holds estimation error at CV $= 1.0$ and compares expected losses for counties ($J = 58$), incorporated Census places ($J = 529$), and census tracts ($J = 9{,}109$, the unit at which CalEnviroScreen designates disadvantaged communities). Under undifferentiated taxes or cap-and-trade (Regimes I and III), the expected loss at CV $= 1$ falls from \$38.2 million per year at the county scale to \$28.4 mil-

lion at places and $28.3 million at tracts. Partitioning receptors more finely averages the uniform price over more independent errors and keeps it from jumping across menu breakpoints.

Under regulator-differentiated taxes or cap-and-trade (Regimes II and IV), expected losses under independent errors fall from $11.8 million per year at the county scale to $2.6 million at the place scale and $0.2 million at the tract scale. That near-zero tract loss, however, is an artifact of assuming that estimation errors are statistically independent across neighboring census tracts. When errors are independent, partitioning a facility's plume into 9,109 tracts averages out idiosyncratic draw noise in $\widehat{\mathrm{MD}}_i = \sum_j \hat{d}_j h_{ji}$ and shrinks the population weight on any single mistaken draw. In practice, top-down tract damage estimates multiply census populations by a shared value of a statistical life, a shared epidemiological concentration-response slope, and a shared atmospheric chemistry model. Errors in those common inputs are positively correlated across tracts, while neighborhood-specific variation in willingness to pay or baseline morbidity goes unobserved by a uniform VSL. Because a facility's true marginal damage $\mathrm{MD}_i = \sum_j d_j h_{ji}$ does not depend on how receptors are partitioned, a common multiplicative error scales $\widehat{\mathrm{MD}}_i$ identically at the tract and county scales. When we re-run the scale comparison at $\mathrm{CV} = 1$ with half of the error's log-variance common across communities ($\rho = 0.5$), the differentiated tax's expected loss is $13.8 million at counties, $9.5 million at places, and $9.0 million at tracts. When four-fifths of the log-variance is common ($\rho = 0.8$), the loss is $14.1 million at counties, $13.0 million at places, and $12.7 million at tracts. Under a pure common scaling error ($\rho = 1$), the loss is $15.1 million at all three scales.[27] At the county scale, the estimated subdivision $\hat{n}_j$ loses $8 million from the price-taking start ($7 million to $12 million across equilibria), and the uniform subdivision $\bar{n} = 10$ loses $15 million without using damage estimates. County-level community permits therefore remain competitive with or outperform even a tract-level differentiated tax once common valuation errors are present.

We do not compute the strategic regimes at the place or tract scale, because the exact game on the engineering abatement cost menus is not computable at 529 or 9,109 communities (Appendix H). Decentralized community permit markets are therefore best suited to general-purpose jurisdictions such as counties or incorporated municipalities. Those jurisdictions already possess administrative capacity and can be divided among a modest number of local sub-sellers, such as council districts, school boards, or neighborhood associations. Top-down differentiated taxes, by contrast, gain from defining receptors finely only to the extent that estimation errors are idiosyncratic across neighborhoods rather than driven by common epidemiological or valuation uncertainty.

As in any empirical evaluation of spatially differentiated regulation, our baseline damages d_j are researcher estimates that serve both as the true damage benchmark and as the mean around which regulator errors are drawn. Scaling d_j up or down changes the dollar welfare losses of estimate-dependent policies, which we examine by varying the value of a statistical life in Sec-

[27] Under the cap framing of Table G3, the county-level differentiated tax loss at $\mathrm{CV} = 1$ rises from $11.8 million at $\rho = 0$ to $13.8 million at $\rho = 0.5$ and $14.1 million at $\rho = 0.8$. The county-level uniform tax loss rises from $37.8 million to $38.7 million and $40.5 million.

tion 5.5. The price-taking and uniform-subdivision equilibria do not depend on the regulator's beliefs.

5.5 Robustness and Discussion

Appendix G reports sensitivity to the tuning of the carbon-side menu and to the allowance price (Table G4). Cross-regime welfare differences come from local NO_x pricing rather than from the carbon side.

We next examine sensitivity to the cost of NO_x-only abatement. In Table 4, we apply a multiplier to every control's cost. The multipliers we use are 1.85 and 2.39, the two ratios by which the costs Fowlie and Muller (2019) recovered from observed compliance in the NO_x Budget Program exceeded their engineering menu. We also separately replace the midpoint cost estimate for cement-kiln controls with the low and high estimates in the South Coast staff reports. As NO_x controls become more expensive across the first three columns, the first-best gain from community permits falls from \$66.5 million per year at engineering costs to \$40.5 million and \$32.0 million under the two multipliers. Fewer controls cost less than the damage they avoid. The uniform-price loss also falls from \$28.3 million to \$11.6 million and \$7.4 million, because fewer facilities are mispriced across a step. The loss under the estimated subdivision moves the other way, from \$11.3 million to \$28.2 million and \$40.0 million, and the loss at $\bar{n} = 3$ from \$30.5 million to \$60.7 million and \$91.0 million. Steeper menus leave each county's residual demand with fewer and larger steps, so withholding across a step costs more. This means that the market-power losses that subdivision must correct grow as the first-best gain shrinks. At the revealed costs, then, the estimated subdivision no longer beats the differentiated instruments at $\mathrm{CV} = 1$, and the mechanism beats them only when the regulator has no usable estimates. Across the low and high cement-kiln cost cells, the first-best gain ranges from \$63.7 million to \$72.0 million and the uniform-tax loss from \$19.2 million to \$36.6 million. Cement kilns are among the largest NO_x sources, and their single control straddles the median marginal damage. The loss under the estimated subdivision moves by less than \$2 million and the loss at $\bar{n} = 3$ from the price-taking start spans \$18 million to \$31 million.

Table 5 examines sensitivity to the valuation of local pollution damages. Across columns, we vary the value of a statistical life from our \$10.3 million headline to the 2.5th and 97.5th percentiles of the distribution used in AP4, which scale every community's marginal damage d_j by 0.09 and 2.5. The last column applies the steeper concentration-response slope of Di et al. (2017) to the population aged 65 and over. At the low value almost no control on the frontier costs less than the damage it avoids, so the first-best gain over the carbon-only status quo is \$0.6 million per year and the local layer has little to add. At the high value the first-best gain is \$311 million and the best uniform price rises from \$14,528 to \$64,032 per ton with a loss of \$111 million. This loss is driven by the increase in the dispersion of marginal damages across facilities, which is what a uniform price ignores. The loss under the estimated subdivision changes little, from \$11.3 million to \$9.4 million and \$8.3 million, because higher damages bring each county's efficient price closer to the

Table 4. Sensitivity to the Cost of NO_x Controls

	Multiplier on every control's cost			Cement kiln cost cell		
	Engineering costs	×1.85	×2.39	low	mid (headline)	high
$M per year unless noted						
Price-taking gain vs. status quo	66.46	40.52	31.96	71.98	66.46	63.66
Uniform-tax loss vs. price-taking	28.31	11.63	7.40	36.56	28.31	19.16
Uniform tax t^* ($/ton NO_x)	14,528	14,025	17,780	6,999	14,528	18,226
$\hat{n}_j$ loss, price-taking start	11.30	28.16	40.03	13.25	11.30	11.42
NO_x removed at price-taking (t)	6,522	4,764	4,292	6,522	6,522	5,768
CO_2 abated at price-taking (t)	6,800,184	6,800,843	6,800,843	6,800,184	6,800,184	6,800,184

Notes: The first three columns multiply the cost of every NO_x control by one (the engineering estimates, the headline), by 1.845 and by 2.393. The multipliers are the ratios by which the costs Fowlie and Muller (2019) recovered from observed compliance in the NO_x Budget Program exceeded their engineering menu, at the optimal cap (1.845) and at the observed cap (2.393). The last three columns keep the engineering costs for every other control and set the cement-kiln controls at the low, midpoint and high cost estimates in the South Coast staff reports; the midpoint is the headline, so the engineering-cost column and the midpoint column hold the same numbers. Every other input is at its headline value, with the allowance price held at $44.28. Rows are defined as in the carbon-side menu table. County communities, 2019 dollars.

top of its residual demand (leaving less room for withholding). This is the mechanism behind Proposition 7. The Di slope results in a first-best gain of $54.0 million, a uniform loss of $22.5 million and a subdivision loss of $6.9 million.

Table 5. Sensitivity to the Valuation of Local Damages

	Value of a statistical life ($M)			Mortality slope
	0.94	10.29	25.59	Di 65+
$M per year unless noted				
First-best (price-taking) gain vs. status quo	0.62	66.46	311.22	53.97
Uniform-tax loss vs. price-taking	0.11	28.31	111.12	22.49
Uniform tax t^* ($/ton NO_x)	1,964	14,528	64,032	14,474
$\hat{n}_j$ loss, price-taking start	9.36	11.30	8.30	6.92
Baseline NO_x damages at status quo	39.51	430.52	1,070.90	384.67

Notes: County communities, allowance price held at $44.28, headline cost reading, 2019 dollars. The first three columns rebuild d_j with the AP4 concentration–response slope at the 2.5th percentile, the central value and the 97.5th percentile of the value of a statistical life; the central column is the headline. The last column keeps the headline VSL and replaces the slope by the Di et al. (2017) estimate for ages 65 and over. The $\hat{n}_j$ row is the loss at the equilibrium reached from the price-taking start.

Figure G5 in Appendix G evaluates a counterfactual in which post-combustion carbon capture and sequestration (CCS) becomes economically viable. Because no covered California facility operated CCS in 2019, all carbon abatement in our baseline occurs through the joint CO_2–NO_x abatement channel (efficiency, fuel and dispatch measures), which reduces local NO_x emissions as a byproduct. In Figure G5, we introduce the CO_2-only capture margin (c_i^g) and vary its cost from four times the sector-specific engineering estimates of National Energy Technology Laboratory (2022) down to one-quarter of those estimates. Each facility's capture option removes 90 percent of its remaining CO_2. Cheaper capture does not displace the measures that produce the co-benefit. Capture is adopted at the 42 oil and gas facilities once its cost reaches the engineering

estimate, and at every facility if costs fall to one quarter of the engineering estimate. The efficiency and fuel measures cost less than capture and are adopted regardless. With the allowance price held fixed, incidental NO_x co-abatement is therefore unchanged at every capture cost. Under a fixed cap it falls by 2 percent, because cheaper capture lowers the clearing price to \$21. The efficiency gain from adding the community permit layer moves by less than \$1 million. As established in Lemma 2(c), lower carbon-capture costs weaken the cost complementarity $|\gamma_i|$ between CO_2 and NO_x abatement in the quadratic economy, where firms substitute smoothly between margins. On the discrete menus that substitution does not occur. The environmental-justice concern that capture would strip carbon policy of its local co-benefits (Batres et al., 2021) does not arise at these costs, and the value of a dedicated local permit layer is unchanged.

To put our welfare magnitudes in perspective, it is useful to compare them to the scale of the local externality rather than the aggregate carbon market. The carbon abatement that clears at the social cost of carbon is worth roughly \$300 million per year at that price. By comparison, the 279 covered facilities generate \$465 million per year in baseline local NO_x damages (\$431 million under the carbon-only market). The \$66 million efficiency gain under competitive community pricing is a large share of the achievable local surplus, and so are the \$55 million net gain under the estimated subdivision and the \$51 million under ten uniform sub-sellers. Price-taking community receipts of \$197 million per year, \$75 million in permit revenue plus \$122 million in avoided damages, are two-fifths of the baseline damage burden.[28] Moreover, because efficiency gains from spatial differentiation depend on the spatial dispersion of marginal damages across sources, they are smaller within a single state than across the entire United States (Muller and Mendelsohn, 2009; Fowlie and Muller, 2019; Hollingsworth et al., 2026). Incorporating additional co-pollutants beyond NO_x or morbidity endpoints beyond adult mortality would further increase the welfare gains reported here.

6 Conclusion

This paper proposes and evaluates a stacked environmental market design in which the communities that bear localized damages own and price the permits governing their exposure. Firms covered by an aggregate emissions cap must also hold community-specific ambient permits in proportion to their local impact, and communities receive slack permit endowments that they sell to firms at posted prices. Rather than attempting to elicit private damage valuations through a centralized mechanism (Myerson and Satterthwaite, 1983), assigning permit endowments to communities lets local prices reflect marginal damages directly. While the competitive equilibrium of the stacked market attains the constrained first best without requiring the regulator to estimate local damages, vesting exclusive selling authority in individual communities creates local

[28] This contrast between net efficiency gains and permit revenues follows the standard distinction between second-order Harberger deadweight-loss triangles and first-order revenue rectangles (Harberger, 1964; Hines, 1999). Consequently, the distributional incidence of community permit endowments can be of first-order political and economic importance even when net deadweight-loss reductions are modest.

monopoly power and an anticommons externality across communities (Heller, 1998; Buchanan and Yoon, 2000).

Dividing each community's selling authority among independent local sub-sellers introduces a within-community common-pool externality that counteracts these between-community markups. The exact subdivision $n_j^\star$ that restores the first best serves as a theoretical benchmark because it depends on true damages. It motivates two practical policy designs. An estimated subdivision ($\hat{n}_j$) uses noisy damage estimates to set local governance rather than prices, and a uniform subdivision ($\bar{n}$) assigns the same number of sub-sellers to every community without using damage estimates.

Our simulation on California's cap-and-trade program illustrates the quantitative stakes of this tradeoff. Covered facilities generate $465 million per year in local NO_x mortality damages at 2019 emissions ($431 million under the carbon-only market). Relative to the carbon-only market, competitive community pricing gains $66 million per year and removes three times as much NO_x. When each county prices as a single seller, the state closest to an equilibrium loses $113 million per year relative to the first best ($46 million worse than the carbon market alone). Both practical subdivisions restore equilibria and remove most of that loss. Under baseline damage-estimation error ($\text{CV} = 1$), the estimated subdivision $\hat{n}_j$ loses $8 million per year relative to the first best from the price-taking start ($7 million to $12 million across equilibria). By comparison, differentiated taxes or cap-and-trade lose $11.8 million and a uniform tax loses $38.2 million. The estimated subdivision is the better instrument once the estimates' coefficient of variation exceeds 0.6 to 1.0, depending on the equilibrium. A uniform subdivision of ten sub-sellers per county loses $15 million but does not require any damage estimate. On the distributional margin, price-taking concentrates welfare gains in the counties that contain disadvantaged communities, which receive almost all of the $122 million in avoided damages and the $75 million in permit revenue (Fullerton and Metcalf, 2001; Fullerton, 2011). Because communities actively set the terms of local exposure rather than passively receiving ex post compensation, the mechanism also addresses procedural equity concerns without imposing trade restrictions that erode market efficiency (Schlosberg, 2007; Kroetz et al., 2015; Banzhaf et al., 2019).

Implementing a community permit system in practice requires pairing community endowments with an appropriate geographic and governance scale. Partial precedents exist in rights-based fisheries management, such as Alaska's Community Development Quota (CDQ) program, which allocates tradable Bering Sea harvest quotas to 65 western Alaska villages that lease quota for royalties (Ginter, 1995; Haynie, 2014; Sutherland and Edwards, 2022). In our pollution setting, defining receptors at the census-tract level reduces the differentiated tax's loss only when estimation errors are independent across neighboring tracts. When errors instead share a common component from the VSL, the concentration-response slope, or regional atmospheric modeling, fine geographic partitioning cannot average out the common bias. The tract-level differentiated tax then still loses $9 million to $13 million per year at our baseline error level. At the county scale, communities price on their own local damages, and dividing selling authority among about ten

local bodies (such as council districts, school boards, or neighborhood associations) curbs market power once unabatable tons are exempted. Whether finer jurisdictions sharpen market power faster than they diversify error is a question our county-level game does not answer. We emphasize three caveats to our empirical analysis. First, our calibration focuses on a single local co-pollutant (NO_x-derived $PM_{2.5}$ mortality). Second, our welfare calculations are conducted in partial equilibrium, leaving general-equilibrium interactions with local housing and labor markets for future work. Third, we simulate compliance by cost minimization on engineering control menus, which is the ex ante approach of Fowlie and Muller (2019). They show that an econometric model of observed compliance choices predicts behavior better than cost minimization does, and a version of our simulations built on such a model is in progress.

Finally, while we have focused on local co-pollutants under a carbon cap, the stacked community permit design applies more broadly whenever an aggregate cap-and-trade market generates spatially heterogeneous local externalities. The essential ingredients are a measurable mapping $\mathbf{H}$ from source activities to local receptor impacts and identifiable local jurisdictions capable of managing permit sales. In water-quality trading, watershed-wide nutrient caps coexist with localized damages at downstream drinking-water intakes and estuaries linked by hydrologic routing models (Hung and Shaw, 2005). In groundwater markets, basin-wide extraction caps coexist with localized well-drawdown externalities governed by aquifer response functions (Ayres et al., 2021). Wetland and habitat offset markets pose a greater measurement challenge because site-specific ecosystem services are harder to meter than physical concentrations (Salzman and Ruhl, 2000; Aronoff and Rafey, 2026). Where credible transfer mappings exist, however, endowing affected communities with local exposure permits offers a general way to incorporate local damage information into environmental markets.

References

Agee, M. D., S. E. Atkinson, T. D. Crocker, and J. W. Williams (2014). Non-separable pollution control: Implications for a CO_2 emissions cap and trade system. *Resource and Energy Economics 36*(1), 64–82.

Aronoff, D. and W. Rafey (2026). Conservation priorities and environmental offsets: Markets for Florida wetlands. *American Economic Review 116*(5). forthcoming/May 2026; pages TBD.

Atkinson, S. E. and T. H. Tietenberg (1982). The empirical properties of two classes of designs for transferable discharge permit markets. *Journal of Environmental Economics and Management 9*(1), 101–121.

Ayres, A. B., K. C. Meng, and A. J. Plantinga (2021). Do environmental markets improve on open access? evidence from california groundwater rights. *Journal of Political Economy 129*(10), 2817–2860.

Banzhaf, S., L. Ma, and C. Timmins (2019). Environmental justice: The economics of race, place, and pollution. *Journal of Economic Perspectives 33*(1), 185–208.

Batres, M., F. Wang, H. Buck, R. Kapila, U. Kosar, R. Licker, D. Nagabhushan, E. Rekhelman, and V. Suarez (2021). Environmental and climate justice and technological carbon removal. *The Electricity Journal 34*(7), 107002.

Baumol, W. J. and W. E. Oates (1971). The use of standards and prices for protection of the environment. *Swedish Journal of Economics 73*(1), 42–54.

Bonilla, J., J. Coria, and T. Sterner (2018). Technical synergies and trade-offs between abatement of global and local air pollution. *Environmental and Resource Economics 70*(1), 191–221.

Borenstein, S., J. Bushnell, F. A. Wolak, and M. Zaragoza-Watkins (2019). Expecting the unexpected: Emissions uncertainty and environmental market design. *American Economic Review 109*(11), 3953–3977.

Buchanan, J. M. and Y. J. Yoon (2000). Symmetric tragedies: Commons and anticommons. *Journal of Law and Economics 43*(1), 1–14.

Burtraw, D., A. Krupnick, K. Palmer, A. Paul, M. Toman, and C. Bloyd (2003). Ancillary benefits of reduced air pollution in the US from moderate greenhouse gas mitigation policies in the electricity sector. *Journal of Environmental Economics and Management 45*(3), 650–673.

Burtraw, D. and N. Roy (2023). How would facility-specific emissions caps affect the california carbon market? Technical report, Resources for the Future.

California Air Resources Board (2019). Summary of california–quebec joint auction settlement prices and results. Technical report, CARB.

California Office of Environmental Health Hazard Assessment (2021). Calenviroscreen 4.0. Technical report, CalEPA OEHHA.

Chong, H. and D. Sunding (2006). Water markets and trading. *Annual Review of Environment and Resources 31*, 239–264.

Coase, R. H. (1960). The problem of social cost. *Journal of Law and Economics 3*, 1–44.

Considine, T. J. and D. F. Larson (2006). The environment as a factor of production. *Journal of Environmental Economics and Management 52*(3), 645–662.

Costello, C. and D. Kaffine (2018). Natural resource federalism: Preferences versus connectivity for patchy resources. *Environmental and Resource Economics 71*(1), 99–126.

Crocker, T. D. (1966). The structuring of atmospheric pollution control systems. In H. Wolozin (Ed.), *The Economics of Air Pollution*, pp. 61–86. New York: W. W. Norton.

Crocker, T. D. (1981). Control of acid precipitation: Nonconvexities and irreversibilities. *Journal of the Air Pollution Control Association 31*(1), 31–37.

Dales, J. H. (1968). *Pollution, Property and Prices: An Essay in Policy-making and Economics*. Toronto: University of Toronto Press.

Dasgupta, P., P. Hammond, and E. Maskin (1980). On imperfect information and optimal pollution control. *Review of Economic Studies 47*(5), 857–860.

Di, Q., Y. Wang, A. Zanobetti, Y. Wang, P. Koutrakis, C. Choirat, F. Dominici, and J. D. Schwartz (2017). Air pollution and mortality in the Medicare population. *New England Journal of Medicine 376*(26), 2513–2522.

Diamond, P. A. (1973). Consumption externalities and imperfect corrective pricing. *Bell Journal of Economics and Management Science 4*(2), 526–538.

Duggan, J. and J. Roberts (2002). Implementing the efficient allocation of pollution. *American Economic Review 92*(4), 1070–1078.

Färe, R., S. Grosskopf, D.-W. Noh, and W. Weber (2005). Characteristics of a polluting technology: Theory and practice. *Journal of Econometrics 126*(2), 469–492.

Fowlie, M. (2010). Emissions trading, electricity restructuring, and investment in pollution abatement. *American Economic Review 100*(3), 837–869.

Fowlie, M., S. P. Holland, and E. T. Mansur (2012). What do emissions markets deliver and to whom? evidence from southern california's NOx trading program. *American Economic Review 102*(2), 965–993.

Fowlie, M. and N. Muller (2019). Market-based emissions regulation when damages vary across sources: What are the gains from differentiation? *Journal of the Association of Environmental and Resource Economists 6*(3), 593–632.

Fowlie, M., R. Walker, and D. Wooley (2020). Climate policy, environmental justice, and local air pollution. Technical report, Brookings Economic Studies.

Freeman, A. Myrick, I., J. A. Herriges, and C. L. Kling (2014). *The Measurement of Environmental and Resource Values: Theory and Methods* (3rd ed.). Washington, DC: RFF Press.

Fullerton, D. (2011). Six distributional effects of environmental policy. *Risk Analysis 31*(6), 923–929.

Fullerton, D. and D. H. Karney (2018). Multiple pollutants, co-benefits, and suboptimal environmental policies. *Journal of Environmental Economics and Management 87*, 52–71.

Fullerton, D. and G. E. Metcalf (2001). Environmental controls, scarcity rents, and pre-existing distortions. *Journal of Public Economics 80*(2), 249–267.

Fullerton, D. and G. E. Metcalf (2002). Cap and trade policies in the presence of monopoly and distortionary taxation. *Resource and Energy Economics 24*(4), 327–347.

Ginter, J. J. C. (1995). The Alaska community development quota fisheries management program. *Ocean & Coastal Management 28*(1–3), 147–163.

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

Hayek, F. A. (1945, sep). The use of knowledge in society. *The American Economic Review 35*(4), 519–530.

Haynie, A. C. (2014). Changing usage and value in the Western Alaska community development quota (CDQ) program. *Fisheries Science 80*(6), 1089–1099.

Heller, M. A. (1998). The tragedy of the anticommons: Property in the transition from Marx to markets. *Harvard Law Review 111*(3), 621–688.

Hernandez-Cortes, D. and K. C. Meng (2023). Do environmental markets cause environmental injustice? evidence from california's carbon market. *Journal of Public Economics 217*, 104786.

Hines, James R., J. (1999). Three sides of Harberger triangles. *Journal of Economic Perspectives 13*(2), 167–188.

Holland, S. P., E. T. Mansur, N. Z. Muller, and A. J. Yates (2016). Are there environmental benefits from driving electric vehicles? The importance of local factors. *American Economic Review 106*(12), 3700–3729.

Holland, S. P., E. T. Mansur, N. Z. Muller, and A. J. Yates (2019). Distributional effects of air pollution from electric vehicle adoption. *Journal of the Association of Environmental and Resource Economists 6*(S1), S65–S94.

Holland, S. P. and A. J. Yates (2015). Optimal trading ratios for pollution permit markets. *Journal of Public Economics 125*, 16–27.

Hollingsworth, A., T. Jaworski, C. Kitchens, and I. Rudik (2026). Economic geography and air pollution regulation in the United States. *Journal of Political Economy Microeconomics*. NBER Working Paper 29845.

Hsiang, S., P. Oliva, and R. Walker (2019). The distribution of environmental damages. *Review of Environmental Economics and Policy 13*(1), 83–103.

Hung, M.-F. and D. Shaw (2005). A trading-ratio system for trading water pollution discharge permits. *Journal of Environmental Economics and Management 49*(1), 83–102.

Kokorsch, M. and K. Benediktsson (2018). Prosper or perish? The development of Icelandic fishing villages after the privatisation of fishing rights. *Maritime Studies 17*, 69–83.

Krewski, D., M. Jerrett, R. T. Burnett, R. Ma, E. Hughes, Y. Shi, et al. (2009). Extended follow-up and spatial analysis of the American Cancer Society study linking particulate air pollution and mortality. Technical report, Health Effects Institute Research Report 140.

Kroetz, K., J. N. Sanchirico, and D. K. Lew (2015). Efficiency costs of social objectives in tradable permit programs. *Journal of the Association of Environmental and Resource Economists 2*(3), 339–366.

Krupnick, A. J., W. E. Oates, and E. Van De Verg (1983). On marketable air-pollution permits: The case for a system of pollution offsets. *Journal of Environmental Economics and Management 10*(3), 233–247.

Kwerel, E. (1977). To tell the truth: Imperfect information and optimal pollution control. *Review of Economic Studies 44*(3), 595–601.

Leonard, B., C. Costello, and G. D. Libecap (2019). Expanding water markets in the western united states: Barriers and lessons from other natural resource markets. *Review of Environmental Economics and Policy 13*(1), 43–61.

Montero, J.-P. (2008). A simple auction mechanism for the optimal allocation of the commons. *American Economic Review 98*(1), 496–518.

Montgomery, W. D. (1972). Markets in licenses and efficient pollution control programs. *Journal of Economic Theory 5*(3), 395–418.

Muller, N. Z. (2012). The design of optimal climate policy with air pollution co-benefits. *Resource and Energy Economics 34*(4), 696–722.

Muller, N. Z. and R. Mendelsohn (2009). Efficient pollution regulation: Getting the prices right. *American Economic Review 99*(5), 1714–1739.

Myerson, R. B. and M. A. Satterthwaite (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory 29*(2), 265–281.

National Energy Technology Laboratory (2022). Cost of capturing CO_2 from industrial sources. Technical report, U.S. Department of Energy, NETL.

Reeling, C., C. Garnache, and R. Horan (2018). Efficiency gains from integrated multipollutant trading. *Resource and Energy Economics 52*, 124–136.

Rennert, K., F. Errickson, B. C. Prest, L. Rennels, R. G. Newell, W. Pizer, C. Kingdon, J. Wingenroth, R. Cooke, B. Parthum, D. Smith, K. Cromar, D. Diaz, F. C. Moore, U. K. Müller, R. J. Plevin, A. E. Raftery, H. Ševčíková, H. Sheets, J. H. Stock, T. Tan, M. Watson, T. E. Wong, and D. Anthoff (2022). Comprehensive evidence implies a higher social cost of CO_2. *Nature 610*, 687–692.

Salzman, J. and J. B. Ruhl (2000). Currencies and the commodification of environmental law. *Stanford Law Review 53*(3), 607–694.

Schlosberg, D. (2007). *Defining Environmental Justice: Theories, Movements, and Nature*. Oxford: Oxford University Press.

South Coast Air Quality Management District (2021). Best available retrofit control technology (BARCT) assessments for reclaim facilities, staff reports. Technical report, SCAQMD.

Sutherland, S. A. and E. C. Edwards (2022). The impact of property rights to fish on remote communities in Alaska. *Land Economics 98*(2).

Tessum, C. W., J. D. Hill, and J. D. Marshall (2017). Inmap: A model for air pollution interventions. *PLoS ONE 12*(4), e0176131.

Tietenberg, T. (1995). Tradeable permits for pollution control when emission location matters: What have we learned? *Environmental and Resource Economics 5*(2), 95–113.

Tietenberg, T. H. (1985). *Emissions Trading: An Exercise in Reforming Pollution Policy*. Washington, DC: Resources for the Future.

U.S. Environmental Protection Agency (2020). Control strategy tool (CoST): Control measures database and cost equations. Technical report, U.S. EPA Office of Air Quality Planning and Standards.

Weitzman, M. L. (1974). Prices vs. quantities. *Review of Economic Studies 41*(4), 477–491.

Yang, P. and D. T. Kaffine (2016). Community-based tradable permits for localized pollution. *Environmental and Resource Economics 65*(4), 773–788.

A Proofs

Notation from the main text is maintained throughout. This appendix uses two objects beyond the main text's displays. The first is the firm's abatement response to prices,

$$\frac{\partial(a_i^A, a_i^L)}{\partial(\tau^A, \tau_i^L)} = \left(\nabla^2 A_i\right)^{-1} \equiv \begin{pmatrix} \psi_i & p_i \\ p_i & \tilde{\beta}_i \end{pmatrix}, \qquad \psi_i, \tilde{\beta}_i > 0, \quad p_i = \frac{-\gamma_i}{\det \nabla^2 A_i} \geq 0, \tag{22}$$

evaluated at the relevant allocation. Under Assumption 2, these entries are the constants defined in Equation (9) of the main text. The second is the closed form that Assumption 2 implies for firm abatement and carbon-market clearing,

$$a_i^A = \psi_i \tau^A + p_i \tau_i^L, \qquad a_i^L = \tilde{\beta}_i \tau_i^L + p_i \tau^A, \qquad \tau^A(\boldsymbol{\sigma}) = \frac{R - \sum_i p_i \sum_j \sigma_j h_{ji}}{\sum_i \psi_i}. \tag{23}$$

All further symbols are defined where they are introduced within each proof. Throughout, we maintain nonempty receptor fields, a binding cap, unbounded marginal cost on the carbon margin, interior permit supplies, and weakly concave residual inverse demand wherever communities price strategically. Table G1 in Appendix G reports numerical checks of each closed form presented below.

A.1 Proof of Proposition 1

Proof. We first characterize the planner's problem. The objective in Equation (3) is strictly convex, because each A_i is strictly convex by Assumption 1 and each D_j is convex in an affine function of abatement. The cap constraint is affine. The planner's allocation a^* is therefore unique, the first-order conditions in Equation (4) are necessary and sufficient, and the multiplier λ on the cap is unique.

We next construct an equilibrium that implements a^*. Let $c_j^* = \sum_i h_{ji}(\bar{e}_i^L - a_i^{L*})$ denote the concentrations at a^*, and set

$$\sigma_j^* = D_j'(c_j^*), \qquad \tau^A = \lambda, \qquad \tau_i^L = \sum_j \sigma_j^* h_{ji} = \mathrm{MD}_i^*. \tag{24}$$

At these prices the firm's first-order conditions $\nabla A_i(a_i) = (\tau^A, \tau_i^L)'$ are the planner's conditions in Equation (4), so they hold at a^*. The community's first-order condition $\sigma_j = D_j'(s_j)$ holds at $s_j = c_j^*$. The bounds are slack because $0 \leq c_j^* \leq c_j^0 \leq \bar{L}_j$. Both markets clear because a^* is feasible. The prices in Equation (24) together with a^* therefore satisfy Definition 1, which establishes existence and parts (a) and (b) for this equilibrium.

We now show that every equilibrium implements a^*. In any competitive equilibrium, the

community condition and permit-market clearing give $\sigma_j = D_j'(c_j)$, so

$$\tau_i^L = \sum_j \sigma_j h_{ji} = \sum_j D_j'(c_j)\, h_{ji}. \tag{25}$$

Substituting Equation (25) into the firm conditions yields the planner's first-order conditions at the equilibrium allocation, which is feasible by cap clearing. Because those conditions are sufficient, the equilibrium allocation is a^*, and $\tau^A = \lambda$ by uniqueness of the multiplier. Parts (a) and (b) therefore hold in every equilibrium.

Finally, we note that the equilibrium conditions involve only $\{D_j'\}$, $\{A_i\}$, $\mathbf{H}$, and R. The regulator's estimates $\hat{d}_j$ enter none of them, which gives part (c). □

A.2 Proof of Proposition 2

Proof. (a) We establish existence as a fixed point of the best-response map. Fix $(\boldsymbol{\sigma}_{-j}, \tau^A)$. Community j's residual demand in Equation (8) is continuous and strictly decreasing in σ_j, with slope $-\sum_i \tilde{\beta}_i h_{ji}^2 < 0$ by Equation (22) and the nonempty receptor field. Inverse demand therefore exists and is C^1. Because inverse demand is weakly concave, revenue $\sigma_j(s)s$ is strictly concave in s, and with D_j convex the objective in Equation (8) is strictly concave. The best response is therefore unique and continuous in $(\boldsymbol{\sigma}_{-j}, \tau^A)$.

We next verify that the carbon market clears at a unique price. Aggregate carbon abatement is continuous and strictly increasing in τ^A with slope $\sum_i \psi_i > 0$, falls short of R at $\tau^A = 0$, and is unbounded. A unique clearing price $\tau^A(\boldsymbol{\sigma})$ therefore exists and is continuous in $\boldsymbol{\sigma}$.

Finally, we bound the price space. No best response exceeds the choke price, which is bounded because inverse demand at zero supply is continuous on the compact set $\prod_j [0, \bar{L}_j]$. The map from price profiles to best responses, with the carbon market re-cleared, is therefore a continuous self-map of a compact convex set. Brouwer's fixed-point theorem yields an equilibrium of the pricing game.

(b) The within-community markup follows from the first-order condition for Equation (8) at an interior supply,

$$\sigma_j + \sigma_j'(s_j)\, s_j = D_j'(c_j), \tag{26}$$

which gives $\mu_j \equiv \sigma_j - D_j'(c_j) = -\sigma_j'(s_j)\, s_j > 0$. This distortion is present even with a single community.

For the cross-community distortion, we compare the Nash first-order condition with the condition for joint profit maximization. Let $\Pi(\boldsymbol{\sigma}) = \sum_k \left[\sigma_k\, s_k(\boldsymbol{\sigma}) - D_k\big(s_k(\boldsymbol{\sigma})\big)\right]$ denote joint community profit. Differentiating with respect to σ_j gives

$$\frac{\partial \Pi}{\partial \sigma_j} = \underbrace{s_j + \mu_j \frac{\partial s_j}{\partial \sigma_j}}_{\text{own marginal profit}} + \sum_{k \neq j} \mu_k \frac{\partial s_k}{\partial \sigma_j}. \tag{27}$$

A coordinated seller sets Equation (27) to zero. Community j sets only its own marginal profit to zero, and it does so using the slope of its residual demand with τ^A held fixed, $-\sum_i \tilde{\beta}_i h_{ji}^2$. The cross terms are the distortion.

To sign them, we differentiate community k's residual demand with respect to σ_j, using $\partial\tau^A/\partial\sigma_j$ from Equation (32),

$$\frac{\partial s_k}{\partial \sigma_j} = -\sum_i h_{ki}\left[\tilde{\beta}_i h_{ji} + p_i \frac{\partial \tau^A}{\partial \sigma_j}\right] = -\left(\sum_i \tilde{\beta}_i h_{ji} h_{ki} - \frac{\left(\sum_i p_i h_{ji}\right)\left(\sum_i p_i h_{ki}\right)}{\sum_i \psi_i}\right). \tag{28}$$

The first term in parentheses is the direct effect of more local abatement at shared sources, which lowers demand at k. The second is the carbon-market feedback, which lowers τ^A and raises demand at k. Their difference is the net cross-demand externality in the statement. Setting $k = j$ in Equation (28) shows that the equilibrium own slope exceeds the fixed-τ^A slope by $\left(\sum_i p_i h_{ji}\right)^2 / \sum_i \psi_i$, so community j's own marginal profit at the Nash price is $\mu_j \left(\sum_i p_i h_{ji}\right)^2 / \sum_i \psi_i \geq 0$ rather than exactly zero. Substituting into Equation (27) yields

$$\left.\frac{\partial \Pi}{\partial \sigma_j}\right|_{\text{Nash}} = \mu_j \frac{\left(\sum_i p_i h_{ji}\right)^2}{\sum_i \psi_i} - \sum_{k \neq j} \mu_k \left(\sum_i \tilde{\beta}_i h_{ji} h_{ki} - \frac{\left(\sum_i p_i h_{ji}\right)\left(\sum_i p_i h_{ki}\right)}{\sum_i \psi_i}\right). \tag{29}$$

When the net cross-demand externalities are positive and outweigh the own feedback term, Equation (29) is negative. Joint profit then rises when community j lowers its price, so every Nash price exceeds the coordinated seller's, and markups stack as in the complements-monopoly comparison of Buchanan and Yoon (2000). The own feedback term is small in the calibration because it scales with the factor $1/\sum_i \psi_i$, which is small when aggregate carbon abatement is responsive to the carbon price, and it vanishes when $p_i = 0$ for every i.

(c) We argue by contradiction. Suppose supplies are interior, so $\mu_j > 0$ for every j, and suppose the strategic-pricing allocation equals the planner's allocation a^*. Then concentrations coincide. At any source with a nonempty receptor field, the firm's local first-order condition gives

$$\frac{\partial A_i}{\partial a_i^L} = \sum_j \sigma_j^C h_{ji} > \sum_j D_j'(c_j^*)\, h_{ji} = \text{MD}_i^*, \tag{30}$$

which contradicts Equation (4). The strategic-pricing allocation therefore differs from a^*. It satisfies the cap, because the carbon market clears at R, so it lies in the planner's feasible set. Since the planner's objective is strictly convex, a^* is its unique minimizer over that set, and every other feasible allocation has strictly higher social cost. Hence $\mathcal{L}^C > 0$. □

A.3 Proof of Proposition 3

Proof. Carbon-market clearing defines $\tau^A(\boldsymbol{\sigma})$ through

$$\Phi(\tau^A, \boldsymbol{\sigma}) \equiv \sum_i a_i^A\Big(\tau^A, \textstyle\sum_j \sigma_j h_{ji}\Big) - R = 0. \tag{31}$$

The clearing price exists, is unique, and is continuous by the proof of Proposition 2. Because $\partial\Phi/\partial\tau^A = \sum_i \psi_i > 0$ and $\partial\Phi/\partial\sigma_j = \sum_i p_i h_{ji} \geq 0$, the implicit function theorem applied to Equation (31) yields

$$\frac{\partial \tau^A}{\partial \sigma_j} \;=\; -\,\frac{\sum_i p_i h_{ji}}{\sum_i \psi_i} \;\leq\; 0 \tag{32}$$

at every price vector. Integrating Equation (32) along the segment from $\mathbf{0}$ to $\boldsymbol{\sigma}$ gives $\tau^A(\boldsymbol{\sigma}) \leq \tau^A(\mathbf{0})$. The inequality is strict when $p_i h_{ji}\sigma_j > 0$ for some pair somewhere on the segment, that is, when a firm with $\gamma_i < 0$ lies in the receptor field of a community with a positive price, since $p_i = -\gamma_i/\det\nabla^2 A_i$. When $\gamma_i = 0$ for every firm, $p_i = 0$ and Equation (32) is zero, so the carbon price does not change. Under Assumption 2 the responses are linear and aggregate to Equation (23). □

A.4 Closed forms and proof of Lemma 1

Lemma 1 (Closed form and uniqueness of the strategic pricing equilibrium)**.** *Residual inverse demand is linear; best responses are*

$$\sigma_j^C = \frac{1}{2}\left(d_j + \frac{c_j^0 - \sum_{k\neq j}\big(\sum_i \tilde\beta_i h_{ji}h_{ki}\big)\sigma_k - \big(\sum_i p_i h_{ji}\big)\tau^A}{\sum_i \tilde\beta_i h_{ji}^2}\right), \qquad \mu_j = \frac{s_j^C}{\sum_i \tilde\beta_i h_{ji}^2},$$

and the equilibrium solves a linear system $\mathbf{G}\boldsymbol{\sigma}^C = \mathbf{g}$*. Under the diagonal-dominance condition in Equation (40), stated below, the equilibrium is unique.*

Proof. We first derive residual demand. Permits demanded from community j are $s_j = \sum_i h_{ji} e_i^L = c_j^0 - \sum_i h_{ji} a_i^L$. Substituting the firm response in Equation (23) with $\tau_i^L = \sum_k \sigma_k h_{ki}$ yields

$$s_j = c_j^0 \;-\; \sigma_j \sum_i \tilde\beta_i h_{ji}^2 \;-\; \sum_{k\neq j}\sigma_k \sum_i \tilde\beta_i h_{ji}h_{ki} \;-\; \tau^A \sum_i p_i h_{ji}. \tag{33}$$

Solving Equation (33) for σ_j gives the linear residual inverse demand

$$\sigma_j \;=\; \frac{\Big[c_j^0 \;-\; \sum_{k\neq j}\Big(\sum_i \tilde\beta_i h_{ji}h_{ki}\Big)\sigma_k \;-\; \Big(\sum_i p_i h_{ji}\Big)\tau^A\Big] \;-\; s_j}{\sum_i \tilde\beta_i h_{ji}^2}. \tag{34}$$

We next derive the best response. Community j chooses s to maximize $\sigma_j(s)\,s - d_j s$, taking $(\boldsymbol{\sigma}_{-j}, \tau^A)$, and hence the intercept of Equation (34), as given. The objective is strictly concave, with

second derivative $-2/\sum_i \tilde{\beta}_i h_{ji}^2$. The first-order condition yields

$$s_j^C = \frac{1}{2}\left[\Big(c_j^0 - \sum_{k\neq j}\Big(\sum_i \tilde{\beta}_i h_{ji}h_{ki}\Big)\sigma_k - \Big(\sum_i p_i h_{ji}\Big)\tau^A\Big) - \Big(\sum_i \tilde{\beta}_i h_{ji}^2\Big)d_j\right]. \tag{35}$$

Substituting Equation (35) into Equation (34) gives the best-response price in the lemma and $\mu_j = \sigma_j^C - d_j = s_j^C / \sum_i \tilde{\beta}_i h_{ji}^2$.

We now assemble the equilibrium system. In equilibrium τ^A is given by Equation (23). We substitute it into the intercept of Equation (34) and separate the $k = j$ term of $\sum_i p_i \sum_k \sigma_k h_{ki} = \sum_k \big(\sum_i p_i h_{ki}\big)\sigma_k$, which yields

$$\text{intercept} = \underbrace{c_j^0 - \frac{\big(\sum_i p_i h_{ji}\big)R}{\sum_i \psi_i}}_{\text{constant in } \boldsymbol{\sigma}} - \sum_{k\neq j}\left(\sum_i \tilde{\beta}_i h_{ji}h_{ki} - \frac{\big(\sum_i p_i h_{ji}\big)\big(\sum_i p_i h_{ki}\big)}{\sum_i \psi_i}\right)\sigma_k + \frac{\big(\sum_i p_i h_{ji}\big)^2}{\sum_i \psi_i}\sigma_j. \tag{36}$$

The own-price term in Equation (36) appears because a higher σ_j lowers τ^A and thereby raises demand for community j's own permits, which community j takes as given. We write the best-response condition as

$$2\Big(\sum_i \tilde{\beta}_i h_{ji}^2\Big)\sigma_j^C = \Big(\sum_i \tilde{\beta}_i h_{ji}^2\Big)d_j + \text{intercept}. \tag{37}$$

Substituting Equation (36) into Equation (37) and collecting price terms gives $\mathbf{G}\boldsymbol{\sigma}^C = \mathbf{g}$ with

$$G_{jj} = 2 - \frac{\big(\sum_i p_i h_{ji}\big)^2}{\big(\sum_i \tilde{\beta}_i h_{ji}^2\big)\big(\sum_i \psi_i\big)}, \qquad G_{jk} = \frac{\sum_i \tilde{\beta}_i h_{ji}h_{ki} - \dfrac{\big(\sum_i p_i h_{ji}\big)\big(\sum_i p_i h_{ki}\big)}{\sum_i \psi_i}}{\sum_i \tilde{\beta}_i h_{ji}^2} \quad (k \neq j),$$

$$g_j = d_j + \frac{c_j^0 - \big(\sum_i p_i h_{ji}\big)R/\sum_i \psi_i}{\sum_i \tilde{\beta}_i h_{ji}^2}.$$

Finally, we establish uniqueness. Let

$$x_j \equiv \frac{\big(\sum_i p_i h_{ji}\big)^2}{\big(\sum_i \tilde{\beta}_i h_{ji}^2\big)\big(\sum_i \psi_i\big)}, \tag{38}$$

so that $G_{jj} = 2 - x_j$. Positive definiteness of Equation (22) gives $p_i^2 < \psi_i \tilde{\beta}_i$. The Cauchy–Schwarz inequality then yields

$$\Big(\sum_i p_i h_{ji}\Big)^2 = \Big(\sum_i \sqrt{\psi_i}\cdot\frac{p_i h_{ji}}{\sqrt{\psi_i}}\Big)^2 \leq \Big(\sum_i \psi_i\Big)\Big(\sum_i \frac{p_i^2 h_{ji}^2}{\psi_i}\Big) < \Big(\sum_i \psi_i\Big)\Big(\sum_i \tilde{\beta}_i h_{ji}^2\Big) \tag{39}$$

whenever the receptor field is nonempty, so $x_j \in [0, 1)$. The diagonal-dominance condition is

$$\frac{1}{\sum_i \tilde{\beta}_i h_{ji}^2} \left[\sum_{k \neq j} \left| \sum_i \tilde{\beta}_i h_{ji} h_{ki} - \frac{\left(\sum_i p_i h_{ji}\right)\left(\sum_i p_i h_{ki}\right)}{\sum_i \psi_i} \right| + \frac{\left(\sum_i p_i h_{ji}\right)^2}{\sum_i \psi_i} \right] < 1 \qquad \text{for all } j. \quad (40)$$

Row j of $\mathbf{G}$ is community j's best-response condition. The diagonal entry $G_{jj} = 2 - x_j$ is the weight on its own price, and the off-diagonal entries G_{jk} are the weights on rivals' prices, which enter through shared sources net of the carbon-market feedback. Equation (40) states that $\sum_{k \neq j} |G_{jk}| < 1 - x_j$, so that $|G_{jj}| > 1 + \sum_{k \neq j} |G_{jk}|$ and $\mathbf{G}$ is strictly diagonally dominant. Each community's price thus depends more on its own demand than on all rivals' prices combined, a condition that holds when receptor fields overlap only partially. A strictly diagonally dominant matrix is invertible, so the equilibrium $\boldsymbol{\sigma}^C = \mathbf{G}^{-1}\mathbf{g}$ exists and is unique.

Equation (40) is sufficient rather than necessary. In our calibration it fails at every community definition because atmospheric transport links all receptor fields, and the game on the discrete engineering menus admits multiple equilibria (Table G2 in Appendix G). □

A.5 Proof of Proposition 4

Proof. We first express welfare in abatement space. Because the cap binds under every regime and for the planner, carbon damages are the same in every allocation and drop out of the comparison. Exchanging the order of summation, local damages are

$$\sum_j d_j c_j = \sum_j d_j \sum_i h_{ji}\left(\bar{e}_i^L - a_i^L\right) = \underbrace{\textstyle\sum_j d_j \sum_i h_{ji} \bar{e}_i^L}_{\text{constant}} - \sum_i a_i^L \, \mathrm{MD}_i. \quad (41)$$

Up to a constant, welfare is therefore $W = \sum_i \left[\mathrm{MD}_i \, a_i^L - A_i(a_i^A, a_i^L)\right]$ subject to $\sum_i a_i^A = R$.

We next expand W around the planner's allocation. Under Assumption 2, W is quadratic with Hessian $-\nabla^2 A_i$ at each firm. For any deviation $\Delta a_i = (\Delta a_i^A, \Delta a_i^L)'$ from a^* we therefore have

$$W(a) = W(a^*) + \nabla W(a^*)' \, \Delta a - \frac{1}{2} \sum_i \Delta a_i' \left(\nabla^2 A_i\right) \Delta a_i, \quad (42)$$

with no remainder. By Equation (4), $\partial W / \partial a_i^A = -\lambda$ and $\partial W / \partial a_i^L = 0$ at a^*. Both allocations meet the cap, so $\sum_i \Delta a_i^A = 0$ and the gradient term in Equation (42) vanishes.

We now express the deviations in terms of price wedges. Differencing the firm conditions $\left(\nabla^2 A_i\right) a_i = (\tau^A, \tau_i^L)'$ and the planner conditions $\left(\nabla^2 A_i\right) a_i^* = (\lambda, \mathrm{MD}_i)'$ gives $\Delta a_i = \left(\nabla^2 A_i\right)^{-1} (\Delta^A, \Delta_i^L)'$. Substituting into Equation (42) yields

$$\mathcal{L} = W(a^*) - W(a) = \frac{1}{2} \sum_i \Delta a_i' \left(\nabla^2 A_i\right) \Delta a_i = \frac{1}{2} \sum_i \begin{pmatrix} \Delta^A \\ \Delta_i^L \end{pmatrix}' \left(\nabla^2 A_i\right)^{-1} \begin{pmatrix} \Delta^A \\ \Delta_i^L \end{pmatrix}. \quad (43)$$

Writing the inverse Hessian in the constants defined in Equation (9), we obtain

$$\mathcal{L} = \sum_i \left[\frac{\psi_i}{2}(\Delta^A)^2 + p_i\,\Delta^A \Delta_i^L + \frac{\tilde{\beta}_i}{2}(\Delta_i^L)^2 \right], \qquad \Delta^A = -\frac{\sum_i p_i \Delta_i^L}{\sum_i \psi_i}. \tag{44}$$

The second equation in Equation (44) follows from differencing the two market-clearing conditions $\left(\sum_i \psi_i\right)\tau^A = R - \sum_i p_i \tau_i^L$ and $\left(\sum_i \psi_i\right)\lambda = R - \sum_i p_i \mathrm{MD}_i$, both obtained by aggregating Equation (23). Substituting it into the first equation, the $(\Delta^A)^2$ terms contribute $\left(\sum_i p_i \Delta_i^L\right)^2 / \left(2\sum_i \psi_i\right)$ and the cross terms contribute $-\left(\sum_i p_i \Delta_i^L\right)^2 / \sum_i \psi_i$. Their sum is the negative correction in Equation (10).

Finally, we sign the loss. Each firm's term in Equation (43), $\frac{1}{2} v_i' \left(\nabla^2 A_i\right)^{-1} v_i$ with $v_i = (\Delta^A, \Delta_i^L)'$, is nonnegative because the inverse Hessian is positive definite (Assumption 1). Hence $\mathcal{L} \geq 0$, with equality only if $\Delta^A = 0$ and $\Delta_i^L = 0$ at every firm. If every local wedge vanishes, the second equation of Equation (44) forces $\Delta^A = 0$. The loss is therefore zero exactly when $\Delta_i^L = 0$ for all i. □

A.6 Proof of Proposition 5

Proof. (a) With linear damages, $D_j' \equiv d_j$, so Proposition 1(a) gives $\sigma_j = d_j$ for every realization of $\hat{\mathbf{d}}$. Then $\Delta_i^L = \sum_j d_j h_{ji} - \mathrm{MD}_i = 0$ for all i. The second equation of Equation (44) gives $\Delta^A = 0$, and Equation (10) gives $\mathcal{L} = 0$.

We next verify that no other regime in Table 1 has this property. Regimes I/III have $\Delta_i^L = t^*(\hat{\mathbf{d}}) - \mathrm{MD}_i$, which cannot vanish at every source unless MD_i is constant across sources, the degenerate case in which a uniform price is first best. Regimes II/IV have $\Delta_i^L = b_i \neq 0$ for some realizations of beliefs. Regime V under strategic pricing has $\Delta_i > 0$ at interior supplies by Lemma 1.

(b) By Lemma 1, the equilibrium solves $\mathbf{G}\boldsymbol{\sigma}^C = \mathbf{g}$. Every entry of $\mathbf{G}$ and $\mathbf{g}$ is a function of $\mathbf{H}$, $\{\psi_i, p_i, \tilde{\beta}_i\}$, $\{d_j\}$, and R alone. The markups $\mu_j = \sigma_j^C - d_j$, the source wedges $\Delta_i = \sum_j \mu_j h_{ji}$, and the loss $\mathcal{L}(\{\Delta_i\})$ therefore depend on demand primitives and true damages only. The estimates $\hat{\mathbf{d}}$ do not enter.

(c) Regimes II/IV post $\tau_i^L = \widehat{\mathrm{MD}}_i$, so $\Delta_i^L = \widehat{\mathrm{MD}}_i - \mathrm{MD}_i = b_i$. The bias enters the local wedge one for one. Regimes I/III post $t^*(\hat{\mathbf{d}})$, which is linear in $\{\widehat{\mathrm{MD}}_i\}$ and hence in $\hat{\mathbf{d}}$ (Appendix D). The bias therefore enters the wedge profile $\{t^*(\hat{\mathbf{d}}) - \mathrm{MD}_i\}$ linearly. In either case, when $\gamma_i < 0$ for some i, the induced carbon wedge

$$\Delta^A = -\frac{\sum_i p_i \Delta_i^L}{\sum_i \psi_i} \tag{45}$$

is nonzero whenever $\sum_i p_i b_i \neq 0$. The cross terms $p_i \Delta^A \Delta_i^L$ in Equation (44) represent the resulting carbon-market misallocation. This channel is absent under separability, where $p_i = 0$.

(d) Each regime's loss is the eliminated form in Equation (10) evaluated at that regime's wedge profile. Regime V (strategic pricing) therefore dominates Regime IV exactly when $\mathcal{L}(\{\Delta_i\}) < \mathcal{L}(\{b_i\})$. We show that neither ordering is universal. At $\hat{\mathbf{d}} = \mathbf{d}$, $\mathcal{L}(\{b_i\}) = 0 < \mathcal{L}(\{\Delta_i\})$ at interior

strategic-pricing supplies. Now hold primitives fixed and scale the bias, with $\hat{\mathbf{d}} = \mathbf{d} + \zeta \mathbf{u}$ for fixed $\mathbf{u}$ such that $\mathcal{L}(\{\sum_j u_j h_{ji}\}) > 0$. Then $\mathcal{L}(\{b_i\}) \to \infty$ as $\zeta \to \infty$, while $\mathcal{L}(\{\Delta_i\})$ is unchanged. The comparison is therefore quantitative, and Section 5 presents an illustrative calibration. □

A.7 Proof of Proposition 6

Proof. The proof uses the fact that the welfare loss is a quadratic function of the price wedges, so its expectation under estimation error can be expressed as a variance. Stacking the source wedges Δ_i^L in the vector $\mathbf{\Delta}^L$, the eliminated form in Equation (10) is

$$\mathcal{L} = \tfrac{1}{2}(\mathbf{\Delta}^L)^T N \mathbf{\Delta}^L, \qquad N \equiv \operatorname{diag}(\tilde{\beta}_i) - \frac{\mathbf{p}\mathbf{p}^T}{\sum_i \psi_i}, \tag{46}$$

where the diagonal of N contains the local-margin curvatures and the rank-one correction represents the carbon-market feedback. N is positive semidefinite because the loss is nonnegative (Proposition 4). Three objects built from N appear below. The matrix $M \equiv \mathbf{H} N \mathbf{H}^T$ is the same loss written as a quadratic form in community-level errors rather than source-level wedges, since a community error ε_j affects source i through h_{ji}. The scalar $\Lambda \equiv \mathbf{1}^T N \mathbf{1}$ is the curvature of the loss in a uniform price, as in the statement. The vector $\mathbf{w} \equiv \mathbf{H} N \mathbf{1}/\Lambda$ collects the weights with which each community's error enters the optimal uniform price.

(a) Regimes II/IV post $\widehat{\mathrm{MD}}_i$, so the wedge is the bias $b_i = \sum_j \varepsilon_j h_{ji}$, or $\mathbf{b} = \mathbf{H}^T \boldsymbol{\varepsilon}$. Substituting into Equation (46) yields $\mathcal{L} = \frac{1}{2}\mathbf{b}^T N \mathbf{b} = \frac{1}{2}\boldsymbol{\varepsilon}^T M \boldsymbol{\varepsilon}$. For mean-zero errors, $\mathbb{E}[\boldsymbol{\varepsilon}^T M \boldsymbol{\varepsilon}] = \operatorname{tr}(M\Sigma)$. Because $\mathbf{b}$ has mean zero, this equals

$$\tfrac{1}{2}\operatorname{tr}(M\Sigma) = \tfrac{1}{2}\sum_i \tilde{\beta}_i \operatorname{Var}(b_i) - \frac{\operatorname{Var}\left(\sum_i p_i b_i\right)}{2\sum_i \psi_i}. \tag{47}$$

(b) Under the uniform instruments the regulator sets one price from its estimates. The error therefore enters the loss through the difference between that price and the perfect-information price. A quadratic loss evaluated at a distance s from its minimum exceeds the minimum by half the curvature times s^2. The argument consists of making those two facts precise.

We first compute the price difference. The optimal uniform price is linear in beliefs (Appendix D), so $t^*(\hat{\mathbf{d}}) = t^*(\mathbf{d}) + s$ with

$$s \;=\; \frac{\mathbf{1}^T N\, \mathbf{b}}{\mathbf{1}^T N\, \mathbf{1}} \;=\; \frac{(\mathbf{H}N\mathbf{1})^T \boldsymbol{\varepsilon}}{\Lambda} \;=\; \mathbf{w}^T \boldsymbol{\varepsilon}. \tag{48}$$

The price set from the estimates therefore differs from the perfect-information price by s, a weighted average of the community errors with weights $\mathbf{w}$.

We next evaluate the loss at that price. As a function of a uniform price t, the Regime I/III loss is

$$\mathcal{L}(t) = \tfrac{1}{2}(t\mathbf{1} - \mathbf{y})^T N (t\mathbf{1} - \mathbf{y}), \qquad y_i = \mathrm{MD}_i. \tag{49}$$

This is a parabola in t with its minimum at the perfect-information price $t^*(\mathbf{d})$ and curvature Λ, so setting the price s away from the minimum raises the loss by $\frac{1}{2}\Lambda s^2$. Taking expectations, and using that s is the price difference so that $\mathbf{w}^T\Sigma\,\mathbf{w} = \mathrm{Var}\big(t^*(\hat{\mathbf{d}})\big)$,

$$\mathcal{L}\big(t^*(\hat{\mathbf{d}})\big) = \mathcal{L}_0^{\text{undiff}} + \tfrac{1}{2}\,\Lambda\, s^2, \qquad \mathbb{E}\big[\mathcal{L}\big] = \mathcal{L}_0^{\text{undiff}} + \tfrac{1}{2}\,\Lambda\,\mathrm{Var}\big(t^*(\hat{\mathbf{d}})\big), \tag{50}$$

which is Equation (12).

We now compare the two expected losses. The differentiated loss $\frac{1}{2}\mathbf{b}^T N\mathbf{b}$ depends on the entire bias vector, whereas the uniform loss $\frac{1}{2}\Lambda s^2$ depends only on the component of the bias along the uniform direction $\mathbf{1}$, measured in the inner product defined by N. Because s is the coefficient of that component, the Cauchy–Schwarz inequality gives $\mathbf{b}^T N\mathbf{b} \geq \Lambda s^2$ for every realization. Equality holds only when $N^{1/2}\mathbf{b}$ is proportional to $N^{1/2}\mathbf{1}$ (when the bias is a uniform shift across sources up to directions that carry no loss). Taking expectations, $\mathrm{tr}(M\Sigma) \geq \Lambda\mathbf{w}^T\Sigma\mathbf{w}$ for every Σ.

Finally, we scale the error. Under $\Sigma = \mathrm{CV}^2\Sigma_0$, both variance terms scale with CV^2, while the perfect-information loss $\mathcal{L}_0^{\text{undiff}}$ does not depend on the error at all. The gain from differentiation is therefore a constant minus a term proportional to CV^2,

$$\mathbb{E}[\mathcal{L}^{\text{undiff}}] - \mathbb{E}[\mathcal{L}^{\text{diff}}] = \mathcal{L}_0^{\text{undiff}} - \tfrac{1}{2}\,\mathrm{CV}^2\,\big[\mathrm{tr}(M\Sigma_0) - \Lambda\mathbf{w}^T\Sigma_0\mathbf{w}\big]. \tag{51}$$

The bracket is positive unless the induced bias is almost surely a uniform shift, so the gain is strictly decreasing in CV^2. It is positive at $\mathrm{CV} = 0$ whenever MD_i varies across sources. It therefore crosses zero exactly once, at

$$\mathrm{CV}^{*2} = \frac{2\,\mathcal{L}_0^{\text{undiff}}}{\mathrm{tr}(M\Sigma_0) - \Lambda\mathbf{w}^T\Sigma_0\mathbf{w}}, \tag{52}$$

which is the linear-quadratic reversal, at which an undifferentiated market yields lower expected losses than the differentiated instruments. Fowlie and Muller (2019) derive the two-firm version of this comparison in their quadratic model and document a different reversal empirically that is driven by underestimated abatement costs under cap-and-trade.

(c) For price-taking, Proposition 5(a) holds at every realization. For monopoly pricing at $n_j = 1$, Proposition 5(b) shows that the loss depends on demand primitives and true damages only. It is therefore constant in ε. □

A.8 Proof of Corollary 1

Proof. We first derive the cap that attains full first best. Full first best requires, in addition to the conditions of Proposition 1, that the cap's shadow price equal the social cost of carbon, $\lambda = \mathrm{SCC}$. From the planner's clearing condition $\big(\sum_i \psi_i\big)\lambda = R - \sum_i p_i \mathrm{MD}_i^*$, the cap that delivers this is

$$R^* = \Big(\sum_i \psi_i\Big)\,\mathrm{SCC} + \sum_i p_i \mathrm{MD}_i^*. \tag{53}$$

Setting the cap in Equation (53) requires a forecast of the co-abatement term $\sum_i p_i \mathrm{MD}_i^*$, which is a function of $\{d_j\}$.

We next evaluate the cap set under beliefs. The regulator sets $\hat{R} = \big(\sum_i \psi_i\big)\mathrm{SCC} + \sum_i p_i \widehat{\mathrm{MD}}_i$, and the realized shadow price is

$$\lambda(\hat{R}) = \frac{\hat{R} - \sum_i p_i \mathrm{MD}_i}{\sum_i \psi_i} = \mathrm{SCC} + \frac{\sum_i p_i b_i}{\sum_i \psi_i}. \tag{54}$$

With price-taking communities the local wedges are zero, so the only loss relative to full first best comes from the cap. The marginal social value of a unit of required abatement is $\mathrm{SCC} - \lambda(R)$, with $\lambda(R) = \big(R - \sum_i p_i \mathrm{MD}_i\big)/\sum_i \psi_i$. Integrating from R^* to $\hat{R}$ yields

$$\mathcal{L}^{\mathrm{cap}} = \int_{R^*}^{\hat{R}} \big(\lambda(R) - \mathrm{SCC}\big)\, dR = \int_{R^*}^{\hat{R}} \frac{R - R^*}{\sum_i \psi_i}\, dR = \frac{(\hat{R} - R^*)^2}{2\sum_i \psi_i} = \frac{\big(\sum_i p_i b_i\big)^2}{2\sum_i \psi_i}, \tag{55}$$

where the last step uses $\hat{R} - R^* = \sum_i p_i b_i$.

Finally, we consider the aggregate tax. A tax at SCC fixes $\tau^A = \mathrm{SCC}$ without a forecast. With $\tau_i^L = \mathrm{MD}_i$ from price-taking communities, the firm conditions $\nabla A_i = (\mathrm{SCC}, \mathrm{MD}_i)'$ are the full-first-best conditions. The aggregate tax combined with community permits therefore attains zero loss for any $\hat{\mathbf{d}}$. □

A.9 Proof of Proposition 7

Proof. We first solve the within-community game at a given n_j. Community j is divided into n_j sub-sellers. Each sub-seller ℓ chooses its own supply $s_{j\ell} \geq 0$, taking the supplies of the other sub-sellers $S_{j,-\ell}$, the prices of other communities $\boldsymbol{\sigma}_{-j}$, and the carbon price τ^A as given. It receives the community price σ_j on its own sales and bears a $1/n_j$ share of the community's damages. Under Assumption 2, damages are linear, so the problem in Equation (15) becomes

$$\max_{s_{j\ell} \geq 0} \;\; \sigma_j\big(s_{j\ell} + S_{j,-\ell}\big)\, s_{j\ell} \;-\; \frac{d_j}{n_j}\big(s_{j\ell} + S_{j,-\ell}\big), \tag{56}$$

where $\sigma_j(\cdot)$ is the linear residual inverse demand expressed in Equation (34), with slope $-1/\sum_i \tilde{\beta}_i h_{ji}^2$. The objective in Equation (56) is strictly concave in $s_{j\ell}$, and its interior first-order condition is

$$\sigma_j + \sigma_j'(s_j)\, s_{j\ell} = \frac{d_j}{n_j} \quad \Longleftrightarrow \quad \sigma_j - \frac{s_{j\ell}}{\sum_i \tilde{\beta}_i h_{ji}^2} = \frac{d_j}{n_j}. \tag{57}$$

Every sub-seller's condition in Equation (57) determines the same offer, so any interior equilibrium is symmetric. Summing Equation (57) over ℓ gives $n_j\sigma_j - s_j/\sum_i \tilde{\beta}_i h_{ji}^2 = d_j$. Substituting demand from Equation (34) yields the price in the proposition. Setting $n_j = 1$ recovers the unitary best response of Lemma 1.

The optimal number of sub-sellers $n_j^\star$ is the number at which the price in Equation (16) equals marginal damage. Setting $\sigma_j(n_j) = d_j$ and solving for n_j gives Equation (17), once rivals' prices

are set to their marginal damages, $\boldsymbol{\sigma}_{-j} = \mathbf{d}_{-j}$, and the carbon price to its first-best value λ.

We now verify that $\boldsymbol{\sigma} = \mathbf{d}$ is the joint equilibrium at $\mathbf{n}^\star$. At that profile, the carbon price from Equation (23) is

$$\tau^A(\mathbf{d}) = \frac{R - \sum_i p_i \sum_j d_j h_{ji}}{\sum_i \psi_i} = \frac{R - \sum_i p_i \mathrm{MD}_i}{\sum_i \psi_i} = \lambda, \tag{58}$$

where the middle step exchanges the order of summation. By Equation (58), every community's intercept takes its first-best value, and each community's aggregated condition holds at $\sigma_j = d_j$ by construction of $n_j^\star$. Hence $\boldsymbol{\sigma} = \mathbf{d}$ is a fixed point of the joint system.

Finally, we establish uniqueness. Under Equation (40) the system remains strictly diagonally dominant at any $\mathbf{n} \geq \mathbf{1}$. Relative to the Nash system of Lemma 1, each diagonal entry changes from $2 - x_j$ to $n_j + 1 - x_j$, with x_j defined in Equation (38), while the off-diagonal entries are unchanged. The fixed point is therefore the unique equilibrium. At this profile $\mu_j = 0$ for every j, hence $\Delta_i = 0$ for every i, so both the within-community markup and the cross-community stacking vanish. □

A.10 Proof of Corollary 2

Proof. The corollary specializes Proposition 7 to a subdivision that gives every community the same number of sub-sellers $\bar{n}$. Throughout, we write

$$I_j \equiv c_j^0 - \Big(\sum_i p_i h_{ji}\Big)\,\tau^A, \qquad B_j \equiv \sum_i \tilde{\beta}_i h_{ji}^2, \tag{59}$$

for the intercept and slope of community j's residual inverse demand once cross-community spillovers are negligible, with τ^A the carbon price that clears the aggregate market under the subdivision.

(a) With negligible spillovers, the cross-community terms in Equation (16) vanish, and Proposition 7(a) at $n_j = \bar{n}$ gives

$$\sigma_j(\bar{n}) = \frac{d_j + I_j/B_j}{\bar{n} + 1}. \tag{60}$$

By the definition of $n_j^\star(\bar{n})$ in the statement, which is Equation (17) with the cross terms dropped and the carbon price at τ^A, we have $I_j/B_j = n_j^\star(\bar{n})\, d_j$. Substituting into Equation (60) gives Equation (21). The markup is

$$\mu_j(\bar{n}) = \sigma_j(\bar{n}) - d_j = d_j\,\frac{n_j^\star(\bar{n}) - \bar{n}}{\bar{n} + 1}, \tag{61}$$

whose magnitude falls at rate $1/(\bar{n}+1)$ and whose sign matches the sign of $n_j^\star(\bar{n}) - \bar{n}$. Community j therefore prices above marginal damage when it has fewer sub-sellers than $n_j^\star(\bar{n})$ and below it when it has more. By Proposition 4, community j's contribution to the welfare loss is proportional to $\mu_j(\bar{n})^2$, which is zero at $\bar{n} = n_j^\star(\bar{n})$ and rises on either side, so the aggregate loss is U-shaped in $\bar{n}$.

The carbon price in I_j must be the one that clears the market under the subdivision rather than λ because, by Proposition 3, the carbon price depends on the community prices, and the community prices depend on $\bar{n}$. The two coincide only when $p_i = 0$ for every i, in which case I_j does not involve τ^A.

(b) At $\bar{n} = 1$, Equation (61) gives the unitary markup $d_j\,(n_j^\star(\bar{n}) - 1)/2$. For a community with $n_j^\star(\bar{n}) \geq \bar{n}$, the ratio of the markup at $\bar{n}$ to the unitary markup is

$$\frac{\mu_j(\bar{n})}{d_j\,(n_j^\star(\bar{n}) - 1)/2} = \frac{2}{\bar{n}+1} \cdot \frac{n_j^\star(\bar{n}) - \bar{n}}{n_j^\star(\bar{n}) - 1} \leq \frac{2}{\bar{n}+1}, \tag{62}$$

because the second factor is at most one when $\bar{n} \geq 1$. Since the loss contribution is proportional to the squared markup, it falls by at least the factor $4/(\bar{n}+1)^2$.

(c) The equilibrium system is a function of $\mathbf{H}$, $\{\psi_i, p_i, \tilde{\beta}_i\}$, $\{d_j\}$ and $(R, \bar{n})$. The regulator's estimates enter nowhere, so the argument of Proposition 5(b) applies unchanged and the loss is deterministic.

(d) At $\bar{n} = 1$, Equation (60) is the unitary best response of Lemma 1. As $\bar{n} \to \infty$, Equation (21) gives $\sigma_j(\bar{n}) \to 0$ for every j. With every local price at zero, firms face no local instrument, which is the carbon-only status quo. The two bounds follow from Equation (61). When the subdivision overshoots ($\bar{n} \geq n_j^\star(\bar{n})$), the price is below marginal damage and $|\mu_j(\bar{n})| = d_j\,(\bar{n} - n_j^\star(\bar{n}))/(\bar{n}+1) < d_j$. When it undershoots ($\bar{n} < n_j^\star(\bar{n})$), the price is above marginal damage and the markup decreases in $\bar{n}$ from its unitary maximum at $\bar{n} = 1$. □

A.11 Proof of Proposition 8

Proof. The regulator chooses the subdivision $\hat{\mathbf{n}}$ from its estimates $\hat{\mathbf{d}}$ using Equation (18). Communities then play the pricing game of Proposition 7 at the true damages $\mathbf{d}$. Part (a) characterizes the resulting prices, part (b) their expected welfare loss, and part (c) the comparison with the differentiated instruments. In this proof only, we write

$$B_j \equiv \sum_i \tilde{\beta}_i h_{ji}^2, \qquad P_j \equiv \sum_i p_i h_{ji}, \qquad \Psi \equiv \sum_i \psi_i. \tag{63}$$

(a) By Proposition 7(a) at $n_j = \hat{n}_j$, community j's price satisfies

$$(\hat{n}_j + 1)\,\sigma_j = d_j + \frac{\text{intercept}_j}{B_j}, \qquad \text{intercept}_j = c_j^0 - \sum_{k \neq j} \Big(\sum_i \tilde{\beta}_i h_{ji} h_{ki} \Big)\,\sigma_k - P_j\,\tau^A. \tag{64}$$

The intercept depends on the other communities' prices directly, through shared sources, and indirectly, through the carbon price τ^A (Proposition 3). Substituting τ^A from Equation (23) and collecting all price terms on the left, as in the proof of Lemma 1, turns Equation (64) into the linear system

$$(\mathcal{N} + F)\,\boldsymbol{\sigma} = \mathbf{d} + \boldsymbol{\kappa}, \qquad \mathcal{N} \equiv \operatorname{diag}(\hat{n}_j + 1), \qquad \kappa_j = \frac{c_j^0 - P_j R/\Psi}{B_j}, \tag{65}$$

where the matrix F collects the spillovers,

$$F_{jj} = -\frac{P_j^2}{B_j \Psi}, \qquad F_{jk} = \frac{\sum_i \tilde{\beta}_i h_{ji} h_{ki} - P_j P_k / \Psi}{B_j} \quad (k \neq j). \tag{66}$$

Under Equation (40), row j of $\mathcal{N} + F$ has diagonal $\hat{n}_j + 1 - P_j^2/(B_j\Psi) \geq 2 - P_j^2/(B_j\Psi) > 1 > \sum_{k \neq j} |F_{jk}|$ for any $\hat{n}_j \geq 1$, so $\mathcal{N}+F$ is strictly diagonally dominant and invertible. The equilibrium prices are therefore $\boldsymbol{\sigma}(\mathbf{d}) = (\mathcal{N} + F)^{-1}(\mathbf{d} + \boldsymbol{\kappa})$, an affine function of the realized damages.

We now take the limit in which spillovers are negligible, so that $\sum_i \tilde{\beta}_i h_{ji} h_{ki} \to 0$ for $j \neq k$ and $P_j^2/\Psi \to 0$. Then $F \to 0$, and the intercept in Equation (64) no longer depends on any community's price. Its value is therefore the same whether damages are $\mathbf{d}$ or $\hat{\mathbf{d}}$. Under the rule in Equation (18), $\hat{n}_j$ is the number of sub-sellers at which the price would equal $\hat{d}_j$ if damages were $\hat{\mathbf{d}}$, which by Equation (64) implies $\text{intercept}_j / B_j = \hat{n}_j \hat{d}_j$. Substituting into Equation (64) gives

$$\sigma_j = \frac{d_j + \hat{n}_j \hat{d}_j}{\hat{n}_j + 1} = \theta_j \, d_j + (1 - \theta_j) \, \hat{d}_j, \qquad \theta_j = \frac{1}{\hat{n}_j + 1}, \tag{67}$$

the convex combination in the statement. The community's own information enters with weight θ_j and the regulator's estimate with weight $1 - \theta_j$, so a finer subdivision puts more weight on the estimate. Since $\hat{n}_j \geq 1$, $\theta_j \in (0, 1)$. The differentiated instruments, which post the estimates directly, are the case $\theta_j = 0$.

For completeness, we record the exact deviation from Equation (67) when spillovers are not negligible. Let $\boldsymbol{\sigma}^\theta \equiv \mathcal{N}^{-1}(\mathbf{d} + \boldsymbol{\kappa} - F\hat{\mathbf{d}})$, whose j-th entry is Equation (67) because $\kappa_j - (F\hat{\mathbf{d}})_j$ is the intercept evaluated at the estimates, which Equation (18) sets to $\hat{n}_j \hat{d}_j$. Subtracting its definition from Equation (65) gives

$$\boldsymbol{\sigma}(\mathbf{d}) - \boldsymbol{\sigma}^\theta = (\mathcal{N} + F)^{-1} F \left(\hat{\mathbf{d}} - \boldsymbol{\sigma}^\theta \right). \tag{68}$$

(b) Write $v_j = \sigma_j - d_j$ for community j's pricing error, so the source wedge is $\Delta_i^L = \sum_j v_j h_{ji}$. Substituting into the loss in Equation (10) and using $\sum_i p_i \Delta_i^L = \sum_j v_j P_j$ gives

$$\mathcal{L}(\mathbf{v}) = \frac{1}{2} \, \mathbf{v}^T M \, \mathbf{v}, \qquad M = \mathbf{H} \, \text{diag}(\tilde{\beta}_i) \, \mathbf{H}^T - \frac{1}{\Psi} \, \mathbf{P}\mathbf{P}^T, \tag{69}$$

which is the matrix M of the proof of Proposition 6, with entries $M_{jk} = \sum_i \tilde{\beta}_i h_{ji} h_{ki} - P_j P_k / \Psi$. M is positive semidefinite because Equation (69) is a welfare loss, which is nonnegative by Proposition 4. Its diagonal entries are strictly positive, because $\mathbf{v} = \mathbf{e}_j$ produces the wedge profile $\Delta_i^L = h_{ji}$, which is nonzero at some source by the nonempty receptor field and so carries a strictly positive loss.

In the negligible-spillover limit, Equation (67) gives the mechanism's pricing errors $v_j = (1 - \theta_j)(\hat{d}_j - d_j)$, or $\mathbf{v}^{\text{mech}} = \bar{W}\boldsymbol{\varepsilon}$ with $\bar{W} = \text{diag}(1 - \theta_j)$, while the differentiated instruments post the estimates and have $\mathbf{v}^{\text{diff}} = \boldsymbol{\varepsilon}$. Because $\boldsymbol{\varepsilon}$ has mean zero and covariance Σ, the identity $\mathbb{E}[\boldsymbol{\varepsilon}^T Q \boldsymbol{\varepsilon}] =$

$\mathrm{tr}(Q\Sigma)$ applied to Equation (69) gives

$$\mathbb{E}[\mathcal{L}^{\mathrm{diff}}] = \tfrac{1}{2}\mathrm{tr}(M\Sigma), \qquad \mathbb{E}[\mathcal{L}^{\mathrm{mech}}] = \tfrac{1}{2}\mathrm{tr}(\bar{W}M\bar{W}\Sigma). \tag{70}$$

The source wedge produced by $\bar{W}\boldsymbol{\varepsilon}$ is the damped bias b_i^{θ} of the statement, so the second expression in Equation (70) is Equation (20).

(c) Subtracting the two expressions in Equation (70),

$$\mathbb{E}[\mathcal{L}^{\mathrm{diff}}] - \mathbb{E}[\mathcal{L}^{\mathrm{mech}}] = \tfrac{1}{2}\mathrm{tr}\big((M - \bar{W}M\bar{W})\Sigma\big), \tag{71}$$

where $M - \bar{W}M\bar{W}$ has (j,k) entry $M_{jk}\,\Xi_{jk}$ with $\Xi_{jk} = 1 - (1-\theta_j)(1-\theta_k)$. We consider three cases.

(i) If the errors are independent across communities, Σ is diagonal with entries ς_j^2, and the trace in Equation (71) is $\sum_j \varsigma_j^2\, M_{jj}\, \theta_j(2-\theta_j)$. Every term is nonnegative because $M_{jj} > 0$ and $\theta_j \in (0,1)$. The sum is strictly positive whenever some $\varsigma_j^2 > 0$, so the mechanism strictly dominates.

(ii) If $\theta_j = \bar{\theta}$ for every j, then $\bar{W}M\bar{W} = (1-\bar{\theta})^2 M$ exactly, so $\mathbb{E}[\mathcal{L}^{\mathrm{mech}}] = (1-\bar{\theta})^2\,\mathbb{E}[\mathcal{L}^{\mathrm{diff}}]$ for any Σ.

(iii) With heterogeneous θ and correlated errors, dominance can fail, because Ξ need not be positive semidefinite. At $\theta = (0.1, 0.9)$, $\Xi = \left(\begin{smallmatrix} 0.19 & 0.91 \\ 0.91 & 0.99 \end{smallmatrix}\right)$ has determinant $0.19 \times 0.99 - 0.91^2 = -0.64 < 0$. For two communities, Equation (71) is

$$\mathrm{tr}\big((M - \bar{W}M\bar{W})\Sigma\big) = \varsigma_1^2 M_{11}\Xi_{11} + \varsigma_2^2 M_{22}\Xi_{22} + 2\rho\varsigma_1\varsigma_2 M_{12}\Xi_{12}, \tag{72}$$

which is negative when $M_{12} > 0$, Ξ_{12} is large relative to Ξ_{11}, and the correlation ρ is sufficiently negative. □

B Microfoundation: Three Abatement Margins

This appendix derives Assumption 1 from a technology with three abatement margins and presents how it is applied in the calibration. The margin-level notation introduced here, $(\phi_i, c_i^q, c_i^c, c_i^g)$ and the margin quantities (a_i^q, a_i^c, a_i^g), is used only in this appendix and in the calibration discussion of Section 5 and Appendix F. Firm i has a *three-margin technology* if it abates through three margins with quadratic costs. Joint CO_2–NO_x abatement a_i^q has cost $\frac{c_i^q}{2}(a_i^q)^2$ and cuts aggregate emissions one-for-one and local emissions at the co-emission ratio $\phi_i \geq 0$. NO_x-only abatement a_i^c has cost $\frac{c_i^c}{2}(a_i^c)^2$ and cuts local emissions only, while CO_2-only capture a_i^g has cost $\frac{c_i^g}{2}(a_i^g)^2$ and cuts aggregate emissions only.

Lemma 2 (Microfoundation). *(a) The induced cost of any three-margin technology, the minimum margin cost of delivering (a_i^A, a_i^L), is exactly the quadratic of Assumption 2, with*

$$\alpha_i^A = \frac{c_i^g(c_i^q + \phi_i^2 c_i^c)}{c_i^g + c_i^q + \phi_i^2 c_i^c}, \qquad \alpha_i^L = \frac{c_i^c(c_i^g + c_i^q)}{c_i^g + c_i^q + \phi_i^2 c_i^c}, \qquad \gamma_i = -\frac{\phi_i\, c_i^c c_i^g}{c_i^g + c_i^q + \phi_i^2 c_i^c} \leq 0.$$

(b) $\gamma_i = 0$ if and only if $\phi_i = 0$ (separability). Shutting down NO_x-only abatement and CO_2-only capture ($c_i^c, c_i^g \to \infty$) gives fixed proportions $a_i^L = \phi_i a_i^A$. In that one-pollutant limit, the composite marginal damage of activity is SCC $+ \phi_i \sum_j d_j h_{ji}$ *(with* SCC *the social cost of carbon, as in Corollary 1), and carbon-only regulation is the undifferentiated instrument of Fowlie and Muller (2019) (Appendix E). (c) $|\gamma_i|$ is strictly increasing in c_i^g, so cheaper capture weakens the carbon-to-local linkage. Moreover, part (a) extends to general strictly convex margin costs, with local curvatures replacing the constant slopes, so any three-margin technology satisfies the sign restriction of Assumption 1.*

Proof. Part (a). Suppressing i, the induced cost is

$$C(a^A, a^L) = \min_{a^g, a^q, a^c} \frac{c^g}{2}(a^g)^2 + \frac{c^q}{2}(a^q)^2 + \frac{c^c}{2}(a^c)^2 \quad \text{s.t.} \quad a^q + a^g = a^A, \;\; \phi a^q + a^c = a^L.$$

Eliminate the constraints by setting $q \equiv a^q$, $a^g = a^A - q$, $a^c = a^L - \phi q$; the objective $f(q)$ is strictly convex with $f''(q) = c^g + c^q + \phi^2 c^c$, so the interior first-order condition $q\,(c^g + c^q + \phi^2 c^c) = c^g a^A + \phi c^c a^L$ gives

$$q^* = \frac{c^g a^A + \phi c^c a^L}{c^g + c^q + \phi^2 c^c}.$$

By the envelope theorem, $C_A = c^g(a^A - q^*)$ and $C_L = c^c(a^L - \phi q^*)$, where letter subscripts on C denote partial derivatives. Differentiating again with $\partial q^*/\partial a^A = c^g/(c^g + c^q + \phi^2 c^c)$ and $\partial q^*/\partial a^L = \phi c^c/(c^g + c^q + \phi^2 c^c)$,

$$C_{AA} = c^g\Big(1 - \frac{c^g}{c^g + c^q + \phi^2 c^c}\Big) = \frac{c^g(c^q + \phi^2 c^c)}{c^g + c^q + \phi^2 c^c} = \alpha^A,$$
$$C_{LL} = c^c\Big(1 - \frac{\phi^2 c^c}{c^g + c^q + \phi^2 c^c}\Big) = \frac{c^c(c^g + c^q)}{c^g + c^q + \phi^2 c^c} = \alpha^L,$$
$$C_{AL} = -c^g \frac{\partial q^*}{\partial a^L} = -\frac{\phi\, c^c c^g}{c^g + c^q + \phi^2 c^c} = \gamma,$$

the stated coefficients. Since $C(0,0) = 0$, $\nabla C(0,0) = \mathbf{0}$, and the Hessian is constant, C is exactly the stated quadratic, and expanding $(c^q + \phi^2 c^c)(c^g + c^q) - \phi^2 c^c c^g = c^q(c^g + c^q + \phi^2 c^c)$ gives

$$\alpha^A \alpha^L - \gamma^2 = \frac{c^g c^c c^q}{c^g + c^q + \phi^2 c^c} > 0.$$

Parts (b) and (c). The equivalence $\gamma = 0 \iff \phi = 0$ and the fixed-proportions limit follow by direct substitution in the part-(a) formulas. The comparative static is the quotient rule, using $\partial(c^g + c^q + \phi^2 c^c)/\partial c^g = 1$:

$$\frac{\partial|\gamma|}{\partial c^g} = \phi c^c \cdot \frac{(c^g + c^q + \phi^2 c^c) - c^g}{(c^g + c^q + \phi^2 c^c)^2} = \frac{\phi c^c\,(c^q + \phi^2 c^c)}{(c^g + c^q + \phi^2 c^c)^2} > 0.$$

For the general-cost extension, let the margin costs be G, Q, K, strictly convex with zero value and zero marginal cost at the origin. The Lagrangian with multipliers $(\vartheta_A, \vartheta_L)$ on the two delivery constraints gives $G'(a^g) = \vartheta_A$, $Q'(a^q) = \vartheta_A + \phi\,\vartheta_L$, $K'(a^c) = \vartheta_L$, and by the envelope theorem

$C_A = \vartheta_A$, $C_L = \vartheta_L$. Differentiating the margin-level first-order conditions, $dx = D^{-1}B^T\, d\boldsymbol{\vartheta}$ with $D = \operatorname{diag}(G'', Q'', K'')$ evaluated at the optimum and $B = \left(\begin{smallmatrix} 1 & 1 & 0 \\ 0 & \phi & 1 \end{smallmatrix}\right)$, so the delivery constraints $Bx = (a^A, a^L)^T$ give $\nabla^2 C = (BD^{-1}B^T)^{-1}$, positive definite because B has full row rank and $D \succ 0$. Direct multiplication gives

$$BD^{-1}B^T = \begin{pmatrix} \dfrac{1}{G''} + \dfrac{1}{Q''} & \dfrac{\phi}{Q''} \\ \dfrac{\phi}{Q''} & \dfrac{\phi^2}{Q''} + \dfrac{1}{K''} \end{pmatrix}, \qquad C_{AL} = -\frac{\phi\, G''\, K''}{G'' + Q'' + \phi^2 K''} \leq 0,$$

with equality iff $\phi = 0$, so the sign restriction of Assumption 1 holds for any three-margin technology, with local curvatures in place of the constant slopes. □

The four margin parameters $(\phi_i, c_i^q, c_i^c, c_i^g)$ are not identified from the three reduced-form parameters $(\alpha_i^A, \alpha_i^L, \gamma_i)$ alone, so our empirical calibration in Section 5 constructs the three abatement margins directly from facility emissions and engineering cost menus. Part (c) shows that in a smooth quadratic economy, cheaper carbon capture weakens cost complementarity $|\gamma_i|$ and reduces carbon-market co-benefits (Batres et al., 2021). On the discrete engineering menus of Section 5.5, however, joint efficiency and fuel measures are cheaper than capture and remain inframarginal as capture costs fall.

C Extensions

Cost uncertainty. Suppose realized cost parameters differ from the regulator's estimates. Regime IV's taxes and Regime V's community prices are unchanged, because neither is built from cost data. With $\tau_i^L = \mathrm{MD}_i$ fixed, the carbon price re-clears at the realized costs, so the firm first-order conditions replicate the planner's ex post conditions given the cap and both regimes are ex post efficient. Regime II fixes quantities from ex ante cost estimates and is not ex post efficient, and the optimal levels in Regimes I/III are built from cost data and therefore inherit estimation error. The community mechanism therefore mirrors the differentiated tax's cost-uncertainty robustness while strictly relaxing its damage-information requirement. Appendix E states the corresponding Fowlie and Muller (2019) results.

Optimal cap adjustment. The cap that restores $\tau^A = \mathrm{SCC}$ once community permits operate solves $R^* = \left(\sum_i \psi_i\right)\mathrm{SCC} + \sum_i p_i \mathrm{MD}_i^*$ (derived in the proof of Corollary 1). Under strategic pricing the co-abatement term is evaluated at $\mathrm{MD}_i + \Delta_i$ instead, so the required adjustment is larger by $\sum_i p_i \Delta_i$. A regulator adjusting the cap should therefore account for the extra co-abatement that markups induce or, following Corollary 1, switch the aggregate instrument to a tax.

Unabatable emissions and hold-up. The existence result in Proposition 2(a) uses the fact that each community's residual demand reaches zero at a finite price. That fails when some firm has

local emissions that no abatement option can remove, which is the case on the engineering menus of Section 5.

Proposition 9 (Hold-up without an exemption)**.** *Suppose firm i's local abatement is bounded by a technology limit $\bar{a}_i^L < \bar{e}_i^L$, leaving $\underline{e}_i^L = \bar{e}_i^L - \bar{a}_i^L > 0$ unabatable tons at every price. Let community j's receptor field contain at least one such firm, with $\underline{s}_j = \sum_i h_{ji}\,\underline{e}_i^L > 0$. (a) Under the permit requirement of Section 2, community j's residual demand satisfies $s_j(\sigma_j) \geq \underline{s}_j$ for every σ_j, its payoff in Equation (8) is unbounded above, and no best response exists. The same holds for each of n_j sub-sellers in Equation (15), because any one of them can withhold enough supply to push the clearinghouse price above any bound. The pricing game therefore has no pure-strategy equilibrium at any finite subdivision. (b) If the regulator instead allocates permits covering $\underline{e}_i^L$ to each firm without charge, communities price only the abatable exposure $s_j - \underline{s}_j$. Residual demand then reaches zero at a finite choke price and a best response exists; when residual demand is continuous, Proposition 2(a) applies, while on discrete menus existence must be verified directly. (c) The exemption leaves the planner's allocation and the price-taking equilibrium unchanged because the exempt tons cannot be abated under either. It changes only the transfer from firms to communities, by $\sum_j \sigma_j \underline{s}_j$.*

Proof. (a) Revenue satisfies $\sigma_j s_j(\sigma_j) \geq \sigma_j \underline{s}_j$ and is unbounded as $\sigma_j \to \infty$, whereas damages are bounded by $D_j(c_j^0)$. The objective in Equation (8) therefore has no maximizer. For a sub-seller, total authorized supply below $\underline{s}_j$ clears at no finite price, so the sub-seller's revenue $\sigma_j(s)\, s_{j\ell}$ is unbounded in the same way. (b) With the exempt tons allocated, the quantity firms must buy from community j is $s_j - \underline{s}_j = \sum_i h_{ji}(\bar{a}_i^L - a_i^L)$, which is zero once every firm in the receptor field abates to its limit. Because marginal abatement cost is finite at the limit, each firm reaches its limit at a finite price, and residual demand hits zero at a finite choke price. A best response therefore exists on the compact price set, and with continuous residual demand the proof of Proposition 2(a) applies with that choke price as the bound. (c) The planner's problem with limits $a_i^L \leq \bar{a}_i^L$ has the same feasible set with or without the exemption, and in the price-taking equilibrium each firm's marginal condition on the abatable tons is unchanged. Both allocations are therefore unchanged, and because the exempt tons bear no permit price, payments change by $\sum_j \sigma_j \underline{s}_j$ and nothing else. □

General-form $n_j^\star$. The optimal subdivision does not require the linear structure of Assumption 2, because its closed form generalizes to the environment of Section 2.

Proposition 10 (Optimal subdivision, general form)**.** *Maintain Assumption 1, weakly concave residual inverse demand, and interior planner supplies $s_j^\star = c_j^\star > 0$. Evaluate each community's residual demand at rivals' efficient prices $\sigma_k = D_k'(c_k^\star)$ and the efficient carbon price λ, and allow n_j real-valued. Define*

$$n_j^\star \;=\; 1 + \frac{\big(-\sigma_j'(s_j^\star)\big)\, s_j^\star}{D_j'(c_j^\star)} \;\in\; [1, \infty).$$

(a) Within each community, the symmetric profile $s_{j\ell} = s_j^\star / n_j^\star$ is the unique within-community equilibrium

of Equation (15) given the efficient environment. The planner's allocation, priced at $\sigma_j = D_j'(c_j^\star)$ for every j, is therefore an equilibrium of the joint game with subdivided communities, and both strategic distortions vanish. (b) The within-community equilibrium price is strictly decreasing in n_j, so $n_j^\star$ is the unique subdivision with this property (with monopoly pricing below $n_j^\star$ and commons pricing above it). (c) Under Assumption 2 the formula reduces to the closed form of Proposition 7.

Proof. We first solve the within-community game for a fixed number of sub-sellers. Fix community j, its environment, and $n_j = n$. Write $\sigma(\cdot)$ for its residual inverse demand and $s = \sum_\ell s_{j\ell}$ for its aggregate supply. Each sub-seller's objective in Equation (15) is strictly concave in its own offer, because the revenue term has second derivative $2\sigma'(s) + \sigma''(s)\, s_{j\ell} < 0$ when inverse demand is strictly decreasing and weakly concave, and the damage share $\frac{1}{n} D_j$ is convex. Interior best responses therefore satisfy

$$\sigma(s) + \sigma'(s)\, s_{j\ell} = \frac{D_j'(s)}{n}. \tag{73}$$

Given the aggregate s, Equation (73) determines the same offer for every ℓ, since $\sigma'(s) \neq 0$, so any interior equilibrium is symmetric. Summing Equation (73) over ℓ gives the aggregate equilibrium condition

$$\Phi(s) \equiv \sigma(s) + \frac{\sigma'(s)\, s}{n} - \frac{D_j'(s)}{n} = 0, \tag{74}$$

and

$$\Phi'(s) = \sigma'(s) + \frac{\sigma'(s) + \sigma''(s)\, s}{n} - \frac{D_j''(s)}{n} < 0 \tag{75}$$

because $\sigma' < 0$, $\sigma'' \leq 0$, and $D_j'' \geq 0$. The within-community equilibrium is therefore unique.

We next show that the efficient point solves Equation (74) at $n = n_j^\star$. At rivals' efficient prices and $\tau^A = \lambda$, the price-taking equilibrium of Proposition 1 clears community j's market at $\big(s_j^\star, D_j'(c_j^\star)\big)$, so the residual demand passes through that point and $\sigma(s_j^\star) = D_j'(c_j^\star)$. Substituting into $\Phi(s_j^\star) = 0$ and solving for n,

$$D_j'\Big(1 - \frac{1}{n}\Big) = \frac{-\sigma'(s_j^\star)\, s_j^\star}{n} \quad \Longleftrightarrow \quad n = 1 + \frac{\big(-\sigma'(s_j^\star)\big) s_j^\star}{D_j'(c_j^\star)} = n_j^\star. \tag{76}$$

At $n_j^\star$ the unique within-community equilibrium is therefore $s_j^\star$ at price $D_j'(c_j^\star)$. This holds for every community at once, because each community's environment is then the efficient one and the carbon market clears at λ by Proposition 1. The efficient profile is therefore a fixed point of the joint game. Markups are zero at that profile, so both distortions of Proposition 2 vanish, which is part (a). We do not claim joint uniqueness at this generality. Under Assumption 2 it follows from diagonal dominance, as in Proposition 7.

For part (b), we differentiate Equation (74) implicitly in n. Since $\partial\Phi/\partial n = \big(D_j'(s) - \sigma'(s)s\big)/n^2 > 0$ and $\Phi'(s) < 0$ by Equation (75), $ds/dn = -(\partial\Phi/\partial n)/\Phi'(s) > 0$. The equilibrium price $\sigma(s(n))$ is therefore strictly decreasing in n.

For part (c), under Assumption 2 we have $\sigma' = -1/\sum_i \tilde{\beta}_i h_{ji}^2$ and $D_j' = d_j$. The efficient supply is $s_j^\star = \widetilde{Q}_j - d_j \sum_i \tilde{\beta}_i h_{ji}^2$, where $\widetilde{Q}_j$ is the intercept of Equation (34) at the efficient environment.

Substituting into Equation (76) gives $n_j^\star = 1 + s_j^\star/(d_j \sum_i \tilde{\beta}_i h_{ji}^2) = \widetilde{Q}_j/(d_j \sum_i \tilde{\beta}_i h_{ji}^2)$, the closed form of Proposition 7. □

The numerator of the fraction in $n_j^\star$ is the markup a single seller would charge if it sold the efficient quantity, and the denominator is marginal damage at that quantity. The number of sub-sellers a community needs is therefore one plus its monopoly markup at the efficient point, expressed as a multiple of marginal damage. A community whose residual demand is steep relative to its damages needs more sub-sellers. Proposition 8 does not extend to this setting, because its argument relies on prices being linear in damages.

D Regime Derivations for Table 1

Each entry is Proposition 4 evaluated at the regime's wedge profile; the wedge profiles are derived here, and the loss algebra is in the proof of Proposition 5 (Appendix A).

Regimes I/III. A uniform local price t gives $\Delta_i^L = t - \mathrm{MD}_i$. Under a local cap $\bar{E}^L$ (so that required local abatement is $\sum_i \bar{e}_i^L - \bar{E}^L$), joint clearing of the two markets yields

$$t = \frac{\left(\sum_i \bar{e}_i^L - \bar{E}^L\right) - \left(\sum_i p_i\right) R / \sum_i \psi_i}{\sum_i \tilde{\beta}_i - \left(\sum_i p_i\right)^2 / \sum_i \psi_i}.$$

The welfare-optimal uniform price is the value of t that minimizes Equation (10). The first-order condition

$$\sum_i \tilde{\beta}_i (t - \mathrm{MD}_i) - \frac{\left(\sum_i p_i\right) \sum_i p_i (t - \mathrm{MD}_i)}{\sum_i \psi_i} = 0$$

is linear in t and solves to

$$t^* = \frac{\left(\sum_i \psi_i\right)\left(\sum_i \tilde{\beta}_i \mathrm{MD}_i\right) - \left(\sum_i p_i\right)\left(\sum_i p_i \mathrm{MD}_i\right)}{\left(\sum_i \psi_i\right)\left(\sum_i \tilde{\beta}_i\right) - \left(\sum_i p_i\right)^2},$$

a flexibility-weighted ($\tilde{\beta}_i$-weighted) mean of marginal damages, corrected toward the co-abatement-weighted (p_i-weighted) mean, because a uniform local price partially delegates local abatement to the carbon market and must offset that channel's average error. At $\gamma_i = 0$ the correction vanishes and $t^* = \sum_i \tilde{\beta}_i \mathrm{MD}_i / \sum_i \tilde{\beta}_i$, the Diamond (1973) weighted average recovered in Appendix E. Under beliefs, $t^*(\hat{\mathbf{d}})$ replaces MD_i with $\widehat{\mathrm{MD}}_i$ (the formula is linear). At $b = 0$ the residual loss is the dispersion term $\frac{1}{2} \sum_i \tilde{\beta}_i (t^* - \mathrm{MD}_i)^2$ net of the rebate correction, positive whenever MD_i varies across sources.

Regimes II/IV. Differentiated instruments computed from beliefs support $\tau_i^L = \widehat{\mathrm{MD}}_i$, so $\Delta_i^L = b_i$, with the equivalence of quantity and price versions holding under certainty and breaking under cost uncertainty (Appendix C).

Regime V. Price-taking: $\Delta_i^L = 0$ (Proposition 1). Strategic pricing: $\Delta_i^L = \Delta_i = \sum_j \mu_j h_{ji}$ from Lemma 1, with carbon wedge $\Delta^A = -\big(\sum_i p_i \Delta_i\big) / \sum_i \psi_i \leq 0$, strict when some $p_i > 0$; the markup-induced over-abatement raises carbon co-abatement and depresses the carbon price, and the rebate term of Equation (10) partially offsets the loss.

E Mapping to Fowlie and Muller (2019)

This appendix states the exact relation to Fowlie and Muller (2019) (henceforth FM). We show that their one-pollutant model is a limit of ours, that each of their analytical results is recovered inside our framework, and that the cost interaction $\gamma_i < 0$ is what our environment adds (Result 4). Throughout, FM objects are stated in our notation, with their symbols noted parenthetically on first appearance.

E.1 The FM environment in our notation

FM study a single non-uniformly mixed pollutant. Sources $i = 1, \ldots, I$ emit e_i, abatement costs are quadratic in emissions (their $C_i(e_i)$, with marginal-cost slope $C_i'' = 2\beta_i$), and damages are linear with a constant source-specific marginal damage (their δ_i), so total damages are $\sum_i \delta_i e_i$. The regulator compares four designs, namely a differentiated tax ($\tau_i = \delta_i$), an undifferentiated tax, undifferentiated cap-and-trade, and differentiated cap-and-trade implemented by compliance ratios (their $r_i \propto \delta_i$). Their first main analytical finding (FM-R1) is that only the differentiated tax requires no cost information. Under certainty, the gains from differentiation grow with the dispersion of marginal damages and shrink with the slope of marginal abatement costs (FM-R2, FM-R3, their eq. 11). The optimal undifferentiated tax is the Diamond (1973) weighted average of marginal damages, weighted by inverse marginal-cost slopes (their eq. 8b). Under damage uncertainty with policies committed at estimated damages, realized gains from differentiation can turn negative when the realized damage ranking reverses the expected one (their eq. 14). Under cost uncertainty (FM-R5, their eq. 15 and §2.3.3), the differentiated *tax* still weakly dominates the undifferentiated tax. Differentiated *trading*, however, can be dominated by undifferentiated trading when the cap is set on mistaken cost estimates, because compliance ratios fix relative allocations that are optimal only at the intended cap.

Two structural differences distinguish our setup. First we model two pollutants, a uniformly mixed aggregate pollutant and a non-uniformly mixed local one, whereas FM model a single non-uniformly mixed pollutant whose marginal damage δ_i bundles all sources of spatial variation. Second, our pollutants are linked on the cost side by $\gamma_i \leq 0$; FM have a single abatement margin, so no such link can arise. The nesting results below establish how these differences disappear in the limit.

E.2 Nesting

Result 1 (Nesting). *(a)* Separable limit. *At $\gamma_i = 0$ for all i, we have $p_i = 0$, the carbon market decouples from the local layer, and the local layer of our model is exactly FM's model with the identifications*

$$e_i \leftrightarrow e_i^L, \qquad \delta_i \leftrightarrow \mathrm{MD}_i = \sum_j d_j h_{ji}, \qquad 2\beta_i \leftrightarrow \alpha_i^L,$$

and with our **H** *supplying the microstructure of their δ_i. (b)* Fixed-proportions limit. *When NO_x-only abatement and CO_2-only capture shut down ($c_i^c, c_i^g \to \infty$ in Appendix B, so $a_i^L = \phi_i a_i^A$), our model collapses to a single emission whose marginal damage at source i is $\mathrm{SCC} + \phi_i \mathrm{MD}_i$. FM's non-uniform marginal damage then decomposes into a spatially uniform component and a spatially varying one. Given a binding cap, the uniform component contributes the same constant to every regime's damages and cancels from all welfare comparisons, so all heterogeneity relevant to the efficient spatial allocation comes from $\phi_i \mathrm{MD}_i$, our local component.*

Proof. (a) At $\gamma_i = 0$, Equation (9) gives $\psi_i = 1/\alpha_i^A$, $p_i = 0$, $\tilde{\beta}_i = 1/\alpha_i^L$, and Equation (23) decouples into $a_i^A = \tau^A/\alpha_i^A$ and $a_i^L = \tau_i^L/\alpha_i^L$. The local layer's firm problem, market structure, and damages then coincide term by term with FM's model under the stated identifications. Quadratic local abatement cost $\frac{\alpha_i^L}{2}(a_i^L)^2$ has marginal-cost slope $\alpha_i^L = 2\beta_i$, and total local damages are $\sum_j d_j c_j = \text{const} - \sum_i \mathrm{MD}_i a_i^L$ (see the proof of Proposition 4), matching their $\sum_i \delta_i e_i$ with $\delta_i = \mathrm{MD}_i$. (b) In the fixed-proportions limit, one unit of activity abatement removes one unit of the aggregate pollutant and ϕ_i units of the local one, so its marginal social value is $\mathrm{SCC} + \phi_i \mathrm{MD}_i$ (Lemma 2(b)). With the cap binding, aggregate emissions are $\bar{E}^A$ in every regime compared, so the uniform component contributes the same constant $\mathrm{SCC} \cdot \bar{E}^A$ to damages in every regime and differences out, leaving $\phi_i \mathrm{MD}_i$ as the operative source-specific marginal damage. □

E.3 Recovery of FM results under certainty

Result 2 (FM under certainty, recovered at $\gamma = 0$). *At $\gamma_i = 0$, with the identifications of Result 1(a): (a) the differentiated tax $\tau_i^L = \mathrm{MD}_i$ attains the first best (their eq. 8a), and is our Regime IV at $b_i = 0$; (b) the optimal uniform price is the inverse-cost-slope-weighted average*

$$t^*\big|_{\gamma=0} = \frac{\sum_i \mathrm{MD}_i/\alpha_i^L}{\sum_i 1/\alpha_i^L},$$

their eq. 8b and the Diamond (1973) formula; with two sources it is $(\beta_H \delta_L + \beta_L \delta_H)/(\beta_H + \beta_L)$ in their notation; (c) the gains from differentiation are the flexibility-weighted dispersion of marginal damages about t^*,

$$\mathcal{L}^{\text{undiff}} - \mathcal{L}^{\text{diff}}\Big|_{b=0,\,\gamma=0} = \frac{1}{2}\sum_i \frac{\left(t^* - \mathrm{MD}_i\right)^2}{\alpha_i^L},$$

which with two sources equals $(\delta_L - \delta_H)^2/\left[4(\beta_L + \beta_H)\right]$, their eq. 11; their comparative statics FM-R2 (increasing in damage dispersion) and FM-R3 (decreasing in cost slopes) follow directly; (d) aggregate local

emissions coincide across all four optimally parameterized designs, their emissions-equivalence result; (e) differentiated cap-and-trade with compliance ratios $r_i \propto \mathrm{MD}_i$ and an optimal cap supports the same source prices as the differentiated tax (their eq. 7), matching our Regime II. (f) Among Regimes I–IV, only the differentiated tax requires no cost data (FM-R1), as noted in Table 1 and Appendix C.

Proof. (a) Immediate from Table 1, where at $b_i = 0$ Regime IV's wedge and loss are zero. (b) Set $\gamma_i = 0$ in the t^* formula of Appendix D, where $p_i = 0$ removes the correction term and $\tilde{\beta}_i = 1/\alpha_i^L$, giving the display; with two sources and their $2\beta_i = \alpha_i^L$, multiply numerator and denominator by $2\beta_L\beta_H$ to get $(\beta_H\delta_L + \beta_L\delta_H)/(\beta_H + \beta_L)$. (c) At $\gamma = 0$ the rebate term of Equation (10) vanishes, so $\mathcal{L}^{\text{undiff}} = \frac{1}{2}\sum_i (t^* - \mathrm{MD}_i)^2/\alpha_i^L$ while $\mathcal{L}^{\text{diff}} = 0$ at $b = 0$. With two sources, $t^* - \delta_L = \beta_L(\delta_H - \delta_L)/(\beta_L + \beta_H)$ and $t^* - \delta_H = \beta_H(\delta_L - \delta_H)/(\beta_L + \beta_H)$, so

$$\frac{1}{2}\left[\frac{\left(t^* - \delta_L\right)^2}{2\beta_L} + \frac{\left(t^* - \delta_H\right)^2}{2\beta_H}\right] = \frac{(\delta_H - \delta_L)^2}{4(\beta_L + \beta_H)^2}\left(\beta_L + \beta_H\right) = \frac{(\delta_L - \delta_H)^2}{4(\beta_L + \beta_H)},$$

their eq. 11; it is increasing in $|\delta_L - \delta_H|$ and decreasing in $\beta_L + \beta_H$. (d) At $\gamma = 0$, aggregate local abatement under a uniform price t is $t\sum_i 1/\alpha_i^L$ and under the differentiated prices MD_i is $\sum_i \mathrm{MD}_i/\alpha_i^L$; these coincide at $t = t^*$ by (b). The quantity designs support the same prices by construction of their optimal totals. (e) Under compliance ratios r_i, source i's cost minimization equates marginal abatement cost to $r_i\tau$ at the single permit price τ. With $r_i = \mathrm{MD}_i/\bar{\delta}$ for any normalization $\bar{\delta} > 0$, the supported prices are $\tau_i^L = \tau\mathrm{MD}_i/\bar{\delta}$. Choosing the cap so that $\tau = \bar{\delta}$ gives $\tau_i^L = \mathrm{MD}_i$, matching the differentiated-tax prices and allocation. (f) The differentiated tax posts $\widehat{\mathrm{MD}}_i$ from $\mathbf{H}$ and beliefs alone. The uniform designs require t^* or its supporting cap (weighting damages by cost slopes), and the quantity designs require cost data to set totals. □

E.4 Recovery of FM results under uncertainty

Result 3 (FM under uncertainty, recovered at $\gamma = 0$)**.** *At $\gamma_i = 0$, with policies set at the estimated means m_j (so sources face $\widehat{\mathrm{MD}}_i = \sum_j m_j h_{ji}$ under differentiation and $t^*(\mathbf{m})$ under uniformity): (a) for any damage realization $\mathbf{d}$, the realized gain from differentiation is*

$$\mathcal{L}^{\text{undiff}}(\mathbf{d}) - \mathcal{L}^{\text{diff}}(\mathbf{d}) = \frac{1}{2}\sum_i \frac{\left(t^*(\mathbf{m}) - \mathrm{MD}_i\right)^2 - \left(\widehat{\mathrm{MD}}_i - \mathrm{MD}_i\right)^2}{\alpha_i^L},$$

which with two sources reduces to their eq. 14; it is positive when damages realize as expected and can turn negative when the realized ranking of source damages reverses the expected one, so differentiation can reduce welfare under damage uncertainty. (b) Under cost uncertainty with known damages, the differentiated tax remains ex post optimal for any realized costs because its rates require no cost data. Its gain over the undifferentiated tax is therefore nonnegative for every realization (their eq. 15 and FM-R5). Under quantity-based designs, a cap set on mistaken cost estimates no longer supports the intended prices, and the compliance ratios of Result 2(e) fix relative source prices proportional to MD_i even when the cap-constrained optimum calls for a different allocation. Differentiated trading can therefore be welfare-dominated by undifferentiated

trading (their §2.3.3 and Appendix C).

Proof. (a) At $\gamma = 0$ the rebate term vanishes, and evaluating Equation (10) at each regime's wedge profile ($\Delta_i^L = t^*(\mathbf{m}) - \mathrm{MD}_i$ for I/III and $\Delta_i^L = \widehat{\mathrm{MD}}_i - \mathrm{MD}_i = b_i$ for II/IV) gives the display. For the two-source case, write $\hat{\delta}_i = \widehat{\mathrm{MD}}_i$ and $\delta_i' = \mathrm{MD}_i$ in their notation ($\alpha_i^L = 2\beta_i$), and expand both squares around the estimated means. With $t^*(\mathbf{m}) = (\beta_H \hat{\delta}_L + \beta_L \hat{\delta}_H)/(\beta_L + \beta_H)$,

$$\frac{1}{2} \sum_{i \in \{L,H\}} \frac{\left(t^*(\mathbf{m}) - \delta_i'\right)^2 - \left(\hat{\delta}_i - \delta_i'\right)^2}{2\beta_i} = \frac{1}{2} \frac{\left(\delta_L' - \delta_H'\right)\left(\hat{\delta}_L - \hat{\delta}_H\right)}{\beta_L + \beta_H} - \frac{1}{4} \frac{\left(\hat{\delta}_L - \hat{\delta}_H\right)^2}{\beta_L + \beta_H},$$

their eq. 14 (Table G1). The expression is positive at $\mathbf{d} = \mathbf{m}$ (it reduces to their eq. 11), increases in the realized spread when the ranking is as expected, and turns negative when $(\delta_L' - \delta_H')$ reverses sign relative to $(\hat{\delta}_L - \hat{\delta}_H)$ strongly enough. (b) The differentiated tax rates $\widehat{\mathrm{MD}}_i$ involve no cost parameters, so for any realized costs the induced allocation solves the planner's problem at believed marginal damages. With known damages ($\widehat{\mathrm{MD}}_i = \mathrm{MD}_i$) the differentiated tax is ex post first-best with zero loss at every cost realization, whereas the undifferentiated tax has nonnegative loss (FM-R5). For the quantity designs, realized-cost market clearing under a cap fixed at the ex ante optimum yields a price level different from the intended one. Under differentiated trading, the compliance ratios continue to force $\tau_i^L \propto \mathrm{MD}_i$ at whatever level clears the mis-set cap, whereas the welfare-optimal allocation *conditional on that cap* equates a different pattern of marginal costs. Undifferentiated trading equates marginal costs across sources and can therefore lie closer to the cap-constrained optimum, as FM document empirically. □

E.5 Cost-interaction term

Result 4 (Cost-interaction channels absent in FM)**.** *Relative to the $\gamma = 0$ benchmark that reproduces FM, cost complementarity adds three channels that vanish continuously as $\gamma_i \to 0$. (a) Under* the carbon-market rebate*, every regime's loss includes the correction $-\left(\sum_i p_i \Delta_i^L\right)^2 / \left(2 \sum_i \psi_i\right)$ of Equation (10). Uniform-instrument losses are therefore no longer a pure weighted dispersion of marginal damages, and t^* is shifted toward the co-abatement-weighted mean (Appendix D). (b) Through* an endogenous aggregate price*, local policy moves the carbon price (Proposition 3), a channel with no counterpart in FM's single market. (c) Through* the cap-level channel*, the optimal cap requires a damage forecast, so regulator bias enters through the cap in an otherwise robust design (Corollary 1). Consequently, every entry of Table 1 converges to its FM counterpart as $\gamma_i \to 0$ for all i.*

Proof. Each channel is proportional to $\{p_i\}$, since the rebate term and the t^* correction carry p_i factors explicitly, the interaction derivative of Proposition 3 has numerator $\sum_i p_i h_{ji}$, and the cap-level loss of Corollary 1 is $\left(\sum_i p_i b_i\right)^2 / \left(2 \sum_i \psi_i\right)$. Since $p_i = -\gamma_i / (\alpha_i^A \alpha_i^L - \gamma_i^2) \to 0$ and $\psi_i, \tilde{\beta}_i$ converge to $1/\alpha_i^A, 1/\alpha_i^L$ as $\gamma_i \to 0$, all three channels vanish continuously, and the loss formulas of Table 1 converge to the $\gamma = 0$ expressions recovered in Results 2 and 3 (Table G1). □

F Calibration Details

Measurement

Transport. We construct $\mathbf{H}$ from facility-specific InMAP runs (Tessum et al., 2017). For each source we perturb NO_x emissions by 1,000 short tons released at ground level, hold other precursors fixed, and record the change in ambient $PM_{2.5}$ at each of 9,109 California census-tract receptors per ton emitted. For each tract, we compute the mean over the InMAP grid cells that intersect it, which agrees with the centroid rule of AP4 within 1 percent at the facility level. We count California receptors only; the median facility has 86 percent of its population exposure in state (59 to 97 percent between the 10th and 90th percentiles). We measure GHG quantities in metric tons and criteria pollutants in short tons, following the reporting conventions of the underlying data, and we report every dollar figure in 2019 dollars.

Damages. We estimate tract damages with the mortality convention of AP3 and AP4 (Muller and Mendelsohn, 2009). Deaths at risk are adults aged 30 and over and infants, computed from each tract's 2020 Census population and county death rates by nineteen age groups. The concentration-response slope is $\ln(1.06)/10$ per $\mu g/m^3$ from Krewski et al. (2009), applied linearly with no threshold. The value of a statistical life is \$10.29 million in 2019 dollars, EPA's \$7.4 million in 2006 dollars carried forward with the consumer price index and with real income per capita at an elasticity of 0.7. Table 5 in Section 5.5 reports results at the 2.5th and 97.5th percentiles of the AP4 VSL distribution. It also reports results using the adult-mortality slope of Di et al. (2017), $\ln(1.073)/10$, which is 21 percent steeper than the Krewski slope and applies to the population aged 65 and over. Tracts with zero 2020 population are assigned zero damage; there are thirty-two in the state and twelve on the transport grid.

Co-emission ratios. We measure ϕ_i as the facility's 2019 NO_x from CARB's criteria-pollutant inventory divided by its 2019 covered greenhouse-gas emissions from mandatory reporting, in short tons of NO_x per thousand metric tons of CO_2e.[29] Eight facilities never report NO_x; we keep them in the carbon market with zero local emissions and zero co-abatement. Because the 2019 NO_x vector predates the South Coast district's BARCT rules adopted between 2018 and 2021, the abatement frontier below is conditioned on the controls facilities had in 2019.

NO_x-only abatement costs. We build a discrete frontier from the end-of-pipe control measures in EPA's Control Strategy Tool and SCAQMD's BARCT technology assessments (U.S. Environmental Protection Agency, 2020; South Coast Air Quality Management District, 2021). For each facility we identify the controls it operated in 2019 from its CO_2 ratio, the staff reports, and the applicable rules. We then keep the deepest listed control beyond its 2019 equipment and price it over the incremental tons removed. The frontier has at least one option for 179 of the 270 facilities with NO_x emissions (one step for 98, two for 64 and ladders of eight for 17) and none for the other 91. Those 91 facilities already operate the deepest control listed for their equipment, with a median

[29] The ratio is capped at the uncontrolled emission factor for the sector's equipment in EPA's AP-42 compilation, but no facility in the sample reaches that cap.

removal of 97 percent. Their remaining 2,226 tons (7 percent of the sample) are entirely exempt. They abate nothing in any regime and buy no community permits. The frontier is capable of removing at most 63 percent of the sample's NO_x (7 percent at hydrogen plants to 79 percent at cement kilns). Costs are in 2019 dollars by the chemical engineering plant cost index and are levelized at the district's 4 percent over 25 years. The staff reports give a range of cost cells for cement kilns, and we use the midpoint cell.

Engineering menus of this kind understate the compliance costs observed in permit markets (Fowlie, 2010; Fowlie and Muller, 2019). Fowlie and Muller (2019) set the scale of their menu by iterating on the permit price until the engineering model reproduced the observed cap, which requires a baseline from which the observed price induced a known amount of abatement. Our baseline is 2019 emissions chosen while facilities already faced the RECLAIM price. Because the abatement induced by that price is already in the baseline, the same scaling route is unavailable here (Fowlie et al., 2012).

We therefore report the engineering costs alongside frontiers scaled by 1.845 and 2.393. Those two multipliers equal the ratios by which the compliance costs Fowlie and Muller (2019) recovered in the NO_x Budget Program exceeded their engineering menu at the optimal and observed caps. Those ratios are borrowed from a different program. As corroboration, the nine RECLAIM credit trades after the 2015 reduction in the program's cap span 0.58 to 2.75 of the frontier's marginal cost at the traded quantities. Table 4 in Section 5.5 reports results under the engineering costs, the two multipliers and the kiln cost cells.

Joint CO_2–NO_x abatement costs. The carbon side entails a per-facility menu of 2,672 facility-option rows. For electricity generators the options are heat-rate improvements and unit replacement priced from each plant's EIA-923 heat rate relative to published benchmarks. Generators can also exhibit a dispatch response that we take from the demand elasticity of Borenstein et al. (2019) at -0.2, with -0.1 as the other edge of the band. We scale this by 0.461 for the share of marginal load served by in-state covered plants and normalize to the 2019 retail price and sales. This dispatch response is allocated across generators in proportion to their CO_2. For refineries and hydrogen plants the options are CARB's documented measures, with pressure-swing adsorption capped at 1.1 million tons across the sector. For cement, oil and gas, cogeneration and other combustion they are sector cost curves from engineering studies allocated to facilities by emissions. Each option records its cost per ton in 2019 dollars, its share of the facility's CO_2, and an option-specific NO_x factor (zero for alternative fuels at cement kilns and one for dispatch reductions at in-state plants). We price the menu in three ways. The headline uses each measure's engineering cost net of the energy it saves, floored at zero. We bound this approach with two alternatives. The first is the gross outlay with no credit for savings. The other drops every measure already profitable at \$16.84, on the view that firms would have adopted those measures without a policy. Every cost is shifted by \$16.84 so that marginal abatement cost at 2019 emissions equals the 2019 allowance price. The menu has no output-reduction options outside electricity and cement, and electricity's sourced stock is small (0.7 of the 6.8 million tons removed at \$44.28) because gas-unit

heat-rate costs are not well documented.

CO_2-only capture costs (c_i^g). We treat capture as unavailable in the 2019 baseline, matching the absence of CCS from the compliance record. In the capture counterfactual of Appendix G, we use sector-specific retrofit costs from National Energy Technology Laboratory (2022) adjusted to 2019 dollars. These equal \$62 per ton for hydrogen and refineries (based on the hydrogen stream), \$64 for cement and other combustion, \$16 for oil and gas (based on gas processing), and \$65 for electricity and cogeneration.

Solution and conventions

We compute every equilibrium in three nested steps. First, given a carbon price and a vector of community prices, each firm adopts every menu option whose cost per ton (including the \$16.84 intercept) is at or below the price it faces. Each NO_x control removes its rated share of the NO_x remaining after the facility's joint abatement, and a firm indifferent at a step keeps its permits. Second, given the community prices, we find the carbon price at which aggregate carbon abatement equals the cap by bisection on the menu's breakpoints, or hold the price at \$44.28 under the exogenous-price framing. Third, in the strategic-pricing regimes, each community chooses its price given the other communities' prices and the carbon price. Its residual demand is a step function, so we find its best response by enumerating the breakpoints of that function with the carbon market re-cleared at each candidate. In the subdivided regimes we enumerate, for each community given the others' prices, the symmetric profiles of its sub-sellers that no sub-seller can improve on, and select the one nearest the previous round's price. We iterate across communities, re-solving the first two steps after each round, until no price changes or the iteration cycles, and we then test every seller in the end state for a profitable unilateral deviation over the full step set. A profile is reported as an equilibrium only if that check passes, whereas a profile that fails it is reported as unverified and excluded.

The diagonal-dominance condition in Equation (40) fails at every community definition, because atmospheric transport makes every receptor field overlap with every other, and on the discrete menus the pricing game has several equilibria. We start the iteration from the price-taking profile and from four random profiles for every regime and report the set of verified equilibria, with the lowest, the highest and the price-taking-start welfare loss. Table G2 in Appendix G reports, for each regime, the number of starts, the number that verified and the number of distinct equilibria found. Without the exemption of Proposition 9, no start verifies in any regime, because residual demand never reaches zero.

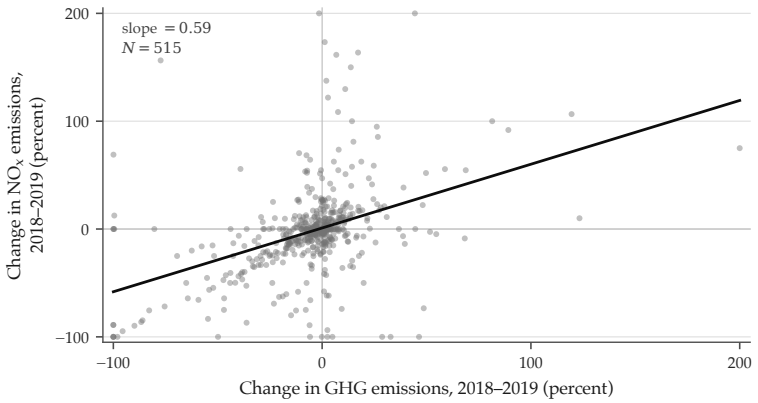
We use consistent conventions throughout. We hold the allowance price at \$44.28 per ton in the headline regimes and, in the robustness runs, fix the cap at the 6.77 million tons that clears at that price without community permits and let the price re-clear. We verify in every regime that the clearing price stays above the 2019 reserve price. We exempt from the community permit requirement the NO_x remaining after the South Coast district's BARCT controls for each facility's equipment. Allocating those exempt tons to facilities without charge ensures that communities

price only the abatable tons. Facilities that report no NO_x remain in the carbon market with no local emissions. In the uncertainty experiments, we draw each community's estimated damage as its true damage times a lognormal error with mean one, independent across communities, whose standard deviation is the coefficient of variation reported with each exhibit. We also report a variant in which part of the error is common to all communities.

G Additional Calibration Exhibits

Figure G1 supports the co-abatement assumption of Section 5, and Figure G6 compares the two taxes' expected losses under estimation error across community definitions. Figure G2 shows why the integer threshold of Proposition 7 is degenerate on the engineering menus and how the linear-quadratic formula distributes across counties. Figure G3 plots the uniform-subdivision loss across $\bar{n}$, and Figure G4 shows one county's residual demand with and without the exemption of Proposition 9. Table G1 reports numerical checks of the closed forms in Appendices A, B, and E, and Table G3 repeats Table 3 under the cap framing.

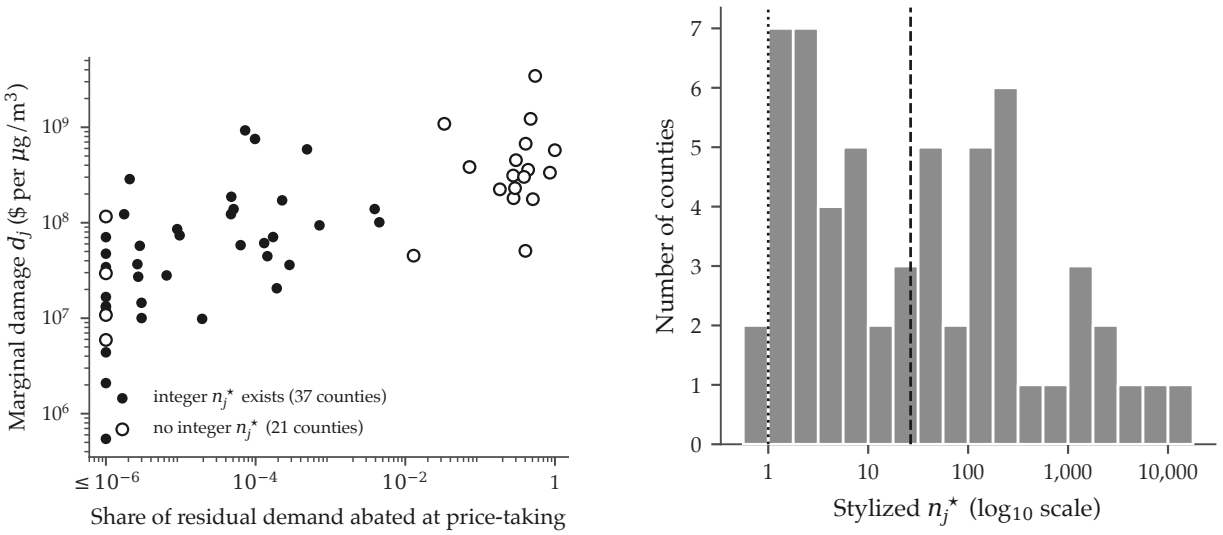
Figure G1. Changes in GHG and NO_x emissions, 2018–2019



Notes: Each point is a facility covered by the cap-and-trade program, with the percentage change in its total GHG emissions and in its NO_x emissions between 2018 and 2019 from CARB's facility emissions data. Facilities whose change in either series exceeds 200 percent are omitted, leaving 515. The line is the least-squares fit, with slope 0.59.

Each row of Table G1 reports the largest error observed when one closed form or identity from the appendices is recomputed on a random test economy by a method that does not use the closed form. Table G2 reports the number of starting profiles, the number passing the unilateral-deviation check, and the number of distinct verified equilibria and their welfare-loss range for each strategic regime. For the unitary game, where no start passed, it reports the end states after 200 rounds with the largest gain any county could still obtain by deviating.

Figure G2. The subdivision threshold $n_j^\star$ across counties

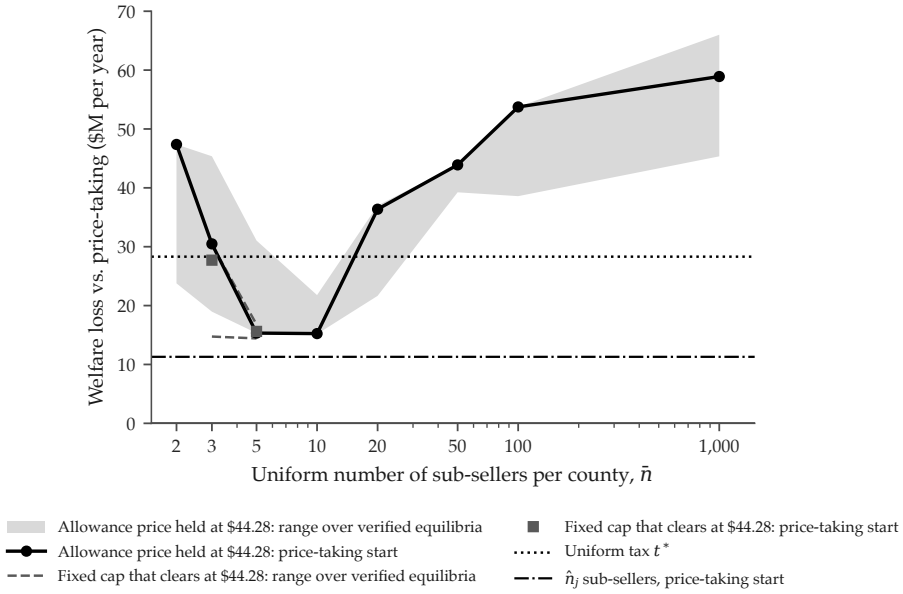


(a) Exact game on the menus

(b) Linear-quadratic economy

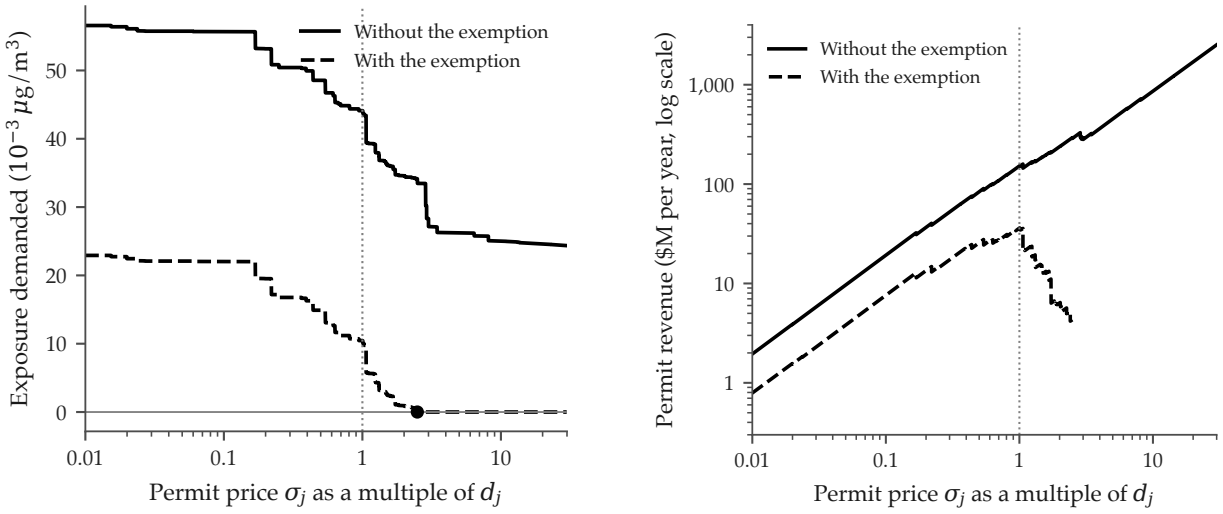
Notes: Panel (a) plots each county's first-best abatement share of net residual demand (log scale). Markers indicate whether an integer subdivision supports marginal-damage pricing as an equilibrium of the exact game on the engineering menus. The threshold exists only in the counties where that share is below one half of one percent. Panel (b) is the histogram of $n_j^\star$ from Proposition 7 in the linear-quadratic economy calibrated to the same menus, on a $\log_{10}$ scale. The dashed line marks the median, 26.6 under the cap framing shown and 26.1 under the price framing.

Figure G3. Welfare loss under a uniform subdivision



Notes: The shaded band spans the verified equilibria of the subdivided game at each $\bar{n}$ (county unit, allowance price $44.28, five starting profiles), and the solid curve marks the equilibrium reached from the price-taking profile. The dashed band at $\bar{n} \in \{3, 5\}$ shows the corresponding range under the fixed-cap framing. Horizontal reference lines mark the full-information uniform tax ($28.3M), the estimated subdivision $\hat{n}_j$ from the price-taking start ($11.3M), and the constrained first best (0), all in $M per year relative to the constrained first best.

Figure G4. Residual demand with and without the exemption of unabatable tons

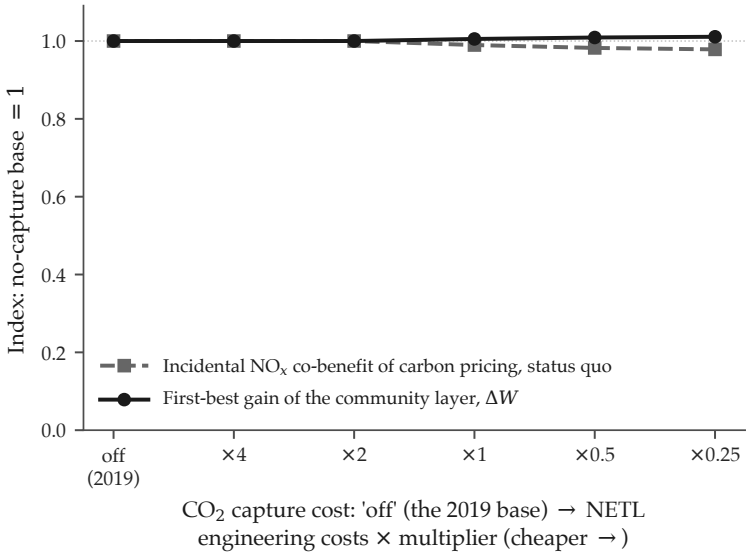


(a) Exposure demanded

(b) Permit revenue

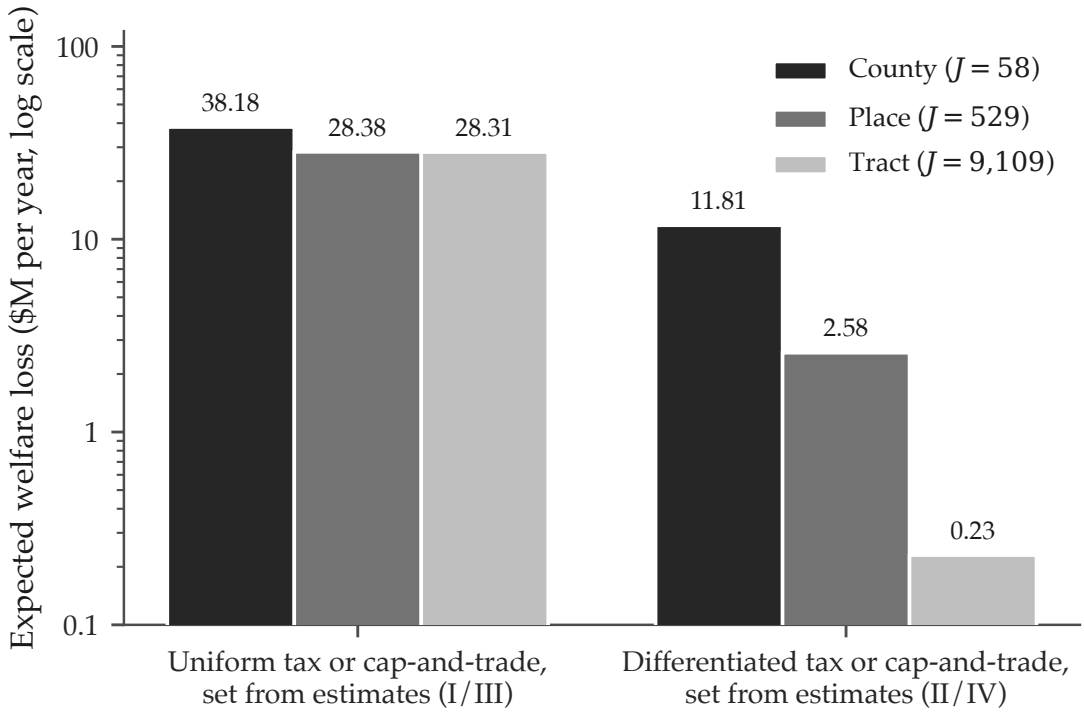
Notes: Los Angeles County's residual demand for exposure permits as a function of its posted price, as a multiple of its marginal damage, with rivals at marginal damage and the allowance price at $44.28. Without the exemption, demand has a positive floor at the exposure that no listed control can remove, revenue rises without bound in the price and no best response exists (Proposition 9). With the exempt tons allocated to facilities, community permit demand equals exposure above the exempt allocation and reaches zero at a finite choke price. Beyond that price, facilities abate below their allocation and hold unused exempt permits.

Figure G5. The capture counterfactual



Notes: Both series are indexed to the base in which capture is unavailable. Moving right, the capture cost falls from four times the NETL engineering estimates (National Energy Technology Laboratory, 2022) to one quarter of them, with the allowance price held at $44.28. The co-benefit series is the incidental NO_x reduction from carbon pricing alone, and ΔW is the first-best gain of the community permit layer over the carbon-only status quo. Under a fixed cap the co-benefit falls by 2 percent at the lowest cost as the clearing price falls to $21.

Figure G6. Expected Welfare Losses of the Two Taxes at Different Community Definitions



Notes: Markers report expected welfare losses relative to the constrained first best (log scale, $M per year) at CV $= 1.0$ over 200 independent lognormal draws for three community definitions, with the allowance price held at $44.28. When half or four-fifths of the error's log-variance is common across communities ($\rho = 0.5$ or 0.8), the differentiated tax loss rises to $9.5 million and $13.0 million at places and $9.0 million and $12.7 million at tracts (Section 5). Regime V price-taking is exactly zero at every definition and therefore does not appear on the log axis. The strategic regimes are computed for counties only.

Table G1. Numerical checks of the closed forms

Check	Where in the paper	Test economy	Tolerance	Max. abs. error
Appendix A: propositions and lemmas (Sections 2–4)				
Price-taking equilibrium equals the planner's constrained optimum (BFGS on the margin vector)	Prop. 1	random 12-firm, 4-community	10^{-4}	1.1×10^{-7}
Diagonal-dominance bound, Equation (40), holds on the test economy	Lemma 1	random 12-firm, 4-community	pass/fail	passed
Strategic-pricing linear system $\mathbf{G}\boldsymbol{\sigma} = \mathbf{g}$ equals best-response iteration from four starts	Lemma 1; Prop. 2	random 12-firm, 4-community	10^{-7}	1.8×10^{-15}
Direct Nash solve with carbon feedback satisfies the strategic-pricing first-order conditions, $n \in \{1, 3\}$	Lemma 1; Prop. 2	random 20-firm, 15-community	10^{-6}	1.2×10^{-10}
Strategic-pricing supply is optimal against residual demand; wedges and loss match direct welfare	Prop. 2	random 12-firm, 4-community	10^{-6}	1.4×10^{-14}
Strategic-pricing welfare loss equals the loss identity at the strategic-pricing wedge profile (relative)	Prop. 2; Prop. 4	random 20-firm, 15-community	10^{-8}	7.7×10^{-16}
Carbon market clears at the closed-form τ^A; layers decouple at $\phi = 0$	Prop. 3	random 12-firm, 4-community	10^{-8}	0
Carbon market clears at the closed-form τ^A (relative clearing residual)	Prop. 3; Equation (44)	random 20-firm, 15-community	10^{-6}	1.5×10^{-16}
At $\gamma = 0$ the carbon price is invariant to local prices	Prop. 3	random 20-firm, 15-community	10^{-9}	0
Loss identity, Equation (10), versus brute-force welfare difference (20 random price profiles)	Prop. 4	random 12-firm, 4-community	10^{-8}	6.1×10^{-15}
Loss identity, Equation (10), versus brute-force welfare difference (relative; both γ values)	Prop. 4	random 20-firm, 15-community	10^{-9}	2.8×10^{-14}
Carbon-wedge linkage $\Delta^A = -\sum_i p_i \Delta_i^L / \sum_i \psi_i$; loss nonnegative on 50 random profiles	Prop. 4	random 12-firm, 4-community	10^{-8}	1.4×10^{-16}
Loss nonnegative on 40 random profiles (most negative value found)	Prop. 4	random 20-firm, 15-community	10^{-12}	0
Level + allocation − rebate decomposition equals the loss, Equation (44)	Prop. 4; Equation (44)	random 12-firm, 4-community	10^{-8}	8.9×10^{-16}
Optimal number of sub-sellers $n^\star$ delivers $\sigma_j = d_j$ at the joint fixed point	Prop. 7	random 12-firm, 4-community	10^{-7}	2.2×10^{-16}
Optimal number of sub-sellers $n^\star$ delivers $\sigma_j = d_j$ (relative to mean d_j)	Prop. 7	random 20-firm, 15-community	10^{-6}	5.3×10^{-16}
Affine price map under committed $\hat{n}$ equals the simulated equilibrium; residual identity; $\theta_j = 1/(\hat{n}_j + 1)$	Prop. 8(a); Equation (19)	random 12-firm, 4-community	10^{-6}	3.5×10^{-15}

continued on next page

Table G1, continued

Check	Where in the paper	Test economy	Tolerance	Max. abs. error
Scalar-limit blend Equation (19) predicts realized prices (median relative error)	Prop. 8(a)	random 20-firm, 15-community	5×10^{-2}	6.5×10^{-5}
Trace forms of expected losses (Monte Carlo, 200,000 draws); diagonal-Σ dominance; $(1-\bar{\theta})^2$ ratio; counterexample	Prop. 6; Prop. 8(b)–(c)	random 12-firm, 4-community	2×10^{-2}	2.3×10^{-3}
Trace-form expected losses Equation (11), Equation (20) versus direct evaluation (relative)	Prop. 6; Prop. 8(b)	random 20-firm, 15-community	10^{-9}	0
Uniform-θ ratio $(1-\bar{\theta})^2$ (relative)	Prop. 8(c)	random 20-firm, 15-community	10^{-9}	1.7×10^{-16}
Uniform-subdivision price $\sigma_j(\bar{n}) = d_j(n_j^\star(\bar{n})+1)/(\bar{n}+1)$, $\bar{n} \in \{2,5,20\}$ (relative)	Cor. 2; Equation (21)	block-diagonal 12-firm, 6-community	10^{-10}	3.8×10^{-16}
Appendix B: microfoundation				
Induced cost equals direct minimization over the three margins	Lemma 2(a)	random 12-firm, 4-community	10^{-7}	4.4×10^{-16}
Inverse Hessian $= [\psi\ p;\ p\ \tilde{\beta}]$ and $\det = c^g c^c c^q / K$	Lemma 2(a); Equation (9)	random 12-firm, 4-community	10^{-8}	1.8×10^{-15}
Composites $(\psi_i, p_i, \tilde{\beta}_i)$ equal the entries of $(\nabla^2 A_i)^{-1}$ (relative)	Equation (9); Equation (22)	random 20-firm, 15-community	10^{-12}	1.7×10^{-16}
Margin-level first-order conditions equal the induced-cost firm response	Equation (23); Lemma 2	random 12-firm, 4-community	10^{-8}	4.4×10^{-16}
Firm responses satisfy $\nabla A_i = (\tau^A, \tau_i^L)'$	Equation (23)	random 20-firm, 15-community	10^{-9}	1.7×10^{-13}
$\lvert\gamma_i\rvert$ strictly increasing in c_i^g (200-point grid, every firm)	Lemma 2(c)	random 12-firm, 4-community	pass/fail	passed
General-cost Hessian $(BD^{-1}B')^{-1}$ and $C_{AL} \leq 0$ at quartic margin costs (finite differences)	Lemma 2(c), general-cost extension	random 12-firm, 4-community	10^{-3}	6.4×10^{-8}
Any $\gamma \leq 0$ quadratic is generated by some three-margin technology (round trip)	Lemma 2, non-identification remark	random 12-firm, 4-community	10^{-9}	2.2×10^{-15}
Appendix D: regime derivations				
Regime I: both markets clear at the closed-form (t, τ^A)	Appendix D, Regimes I/III	random 12-firm, 4-community	10^{-8}	0
Optimal uniform price t^* equals the numerical minimizer of the loss	Appendix D, t^*	random 12-firm, 4-community	10^{-6}	0
Optimal uniform price t^* versus an 8,001-point grid minimizer (tolerance = grid step)	Appendix D, t^*	random 20-firm, 15-community	1.9×10^{-2}	8.3×10^{-3}
At $\gamma = 0$, t^* equals the $\tilde{\beta}$-weighted mean of MD_i (relative)	Appendix D; Result 2(b)	random 20-firm, 15-community	10^{-10}	0
Appendix E: mapping to Fowlie and Muller (2019)				

continued on next page

Table G1, continued

Check	Where in the paper	Test economy	Tolerance	Max. abs. error
Diamond t^* at $\gamma = 0$; two-source eq. 8b and eq. 11 of FM	Result 2(b)–(c)	random 12-firm, 4-community	10^{-8}	8.7×10^{-19}
Emissions equivalence; compliance ratios $r_i \propto \mathrm{MD}_i$ support the differentiated-tax prices	Result 2(d)–(e)	random 12-firm, 4-community	10^{-8}	4.4×10^{-16}
FM eq. 14: realized gains from differentiation, two sources (50 draws)	Result 3(a)	random 12-firm, 4-community	10^{-8}	1.7×10^{-16}
$\gamma \to 0$: every Table 1 entry converges to its FM counterpart	Result 1; Result 4	random 12-firm, 4-community	10^{-6}	5.7×10^{-7}

Notes: Each row re-derives one closed form or identity numerically and reports the largest absolute error observed (relative error where the row says so). Rows on the 12-firm, 4-community economy are from an independent base-R implementation of the model (23 checks, T0–T21, all pass); rows on the 20-firm, 15-community and the block-diagonal economies are from the Python solver's test suite (39 tests, all pass), which runs every identity both at $\gamma = 0$ and at strong complementarity ($\gamma_i = -0.4\sqrt{\alpha_i^A \alpha_i^L}$). Where a check is an inequality or a monotonicity statement the tolerance column reads "pass/fail". The random economies draw exposure matrices with one dominant receptor per firm and small spillovers, so the diagonal-dominance condition of Lemma 1 holds by construction (verified in the second row). Table G2 reports the equilibria found for each strategic regime on the calibrated economy.

Our main calibration uses the carbon-side menu of engineering estimates of abatement costs net of energy savings and holds the allowance price fixed at \$44.28 per ton. Table G4 reports the results under the other readings of the menu and at other carbon prices. Pricing every measure at its gross outlay (with no credit for energy savings) raises the first-best gain from \$66.5 million to \$75.3 million per year. Under that reading, the carbon market alone removes less NO_x and leaves more for the local layer. Excluding measures already profitable at the 2019 price (the revealed reading) raises the first-best gain to \$75.6 million. The uniform-price loss moves between \$28.3 million and \$34.0 million across the three readings and the loss under the estimated subdivision between \$11.3 million and \$24.6 million. Holding the allowance price at the 2019 level of \$16.84 instead of the social cost of carbon raises the first-best gain to \$100.5 million, because carbon pricing then removes almost no NO_x on its own. A price of \$65 leaves the first-best gain and the uniform loss within \$0.2 million of the headline and moves the subdivision losses by \$2 million to \$16 million, because the carbon price shifts the steps on which communities withhold. Fixing the cap at the quantity that clears at \$44.28 and letting the price re-clear (the framing of Section 2) changes the first-best gain by \$0.01 million. It lowers the equilibrium carbon price to \$42.82 because community pricing induces co-abatement (Proposition 3). Cross-regime welfare differences are therefore driven by local NO_x pricing rather than by the carbon side.

Table G2. Equilibria of the strategic pricing games on the engineering menus

Pricing game	Starts	Verified	Distinct	Loss vs. price-taking ($M/yr)			Median σ_j/d_j	Residual gain
			equilibria	min	max	PT start	(range)	($M)
Allowance price held at $44.28								
Unitary seller, n_j=1 (end states)	11	0	–	81.50	113.63	112.76	4.0	1.23–2.84
Sub-sellers at $\hat{n}_j$	5	5	5	8.58	13.21	11.30	0.25–0.78	0
Uniform $\bar{n}$=2	5	3	3	23.78	47.37	47.37	3.35–7.05	0
Uniform $\bar{n}$=3	5	5	5	18.97	45.36	30.48	2.34–9.61	0
Uniform $\bar{n}$=5	5	5	5	15.33	31.04	15.33	2.70–6.04	0
Uniform $\bar{n}$=10	5	5	5	15.16	21.78	15.24	4.30–7.10	0
Uniform $\bar{n}$=20	5	5	5	21.66	37.00	36.37	4.45–7.19	0
Uniform $\bar{n}$=50	5	5	5	39.24	43.90	43.90	3.02–6.16	0
Uniform $\bar{n}$=100	5	5	5	38.59	53.74	53.74	2.71–5.20	0
Uniform $\bar{n}$=1,000	5	5	5	45.35	66.00	58.91	0.88–2.78	0
Fixed cap that clears at $44.28								
Unitary seller, n_j=1 (end states)	11	0	–	81.26	113.50	108.66	2.1	1.02–5.71
Sub-sellers at $\hat{n}_j$	5	5	5	8.81	13.02	13.02	0.39–0.82	0
Uniform $\bar{n}$=3	5	4	4	14.74	30.32	27.71	3.04–8.25	0
Uniform $\bar{n}$=5	5	5	5	14.42	16.80	15.61	4.86–8.85	0
Uniform $\bar{n}$=10	5	5	5	15.24	21.86	21.72	3.29–8.79	0

Notes: County communities ($J = 58$). Each game is started from the price-taking profile $\sigma^0 = d$ and from random profiles $\sigma_j^0 = d_j \exp(2z_j)$ (four per game, ten for the unitary seller). A start is verified when a full sweep of exact best responses finds no deviation worth more than the tolerance; distinct equilibria are verified fixed points that differ in some price by more than 10^{-6} relative. The unitary game reached the round limit of 200 from every start under both carbon-market treatments, so its rows report the end states: the loss range across end states, the loss at the price-taking start, and the residual gain, the largest best-response gain any county still had at the end, whose range across end states is the last cell. The approximate (epsilon) equilibrium is the end state with the smallest residual gain, the price-taking start at the held allowance price and a random start under the fixed cap, and its median σ_j/d_j is the one shown. Verified games have a residual gain below the tolerance, written 0. The median σ_j/d_j column is the median across counties of the equilibrium price relative to marginal damage, shown as its range over the verified equilibria. The fixed-cap unitary end states and $\bar{n} = 10$ band were run on 2026-11-03 (out/resolve_2026-11/unitary_eps/).

Table G3. Welfare Losses by Regime under the Cap Framing

Welfare loss relative to the price-taking first best ($M per year; larger = worse)			
Regime	Perfect information	Estimation error (lognormal, CV = 1)	Verified equilibria / starts
Uniform tax or cap-and-trade (I/III)	28.32	37.83	–
Differentiated tax or cap-and-trade (II/IV)	0	11.80	–
Community permits, price-taking (V)	0	0	–
Community permits, unitary, n_j=1 (V); approximate equilibrium, residual gain $1.0M	108.66 (81.26–113.50)	108.66 (81.26–113.50)	0 / 11
Community permits, subdivision at $\hat{n}_j$ (V)	13.02 (8.81–13.02)	7.88 (6.61–11.73)	5 / 5
Community permits, uniform $\bar{n}$=10 (V)	21.72 (15.24–21.86)	21.72 (15.24–21.86)	5 / 5
Community permits, uniform $\bar{n}$=3 (V)	27.71 (14.74–30.32)	27.71 (14.74–30.32)	4 / 5
Community permits, uniform $\bar{n}$=5 (V)	15.61 (14.42–16.80)	15.61 (14.42–16.80)	5 / 5

Notes: Communities are counties ($J = 58$); 2019 dollars. Fixed cap that clears at $44.28 in the status quo ($R = 6,772,930$ t); $\tau^A = 42.82$ at price-taking. Every entry is a loss relative to the price-taking community-permit equilibrium, whose gain over the carbon-only status quo is $66.45M per year. For the subdivided regimes the entry is the loss at the equilibrium reached from the price-taking start, with the range across the distinct verified equilibria beneath, and the last column counts verified equilibria over the starts (the price-taking start plus random starts). No exact equilibrium of the unitary game was found from any start; the row reports the approximate equilibrium, the end state after 200 rounds with the smallest remaining best-response gain (a random start), with the range across all end states beneath. The appendix table reports every end state. The estimation-error column draws $\hat{d}_j$ lognormal around d_j with sd = d_j from the fixed-cap run of 2026-10-29 (200 lognormal draws, shadow solver); the two taxes are re-solved at the estimates, the $\hat{n}_j$ row re-solves the sub-seller game at $\hat{n}_j(\hat{d})$ (carried from the fixed-cap band run of 2026-10-29 through 11-01; the price-taking-start equilibrium, with the best and worst equilibria beneath, averaged over 13 draws with a verified equilibrium), and the regimes that use no estimate repeat the first column.

Table G4. Sensitivity to the Carbon-Side Cost Reading and the Allowance Price

	Reading of the carbon-side menu			Allowance price	
	Engineering, net of energy savings (headline)	Gross outlay	Excluding measures already profitable at $16.84	$16.84	$65
$M per year unless noted					
Price-taking gain vs. status quo	66.46	75.27	75.57	100.49	66.37
Uniform-tax loss vs. price-taking	28.31	30.59	34.00	28.54	28.13
Uniform tax t^* ($/ton NO_x)	14,528	8,924	13,815	14,528	14,528
$\hat{n}_j$ loss, price-taking start	11.30	12.04	24.59	11.92	13.22
NO_x removed at price-taking (t)	6,522	6,108	5,405	6,478	6,555
CO_2 abated at price-taking (t)	6,800,184	3,982,858	920,732	5,878,902	7,186,946

Notes: The first column is the headline calibration: each carbon-abatement measure costs its engineering estimate net of the energy it saves, with the allowance price held at $44.28. The second column charges the gross outlay of each measure and ignores the energy savings. The third column removes from the menu the measures that already pay for themselves at the 2019 auction floor of $16.84, on the view that facilities would have adopted them without the policy. The fourth and fifth columns keep the headline menu and hold the allowance price at $16.84 and at $65 instead of $44.28. In every column every other input is at its headline value. The rows are the gain from price-taking community permits over the carbon-only status quo, the loss a uniform NO_x tax leaves relative to that first best and the tax rate that minimises it, the loss at the $\hat{n}_j$ sub-seller equilibrium reached from the price-taking start, and the NO_x and CO_2 removed at price-taking. County communities, 2019 dollars.

H Scope of the Strategic Simulations

The strategic regimes of Section 5 are solved as exact games on the engineering menus at the county unit only. The game enumerates the breakpoints of each community's step-function residual demand with the carbon market re-cleared at each candidate. Because the number of breakpoints grows with the number of communities and facilities in each receptor field, the game is not computable at the 529 places or the 9,109 tracts within our resources. The non-strategic regimes and the two taxes under estimation error are computed at all three units. Figure G3 in Appendix G plots the uniform-subdivision loss across $\bar{n}$ at the county unit with the range across the equilibria we find. An intermediate real geography such as ZIP codes or the AB 617 community boundaries, solved as an exact game, would show how market power and error diversification trade off as communities are defined more finely.