

Designing Environmental Markets to Address Community Impacts

Danae Hernandez-Cortes*

Bryan Leonard[†]

Sara Sutherland[‡]

September 2026

PRELIMINARY DRAFT. PLEASE DO NOT SHARE WITHOUT PERMISSION.

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

Abstract

Environmental markets offer a way to minimize the costs of meeting environmental goals in the aggregate, but regulated activities often have a spatial incidence that regulators cannot price without knowing community-level damages and preferences. This spatial incidence can affect both the efficiency and political viability of market-based approaches. We propose a stacked permit system in which firms covered by an aggregate cap must also hold community-specific permits in proportion to their measured local impact. Communities receive slack endowments of these permits, so equilibrium prices reveal community damages. The competitive equilibrium attains the constrained first best even when community damages are unknown to the regulator. However, the mechanism endows each community with market power and creates a coordination problem between communities that jointly make a competitive outcome unlikely. We derive closed-form welfare comparisons across market-based instruments and show that the ranking across instruments depends on the relative magnitude of regulator uncertainty vs. market power. Dividing each community's selling authority among independent bodies, where each keeps the revenue from its own sales while bearing a share of the community's exposure, eliminates the market-power distortions. Calibrated to California's carbon market, the competitive equilibrium improves on the status quo by \$4.0 million per year, strategic community pricing generates a \$12 million net loss, and, under damage-estimation error, segmented supply restores a \$2.1 million expected gain, with gains concentrated in disadvantaged communities, largely as permit revenue.

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

[†]University of Wyoming bryan.leonard@uwyo.edu

[‡]U.C. Davis sasutherland@ucdavis.edu

¹Tom 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.

1 Introduction

Environmental markets — approaches that use incentives and property rights to reduce the costs of achieving environmental objectives — have been applied to an increasing number of domains in recent decades, including fisheries, surface water, ground water, habitat, wetlands, air pollution, and climate change mitigation. While these tools have intuitive theoretical appeal and have been favored by many regulators, market approaches often face political opposition and concerns from affected communities (Fowlie et al., 2020). Among these concerns is the fact that regulators are typically focused on an aggregate environmental objective, while communities care about the spatial distribution of the associated changes in behavior (Banzhaf et al., 2019). Most environmental 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). The focus on trading among firms means that communities lack a way to express their preferences because they are typically not market participants.

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 the communities that relied on it (Aronoff and Rafey, 2026), and reducing CO₂ emissions via cap-and-trade can change communities’ exposure to local co-pollutants (Hernandez-Cortes and Meng, 2023). In principle, environmental markets could address these concerns through a variety of mechanisms that reflect variation in the social costs of the regulated activity across communities (e.g., Tietenberg, 1995; Hung and Shaw, 2005). However, these approaches require the regulator to know the full profile of community-level damages and preferences (Holland and Yates, 2015), which is demanding in practice.²

This paper explores an alternative approach to addressing heterogeneous community damages that are unknown to the regulator. We propose a stacked permit system where, in addition to the permits required under the aggregate cap, firms must purchase community-specific permits according to their measured impact on affected communities. Communities are endowed with a slack allocation of these permits that exceeds the business-as-usual baseline, so that the equilibrium outcome is determined by communities’ supply decisions rather than by the regulator’s uncertain damage estimates. The design is fundamentally Coasean in that it assigns rights over the contested activity to the parties who bear the damages (Coase, 1960).

²A related mechanism-design literature elicits what polluters privately know — their abatement costs — through incentive-compatible schemes (Kwerel, 1977; Dasgupta et al., 1980; Montero, 2008; Duggan and Roberts, 2002); but every mechanism in this lineage takes marginal damages as known to the regulator, and private information defeats decentralized bargaining over damages as well (Farrell, 1987; Myerson and Satterthwaite, 1983).

This mechanism also addresses the procedural dimension of environmental justice concerns because communities directly influence their own pollution exposure, rather than merely relying on compensation determined elsewhere (Schlosberg, 2007). There is also a working institutional precedent for this approach in Alaska’s community development quota program, which has endowed coalitions of villages with tradable shares of federal fishery quota since 1992 (Ginter, 1995; National Research Council, 1999; Berman, 2024). We present the mechanism and derive its properties in its most general form before focusing on the specific example of CO₂ abatement with local co-pollutants.

Cap-and-trade markets for CO₂ emissions provide an especially useful context to illustrate the mechanism for several reasons. First, the social benefit of reducing CO₂ emissions is global regardless of where abatement occurs, but this abatement also induces changes in co-pollutants such as NO_x and SO₂ whose welfare implications depend sharply on the spatial pattern of abatement (Burtraw et al., 2003; Agee et al., 2014; Fullerton and Karney, 2018). Second, the behavior of firms, the performance of existing markets, and the atmospheric and engineering relationships underlying the spatial incidence of co-pollutants are especially well-documented relative to some other environmental markets (e.g., Tessum et al., 2017; Holland et al., 2016). Third, the welfare implications of the co-pollutant problem are significant, since pricing local air pollution at source-specific damages would substantially raise the net benefits of US regulation (Muller and Mendelsohn, 2009), and the optimal carbon instrument itself becomes source-specific once spatially varying co-pollutant damages are counted (Muller, 2012).

Current leading proposals for differentiated carbon policy require estimates of community-specific damages from co-pollutants. The challenge is that when the regulator’s damage estimates include error, spatially differentiated regulation 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). Source-specific prices and trading ratios build damage estimates directly into the instrument (Muller and Mendelsohn, 2009; Hung and Shaw, 2005), and even optimally chosen ratios inherit the estimation problem (Holland and Yates, 2015). Zonal designs and offset systems coarsen the information requirement but do not remove it (Atkinson and Tietenberg, 1982; Krupnick et al., 1983). Facility-specific caps layered onto the carbon market ration quantities without prices, and still require the regulator to select facilities and levels (Burtraw and Roy, 2023). Each of these approaches relies on the same missing information, whereas assigning the permits to the communities themselves creates a mechanism through which markets can reveal community damages.

In the CO₂ abatement context, our mechanism consists of an economy-wide cap on CO₂ *emissions*, coupled with a requirement that firms hold permits to cover each unit of *pollution* they distribute to each of J communities. A source-receptor matrix, which we assume is known to the regulator, defines how many permits each firm must obtain from each community to cover one unit of emissions. In a competitive equilibrium, every community prices exposure at its own marginal damage, and the allocation attains the constrained first best even in the presence of uncertainty

because the equilibrium does not depend on community damage estimates. Because each community is the sole seller of an input that every firm in its receptor field must hold, we also explore the more likely non-competitive outcome, where communities behave as monopolists. The resulting equilibrium embeds both a monopoly markup and a cross-community stacking externality that reduces welfare relative to the first best.

We adapt the parameterization of [Fowle and Muller \(2019\)](#) to develop a closed-form comparison of our mechanism with differentiated and undifferentiated carbon taxes and cap-and-trade. Whether or not the non-competitive equilibrium associated with community permits outperforms these other instruments depends on the magnitude of welfare losses due to estimation error vs. market power. We then show that dividing each community’s selling authority among n_j^* independent sub-sellers — each of which under-prices permits because it keeps the revenue from its own sales but bears only a share of the exposure those sales create — can eliminate the market-power distortions exactly. Implementing this solution requires an estimate of community damages, but we show that the welfare consequences of mis-estimating these damages are smaller than those associated with the other instruments. We also show that a simple, no-information uniform partition of communities removes most of the market-power loss.

We calibrate the mechanism to California’s cap-and-trade program, covering 287 large stationary sources, with source-receptor transport from a reduced-complexity atmospheric model and abatement technology recovered from facility-level co-emission ratios and engineering cost data anchored to two revealed market prices. We first show that the competitive equilibrium of the stacked system attains the constrained first best, improving on the existing carbon-only program by roughly \$4.0 million per year. But, when each community prices as a single seller, markups reverse that gain and the mechanism produces a net *loss* of roughly \$12 million per year relative to the carbon market alone.

However, under plausible estimation error, the within-community segmentation of permit supply turns the mechanism’s \$12 million net loss into a \$2.1 million expected net gain over the status quo, roughly double the \$1.1 million delivered by the next-best policy (which is *undifferentiated* cap-and-trade). A uniform partition that gives every community the same number of seats matches the undifferentiated instruments under perfect information and outperforms every regulator-run instrument under estimation error, with a net gain of roughly \$1.3 million per year in both cases. Finally, we find that the mechanism concentrates welfare gains from both emissions reductions and permit revenues in disadvantaged communities. These gains accrue largely through permit revenue rather than avoided damages, because the scarcity rents that other designs allocate elsewhere accrue here to the communities bearing the exposure ([Fullerton, 2011](#)).

Our contribution builds upon and combines elements of the earliest proposals for environmental markets. [Crocker \(1966\)](#) and [Dales \(1968\)](#) proposed markets in environmental use, and [Montgomery \(1972\)](#) formalized the distinction between *emission* licenses and *ambient* licenses before the ambient variant was shelved for its informational and transactional demands ([Tietenberg, 1985, 1995](#)). Our mechanism combines an emissions-permit layer for the aggregate externality

and an ambient-permit layer for the local one, along with careful attention to the initial allocation of permits for the latter. Endowing communities with ambient permits converts Montgomery’s information problem into a revenue stream and allows communities to reveal optimal “receptor caps.” The closest precedent to our design is [Yang and Kaffine \(2016\)](#), whose mechanism assigns tradable permits for a localized pollutant to a single harmed community. We generalize their design to many communities linked through a transport matrix, stack it on a carbon market whose abatement technology links the two layers ([Agee et al., 2014](#)), and characterize the market power that this ownership structure creates, and propose a simple governance design to address it.

The remainder of the paper proceeds as follows. Section 2 presents the general model, defines the stacked permit mechanism, and establishes its core properties, including efficiency under price-taking and the market-power distortions that arise when communities price strategically. Section 3 adapts the Fowle–Muller parameterization to our two-pollutant setting and delivers the closed-form welfare comparison across instruments, with and without damage-estimation error. Section 4 develops the governance remedy for market power and characterizes its performance when damages must be estimated. Section 5 calibrates the mechanism to California’s cap-and-trade program, and a final section concludes.

2 A Stacked Permit Market

This section presents the model environment, defines the efficient level and spatial allocation of abatement, introduces our proposed mechanism, and derives the welfare implications under differing assumptions about the competitiveness of community permit markets. Throughout, our exposition focuses on the case of CO₂ abatement with a local co-pollutant for convenience, but the basic structure here generalizes to any environmental market where the regulator cares about some aggregate level of activity but communities care about the spatial incidence of that activity.

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 :

$$c_j = \sum_i h_{ji} e_i^L, \quad c_j^0 \equiv \sum_i h_{ji} \bar{e}_i^L, \quad (1)$$

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.

The aggregate and local emissions are linked given that firms’ decisions affect both emissions simultaneously, which is reflected in 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:*

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

Equation (2) assumes that increasing abatement of one type of emissions makes abatement of the other weakly cheaper. As we show in Appendix B, this general relationship ($\gamma_i \leq 0$) can be derived from an abatement cost function that separates the three possible types of abatement: activities or technologies that reduce only e_i^A , activities or technologies that reduce only e_i^L , and activities or technologies that reduce both. Reducing both types of emissions can be accomplished by reducing output or increasing fuel efficiency. If the other two classes of activities are independent (e.g., if capturing CO₂ does not raise the cost of reducing NO_x emissions and installing scrubbers does not raise the cost of capturing CO₂), then the two types of abatement must be weak complements due to the shared margin. Empirical estimates of abatement technology for CO₂ and its co-pollutants support treating the two margins as jointly produced (Burtraw et al., 2003; Considine and Larson, 2006; Agee et al., 2014).³

Turning to communities, we assume that each community has a damage function $D_j(c_j)$ that depends only on the ambient concentration of local pollutants in community j (e.g., Muller and Mendelsohn, 2009; Hsiang et al., 2019), and that this damage function is increasing, weakly convex, and known to community j only. Throughout, we assume that damages associated with the global pollutant (CO₂) are separable from local damages and independent of where emissions occur (e.g., Holland et al., 2016). Hence, there is a social cost of carbon (SCC) that we assume is known to the regulator and treated as a constant throughout our analysis (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 a constraint on the total level of aggregate abatement ensuring that it meets the pre-existing cap (Baumol and Oates, 1971; Montgomery, 1972):

$$\min_{\{a_i^A, a_i^L\}} \sum_i A_i(a_i^A, a_i^L) + \sum_j D_j\left(\sum_i h_{ji}(\bar{e}_i^L - a_i^L)\right) \quad \text{s.t.} \quad \sum_i a_i^A = R. \quad (3)$$

We assume there is a fixed emissions cap $\bar{E}^A$ that binds throughout our analysis, so required carbon abatement is pinned down by business-as-usual emissions in excess of the cap, $R \equiv$

³Estimation of jointly produced pollutants in this tradition dates to Färe et al. (2005) and, for interacting pollutants, Crocker (1981). Bonilla et al. (2018) find that CO₂ and NO_x abatement are sometimes substitutes rather than complements. If $\gamma_i > 0$ for all i , then $p_i < 0$ and the sign in Proposition 3 below reverses, so community permits raise rather than lower the carbon price.

$\sum_i \bar{e}_i^A - \bar{E}^A$. 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 every policy regime we compare shares the carbon layer, and 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 returns to this point.

The planner's objective is strictly convex and the constraint is an affine transformation of the choice variables, which guarantees a unique optimum.⁴ λ is the multiplier on the carbon constraint and the first-order conditions are given by

$$\nabla A_i(a_i^A, a_i^L) = \left(\frac{\lambda}{\sum_j D'_j(c_j^*) h_{ji}} \right) \quad \text{for all } i. \quad (4)$$

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 over the communities its plume reaches, which 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 $\text{MD}_i^* \equiv \sum_j D'_j(c_j^*) h_{ji}$.

2.3 The Mechanism and its Properties

Our mechanism is designed to “stack” on top of an existing cap-and-trade scheme that caps total emissions at $\bar{E}^A$, with a resulting equilibrium permit price for e^A of τ^A . Under our mechanism, in addition to needing a permit for every unit of aggregate emissions e_i^A , each firm must also hold a permit to cover its contribution to ambient concentration c_j in each community j , as determined by the source-receptor matrix $\mathbf{H}$. This means that each firm i must 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

$$\tau_i^L = \sum_j \sigma_j h_{ji}. \quad (5)$$

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. 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. Communities sell

⁴Strictness comes entirely from the cost side, because 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, and damages add a (weakly) convex term; the sum of a strictly convex and a convex function is 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.

at posted, take-it-or-leave-it prices that apply uniformly to every firm in their receptor fields so the mechanism operates through arm's-length markets rather than bilateral negotiation between firms and communities. Posted prices also prevent strategic holdup, 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,

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

Strict convexity makes each abatement margin increasing in its own price, and cost complementarity implies that a higher carbon price raises local abatement and vice versa, which corresponds to the co-benefit channel at the center of the co-pollutant literature.

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 revenue against the damage from the exposure its sales admit. Under the assumption of quasi-linear payoffs,⁵ community j solves

$$\max_{s_j \in [0, L_j]} \sigma_j s_j - D_j(s_j). \quad (7)$$

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 $(\tau^A, \{\sigma_j\}_{j=1}^J)$ and an allocation $(\{a_i^A, a_i^L\}_{i=1}^I, \{s_j\}_{j=1}^J)$ such that (i) each firm's abatement solves Equation (6), taking $(\tau^A, \{\sigma_j\})$ 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 $\sigma_j \geq 0$, $s_j \geq \sum_i h_{ji}(\bar{e}_i^L - a_i^L) = c_j$, and $\sigma_j(s_j - c_j) = 0$, 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, $\tau^A = \lambda$; and (c) no equilibrium object depends on the regulator's beliefs about community damages, conditional on the realized cap.*

Proof: [Appendix A.1](#).

⁵Permit 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](#)).

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 shows the cap level is the one channel through which regulator bias can still enter.

The perfectly competitive equilibrium serves as a useful benchmark to illustrate how the combination of community permit allocations with the price mechanism can use market forces to efficiently allocate abatement across space. However, in reality, our proposal grants market power to communities (Fullerton and Metcalf, 2002) while simultaneously exposing firms to an input assembly problem, making the competitive ideal unlikely. To understand how these forces work against the efficiency gains associated with decentralizing decision-making, we next model each community j as a single decision-maker that takes rivals’ prices and the carbon price as given but recognizes its impact on its own price. Each community solves:

$$\max_{s_j \in [0, L_j]} \sigma_j(s_j) s_j - D_j(s_j), \quad s_j(\sigma_j; \boldsymbol{\sigma}_{-j}, \tau^A) = c_j^0 - \sum_i h_{ji} a_i^L(\tau^A, \sum_k \sigma_k h_{ki}). \quad (8)$$

Here $s_j(\cdot)$ is community j ’s *residual demand* for its permits as a function of its own price, other communities’ permit prices, and the aggregate emissions price. This residual demand is strictly decreasing in σ_j , with inverse $\sigma_j(s_j)$. Because each community is the sole seller of its own permit, posting a price and choosing a supply are the same decision — a point on the community’s residual demand curve — so the posted-price institution facing firms is unchanged. We formalize the choice in quantities because that is the form the within-community analysis of Section 4 requires. Price competition in the usual sense cannot arise between communities, because their permits are complements rather than substitutes, so that a firm needs permits from every community in its receptor field and no community can capture another’s demand by undercutting. 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. (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, so that whenever these net cross-demand externalities are positive and together outweigh the community’s own carbon-price feedback, markups stack above even 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, resulting from the fact that each community is a monopolist with respect to its own permits, and so can strategically withhold supply to increase its price. The consequence is prices that are too high and there is too much local abatement. The second inefficiency comes from the fact that firms need to assemble permits from *each* community that they are connected to via **H**. This results in 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 dynamic will drive prices even higher, leading to even more 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.

Whether the informational advantages of the mechanism are dominated by the market-structure inefficiencies cannot be resolved at this level of generality, so quantifying the tradeoff requires parametric structure, which we adopt in the next section. However, before turning to the parameterized version of the model, we derive one more general-form result that does not depend on 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 cap clears at a lower carbon price, as Proposition 3 states formally.

Proposition 3 (Layer interaction). *Hold the cap R fixed. Adding community permit markets with prices $\sigma \succeq \mathbf{0}$ lowers the market-clearing carbon price, $\tau^A(\sigma) \leq \tau^A(\mathbf{0})$. 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, whose goal is to achieve the efficient spatial distribution of local abatement without the regulator needing to know community damages. Because changes in local abatement change each firm’s marginal cost of *aggregate* abatement, however, setting the aggregate cap optimally once again requires the regulator to estimate local damages. The optimal cap equates the marginal cost of aggregate abatement with the SCC, and that marginal cost depends on local abatement, which in turn depends on local damages. A simple workaround is to make the aggregate instrument a *tax* rather than a *cap*, in which case the regulator needs to know only the SCC, which is unaffected by local damages. Corollary 1 in Section 3 formalizes this intuition.

3 Comparison to Other Policy Instruments

Thus far, we have shown that a stacked community pollution-exposure permit system 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 also serves as the basis for the empirical illustration in Section 5.

3.1 Parameterization

We anchor our analysis on a generalization of the functional form assumptions in Fowle and Muller (2019) (FM), so that our policy regime characterizations 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 demonstrates that at $\gamma_i = 0$ the local layer is equivalent their model, and that each of their analytical results is recovered.⁶

For notational convenience, the results below use three derived constants:

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

These are the slopes of firms' linear abatement responses to the prices they face. A higher carbon price raises carbon abatement at rate ψ_i , a higher local price raises local abatement at rate $\tilde{\beta}_i$, and p_i is the rate of abatement co-benefits. On the damage side, Assumption 2 makes the marginal damage of exposure in community j the constant d_j . All equilibrium objects — firm responses, the carbon-clearing price, community residual demands, and the strategic pricing equilibrium — are linear with closed forms, which are derived in Appendix A. Lemma 1 (Appendix A.4) gives a sufficient condition for the strategic pricing equilibrium to be unique. The condition requires that each community's own-price effect on its residual demand outweigh the combined feedback from other communities' prices and from the carbon market. Practically speaking, this hold when the transport matrix is sufficiently sparse so that each community's supply decisions does not affect

⁶The cost linkage adds three channels that are absent from their analysis. Local pricing affects the aggregate permit price (Proposition 3), every regime's welfare loss includes a partial offset because the carbon market re-clears in response to local mispricing, and the optimal level of the aggregate cap depends on local damages (Corollary 1). All three channels vanish as $\gamma_i \rightarrow 0$. 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, which is what allows the damage information to be revealed by communities' decisions.

every other community. Because atmospheric transport connects every source to every receptor, this sufficient condition fails in our calibration, and we instead verify uniqueness computationally (Section 5).

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, which we denote $MD_i \equiv \sum_j d_j h_{ji}$, and the efficient carbon price is the shadow price of the cap, λ . Every regime we consider confronts firms with some pair of prices, so its welfare loss can be written, as Proposition 4 shows, 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 $\Delta_i^L = \tau_i^L - MD_i$, market clearing implies $\Delta^A = -\sum_i p_i \Delta_i^L / \sum_i \psi_i$ and*

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

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. Because firms' co-abatement of carbon is affected by local prices, when those local prices are wrong, the carbon price re-clears based on the error. This re-clearing undoes the common component of the mispricing, so only its dispersion across sources appears in net welfare loss. This is why the entries in Table 1 depend on the dispersion of pricing errors across sources but not on the level of those errors, and why estimation errors that shift all sources in the same direction are less costly.

We compare five regimes under a regulator who holds beliefs $\hat{d}_j$ about community damages, with perfect information as the special case $\hat{d}_j = d_j$ (cf. Fullerton, 2001). Regimes I and II are the undifferentiated and differentiated cap-and-trade systems, respectively, while Regimes III and IV are the undifferentiated and differentiated emissions taxes. Note that firms face uniform prices under regimes I and III, and that regimes II and IV confront firms with localized prices that are based on the regulator's damage estimates.⁷ 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

⁷In 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).

on the regulator's beliefs. Proposition 5 formalizes the welfare comparison across instruments for a given estimate of community damages $\hat{\mathbf{d}}$.

Table 1. Regime comparison under damage beliefs $\hat{d}_j$ (perfect information: $b_i = 0$)

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

Proposition 5 (Damage robustness and regime ranking). *Fix any beliefs $\hat{\mathbf{d}}$ and condition on the realized 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 $\Delta_i^L = t^*(\hat{\mathbf{d}}) - \text{MD}_i$ 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 diverges while the strategic-pricing loss is unchanged.*

Proof: Appendix A.6.

Part (d) weighs the market-power losses relative to the informational advantages of the mechanism, and is of 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, from Lemma 1 (Appendix A.4), community j 's markup is one-half the gap between its demand-to-elasticity ratio and its own marginal damage, so market power is most severe for communities facing large permit demand relative to the damages they suffer, and mildest where damages are high. Finally, the stacking externality grows with the number of communities a firm must obtain permits from, so the market-power wedge rises as communities are defined more finely.

Proposition 5 holds the belief error fixed and therefore cannot provide a general ranking of the regimes against one another, because at any fixed $\hat{\mathbf{d}}$ the comparison between uniform and differentiated instruments depends on the particular error realized. Proposition 6 therefore lets treats the error as a random variable and evaluates the comparison in expectation. To do so, we assume that the regulator must commit to its instrument before damages are known, basing it on an estimate that equals the truth plus an estimation error with mean zero and known covariance. Expected welfare losses then average over the possible errors, so the regimes are compared as gambles on the quality of the regulator's information. Throughout, communities act on the damages they actually experience. The same calculation can be read with the damages themselves as the random object and the regulator's estimate as their mean, which is how Section 4 compares the partition with the differentiated instruments under matched information. 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. 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 commits at beliefs $\hat{\mathbf{d}} = \mathbf{d} + \varepsilon$, where the estimation error ε 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,*

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

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

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

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 $\Sigma = \text{CV}^2 \Sigma_0$, the expected gain from differentiation is strictly decreasing in CV and, provided $\mathcal{L}_0^{\text{undiff}} > 0$, changes sign at a unique threshold $\text{CV}^* > 0$ 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 ε .

Proof: [Appendix A.7](#).

Proposition 6 says that the expected-loss ranking depends on how each instrument absorbs error. A uniform price depends only on a single average of the estimation errors, so idiosyncratic mistakes tend to cancel across sources. In contrast, differentiated prices pass each community's error through one-for-one, so uniform instruments diversify estimation risk while differentiated instruments do not. This is relevant in our context, since recent general-equilibrium estimates find large welfare gains from damage-based differentiation of US air pollution regulation ([Hollingsworth et al., 2026](#)). Crucially, the size of those gains is driven by heterogeneity in community damages, which is the exact same feature that makes implementing differentiated instrument on noisy estimates costly.

Finally, we return to a point first raised in Section 2, which is that damage estimation error also impacts the level of the aggregate cap and can therefore cause welfare losses even under the perfectly competitive version of the stacked community permit market. Here, we use the FM parameterization to formally compare the performance of cap-and-trade vs. 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*

$$R^* = \left(\sum_i \psi_i \right) \text{SCC} + \sum_i p_i \sum_j d_j h_{ji}, \quad (13)$$

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

$$\mathcal{L}^{\text{cap}} = \frac{(\sum_i p_i b_i)^2}{2 \sum_i \psi_i}. \quad (14)$$

(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 places the corollary in the prices-versus-quantities tradition ([Weitzman, 1974](#)), with the difference that the asymmetry here arises from information about a *different* pollutant.

4 Resolving Market Power Distortions

If community permit markets operate under conditions of perfect competition, they result in both the efficient level and spatial allocation of abatement, and they do so without requiring the regulator to estimate community damages. However, our mechanism simultaneously endows each community with market power and creates a price-setting externality between communities, with the likely result that prices will be set inefficiently high. Whether these welfare losses dominate the information gains of the mechanism is an empirical question that we turn to in Section 5. Before doing so, we introduce a refinement to the community permit model designed to counteract the market power distortions, and we show how it can be implemented with estimated damages or with no damage information at all.

4.1 Within-Community Partitions

Instead of treating each community as a unitary decision-maker, this extension devolves the authority to sell to a lower level *within* communities, while maintaining a single community-level price that is quoted to firms. The partition divides the community (which includes the community's residents and their damage claim) rather than a stock of permits. The community's permit endowment $\bar{L}_j$ remains an undivided common pool from which each sub-entity may offer sales. This means that each sub-entity faces an *unbounded* endowment, in the sense that it may authorize any quantity of sales from the pool. Because the community's allocation is slack, supply decisions are determined by trading off permit revenue against damages incurred. The community's total supply available to firms is the sum of the sub-entities' offers, and the single price σ_j clears the market for total supply based on firms' residual demand. None of this alters the decision problem facing firms, which continue to take the community price as given. Specifically, let community j 's supply be chosen by n_j symmetric sub-entities (council seats, neighborhood representatives), each bearing an equal share of community damage and offering supply from a common pool independently:


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

Equation (15) embeds two opposing forces, since each sub-entity earns the full market price on its own sales but bears only a $1/n_j$ share of the damages those sales cause (because it represents $1/n_j$ of the community's residents). Relative to the unitary seller, its incentive to guard against pollution exposure is therefore weaker. Similarly, each sub-entity internalizes the price impact of its own sales alone rather than the pool's, so its incentive to raise prices by withholding supply is also weaker. Nothing in Equation (15) bounds a sub-entity's offer from above. Such an upper bound is unnecessary, because firms demand permits only for the exposure they create, which is at most baseline exposure c_j^0 , while the community's endowment exceeds that level by construction. Any authorization beyond firms' demand therefore goes unsold at a zero price, at which a sub-

entity earns nothing and still bears its share of the damage.

Figure 1 illustrates how partitioning community decision-making erodes the community-level markups. At an interior equilibrium, the first-order condition to Equation (15) for each sub-entity results in the same offer, so the within-community equilibrium is symmetric, and summing the n_j first-order conditions yields a single community-level pricing rule. At $n_j = 1$ that corresponds to the monopoly equilibrium of Section 2. The community as a whole behaves like a Cournot oligopoly of n_j sellers sharing one market, and its price declines in n_j , from the monopoly price at $n_j = 1$ toward giving permits away as n_j grows large and each seller's damage share vanishes. The within-community partition thus balances the monopoly withholding and anticommons stacking that operate *between* communities against the free-riding it creates *within* communities.

Figure 1. Partitioning a community's selling authority



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 seats, $\sigma_j(n_j)$, is 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^* . The two lower curves are the terms of a seat's first-order condition, the share d_j/n_j of marginal damage it internalizes and the markup $-\sigma'_j s_j/n_j$ it earns by withholding its own sales; their sum is the price, and both fall as seats are added.

This approach also maps onto familiar governance structures, such as multi-seat commissions or resident voting bodies whose members authorize supply independently, and communities have long demonstrated the capacity to build and operate institutions of exactly this kind (Ostrom, 1990). For instance, Alaska's community development quota groups and territorial-use-rights fisheries already vest collectively held resource rights in community-scale bodies (Ginter, 1995; National Research Council, 1999; Berman, 2024; Wilen et al., 2012).⁸

⁸The choice of n_j as well as the choice of the community partition itself is also considered part of the governance

Devolving price-setting authority within communities also helps address a potential political economy and environmental justice concern with our mechanism, which is that 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. The within-community partition of sellers or “seats” 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). Each body needs to know only its residents’ damage share and the prices it observes, and bodies that raise or lower their authorizations in response to those prices converge to the equilibrium, provably under the condition of Lemma 1 and computationally in every calibrated case we solve. The administrative requirements are a posted community price and a record of each body’s authorizations and receipts.

Proposition 7 formalizes the equilibrium of the within-community pricing scheme and defines the partition n_j^* at which the two forces exactly offset, restoring the efficient community-level price in the absence of perfect competition between communities.

Proposition 7 (Optimal partition). *Fix a committed partition $\mathbf{n} = (n_1, \dots, n_J)$. (a) Given rivals’ prices σ_{-j} and the carbon price τ^A , community j ’s equilibrium price is*

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

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

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

the same expression with rivals priced at their marginal damages and the carbon market at its first-best price λ . At $\mathbf{n} = \mathbf{n}^*$, 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 proof of Proposition 7 leverages the symmetric first-order conditions of Equation (15). Each seat equates the marginal revenue of its own sales (the price less the price impact on the units it already sells) to its share d_j/n_j of marginal damage. Summing the n_j conditions gives a community-level rule in which the price equals the sum of marginal damage and the choke price of residual demand, σ_j^0 , divided by $n_j + 1$. Setting that price equal to d_j and solving for n_j gives Equation (17), so n_j^* is the ratio of the choke price to marginal damage, the number of seats at

design given that this decision might trade off local heterogeneity against cross-boundary spillovers ([Costello and Kaffine, 2018](#)).

which the seats' collective withholding exactly offsets their diluted damage shares. The remaining step is to verify that pricing at marginal damage in every community is a joint equilibrium. At that profile every rival prices at its own marginal damage and the carbon market clears at the planner's price λ , so each community's residual demand is the one that defines n_j^* , and the pricing rule at n_j^* seats then returns exactly d_j , which confirms the profile as a fixed point of the joint system.

The optimal partition varies across communities, and the formula in Proposition 7 makes n_j^* the ratio of community j 's choke-level permit demand to the product of its marginal damage and the slope of its residual demand. Hence, it is large where market power would otherwise be most severe (e.g. communities facing heavy permit demand relative to the damages they suffer) and small where damages are high relative to demand. Communities whose residual demand cannot support a price as high as their marginal damage (choke price below d_j) have $n_j^* \leq 1$ and need not be partitioned at all. Although we work with the FM parameterization here for convenience, existence of n_j^* does not require linearity. Proposition 9 (Appendix C) derives the general form $n_j^* = 1 + (-\sigma'_j)s_j^*/D'_j$ under the maintained assumptions of Section 2 alone.

4.2 Estimating the Partition

In practice, n_j^* cannot be implemented directly because it is a function of community damages, implicating the same information problem the community permit mechanism is meant to avoid. Because the partition and the differentiated instruments require the same damage estimates, 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 commits at its estimate $\hat{\mathbf{d}} = \mathbf{d} + \varepsilon$, with ε mean zero and covariance Σ , while communities price at the damages they actually experience (as in Proposition 6). The regulator must select a partition $\hat{n}_j$ from its estimates:

$$\hat{n}_j = \frac{c_j^0 - \sum_{k \neq j} (\sum_i \tilde{\beta}_i h_{ji} h_{ki}) \hat{d}_k - (\sum_i p_i h_{ji}) \tau_{\hat{\mathbf{d}}}^A}{\hat{d}_j \sum_i \tilde{\beta}_i h_{ji}^2}, \quad (18)$$

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

Proposition 8 (Information aggregation under damage uncertainty). *Suppose the regulator commits the partition 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*

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

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,

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

(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 the community mechanism is likely to lead to smaller welfare losses than the other differentiated regimes in the presence of uncertainty, because it embeds a partial error correction in which communities price based on their *actual* damages. Under the differentiated instruments, the regulator's estimate (and any associated error) passes into prices on a one-for-one basis. Under our mechanism the regulator's damage estimate only sets the partition, and each community then pulls its own price part of the way back toward its true damage.⁹ Section 5 quantifies this damping in the calibration and shows where it is concentrated.

4.3 A Uniform Partition

Although the partition $\hat{n}_j$ outperforms traditional differentiated mechanisms by dampening the regulator's estimation error, it still requires J damage estimates. Here, we consider a simpler alternative. Instead of trying to calibrate a unique partition for each community, the regulator can give every community the same number of seats, $\bar{n}$, which requires no damage information. Corollary 2 characterizes how a uniform partition would work and why choosing $\bar{n}$ can be done without damage estimates.

Corollary 2 (Uniform partitions). *Fix a uniform partition $n_j = \bar{n}$ for every j , and let $n_j^*(\bar{n})$ denote the expression in Equation (17) evaluated at the carbon price that clears the aggregate market under that partition 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*

$$\sigma_j(\bar{n}) = d_j \frac{n_j^*(\bar{n}) + 1}{\bar{n} + 1}, \quad (21)$$

so its markup over marginal damage, $\mu_j(\bar{n}) = d_j (n_j^*(\bar{n}) - \bar{n})/(\bar{n} + 1)$, falls in magnitude like $1/(\bar{n} + 1)$, changes sign at $\bar{n} = n_j^*(\bar{n})$, and makes the aggregate welfare loss U-shaped in $\bar{n}$; (b) for every community with $n_j^*(\bar{n}) \geq \bar{n}$, the markup is at most $2/(\bar{n} + 1)$ times the unitary markup at the same residual demand, $d_j (n_j^*(\bar{n}) - 1)/2$, so that community's contribution to the welfare loss falls at least at the rate $4/(\bar{n} + 1)^2$; (c)

⁹The dominance of the within-community partition can nonetheless fail when highly heterogeneous θ_j combine with strongly negatively correlated errors, because uneven damping distorts the spatial pattern of the surviving error, and a worse pattern across communities can overwhelm the error-dampening within communities. Whether that case is empirically relevant is a question about the joint distribution of damages, and we address it in the next section.

no element of the equilibrium depends on the regulator’s damage beliefs, so the welfare loss under a uniform partition is deterministic rather than expected; (d) $\bar{n} = 1$ is the monopoly equilibrium and, as $\bar{n} \rightarrow \infty$, 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 $\bar{n} \geq n_j^*(\bar{n})$, whereas the undershoot wedge for $\bar{n} < n_j^*(\bar{n})$ is bounded only by the unitary markup.

Proof: [Appendix A.10](#).

Under the corollary, 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.¹⁰ However, the markup changes sign at the community’s own n_j^* , 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 directions are not symmetric, because adding seats beyond n_j^* pushes prices toward zero and the equilibrium toward the carbon-only status quo, so the cost of allocating too many seats is bounded status quo welfare loss of having no local instrument at all. Too few seats, by contrast, leave a markup that is bounded only by the monopoly loss. A regulator who knows nothing about the distribution of n_j^* should therefore err toward more seats.

Because the free-riding discount takes the same form in every community, we can characterize the gains of implementing the uniform partition without reference to unknown community damages. In any community with $n_j^* \geq \bar{n}$, at most $4/(\bar{n} + 1)^2$ of the unitary market-power loss survives, which is 3.3 percent at $\bar{n} = 10$ and 0.9 percent at $\bar{n} = 20$, so a regulator who can certify only that every community’s n_j^* exceeds twenty knows that a uniform partition of that size removes over 99 percent of the market-power loss everywhere. In practice, losses can exceed this bound in communities with $n_j^* \leq \bar{n}$, a point we return to in the calibration.

5 Empirical Illustration: California’s Carbon Market

5.1 Calibration

We calibrate the model to California’s cap-and-trade program in 2019, covering 287 large stationary sources such as power plants, refineries, cement plants, hydrogen producers, and other industrial facilities. Each facility’s contribution to ambient $\text{PM}_{2.5}$ concentrations comes from individual emission runs of the InMAP atmospheric model ([Tessum et al., 2017](#)), which give the transport matrix $\mathbf{H}$ over a receptor grid of 9,109 census tracts. The carbon layer is the same in every regime, and we set its level to match the 2019 allowance price of \$16.84 per ton. The regimes differ only in how they price the local impact. To estimate community damages d_j , we multiply tract popula-

¹⁰Corollary 2 is stated for the linear parameterization, and 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, use that structure, but the qualitative content does not. Under the maintained assumptions of Section 2 alone, Proposition 9(b) gives a within-community price that is strictly decreasing in the number of seats and crosses marginal damage exactly once, at n_j^* , so a uniform partition still overshoots the communities with $n_j^* < \bar{n}$ and undershoots the rest, still depends on no damage estimate, and is still bracketed by the monopoly outcome at one seat and, as each seat’s damage share and price impact vanish with the number of seats, by the carbon-only status quo.

tion by a concentration-response function for adult mortality (Krewski et al., 2009) and a value of statistical life. We use a conservative \$7 million figure so that estimated damages remain within each facility’s physical abatement limits.¹¹

We focus on NO_x as the CO_2 co-pollutant for three reasons. First it is one of the main local air pollutants emitted by California’s covered facilities. Second, NO_x is a $\text{PM}_{2.5}$ precursor. Third changes in GHG emissions are strongly correlated with changes in NO_x (Burtraw et al., 2003).¹² Damages from other co-pollutants are omitted due to computational constraints, which means that our local damages are likely an underestimate of the full co-pollutant burden.

The cost interaction γ_i , is not directly identified by any available price variation. Instead, we measure the three underlying abatement margins and compute $(\alpha_i^A, \alpha_i^L, \gamma_i)$ from them. Co-emission ratios ϕ_i are measured directly as facility-level NO_x -to- CO_2 ratios, obtained from CARB’s pollution mapping tool for 2019. Cleanup costs c_i^c come from engineering cost curves fit to regulatory technology assessments (U.S. Environmental Protection Agency, 2020; South Coast Air Quality Management District, 2021b), which fix the shape of each sector’s schedule but not its level. We set the level so that the marginal cost of full abatement in the largest sector equals \$5,410 per ton, the 2019 price of a NO_x trading credit in Southern California’s RECLAIM market (South Coast Air Quality Management District, 2021a). The rescaling factor is large, roughly twenty, which is consistent with the documented gap between engineering estimates and revealed compliance costs (Fowlie, 2010; Fowlie and Muller, 2019). Appendix F explains why we do not clear a market as Fowlie and Muller (2019) do, and Section 5.5 reports results across the range of factors implied by the other 2019 credit vintages.

Activity-margin costs c_i^q take their shape from sector-specific abatement cost estimates (e.g., Gillingham and Stock, 2018; Fowlie et al., 2016; California Air Resources Board, 2022). We set the level of these costs so that the model’s carbon market clears at 2019’s average allowance price of \$16.84 per ton (California Air Resources Board, 2019). Community permits did not exist in 2019, so that price came from a market in which firms faced the carbon cap alone. In the model that market clears where aggregate carbon abatement equals the cap. When every firm’s abatement is interior, the clearing price is the cap divided by the sum of firms’ carbon-abatement responses,

$$\tau_{\text{SQ}}^A = \frac{R}{\sum_i \psi_i}, \quad (22)$$

and because every ψ_i scales with the level of the activity-cost curve, a single rescaling of that level sets this price. A few facilities are at their abatement bounds in the status quo, so we solve the bounded clearing condition numerically and choose the rescaling so that it holds at the observed \$16.84 exactly (Appendix F). We assume that carbon capture is unavailable in the 2019 baseline, matching its absence from the compliance record.

The calibration requires several other design choices. First, we set the aggregate cap 30 percent

¹¹Appendix F provides additional details.

¹²For instance, Figure G1 shows a strong correlation between changes in NO_x and GHG emissions between 2018 and 2019.

below business-as-usual emissions so that it binds in every regime we compute. Second, because atmospheric transport gives some communities only trace exposure, we solve all equilibria imposing non-negative values.¹³ Third, our core results use counties to define communities, but we also compute results using census tracts and Census places. Table 2 summarizes the key model objects in the calibration across all three community definitions.

Table 2. Calibration summary

	Tract	Place	County
Communities J	9,109	529	58
Facilities I	287	287	287
Median d_j (\$/ $\mu\text{g}/\text{m}^3$)	141,120	1,624,829	6,272,129
Median MD_i (\$/ton NO_x)	1,256	1,256	1,256
Median α_i^L (\$/ton ²)	385	385	385
Carbon price τ^A (\$/tCO ₂ e)	16.51	16.51	16.51
Cap binds	Yes	Yes	Yes
Uniqueness norm bound, max LHS	179,304	8,397	264
Aggregate-interiority violations	7	7	7

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 and median α_i^L are source-level objects and do not vary with the community definition. The carbon price row reports the equilibrium τ^A under the price-taking mechanism; it lies below the \$16.84 status-quo anchor because community pricing induces carbon co-abatement (Proposition 3). The interior solver of Assumption 2, which ignores the abatement bounds, gives \$16.12. The norm-bound row tests the uniqueness guarantee of Lemma 1. The guarantee applies when a normalized measure of cross-community and carbon-market feedback, defined in Equation (41), is below one for every community. The row reports the largest value of that measure, which exceeds one at every community definition, so the guarantee does not apply and uniqueness is verified computationally instead (Appendix F). The final row counts sources whose carbon abatement sits at its physical bound, namely the seven facilities reporting zero GHG emissions, retained with a one-ton floor, which constitute a negligible share of covered carbon.

We compute the equilibrium of every regime in Table 1, which includes the four regulator-run instruments, the mechanism under price-taking, and the mechanism under Cournot pricing at both the unitary default $n_j = 1$ and the estimated partition $\hat{n}_j$, as well as two uniform partitions $\bar{n} = 20$ and $\bar{n} = 50$. We evaluate each regime’s welfare loss relative to the constrained first best, under perfect information and under damage-estimation error. The remaining subsections report the aggregate welfare comparison, its distribution across communities, and its sensitivity to the calibration’s key inputs.

¹³This would imply that communities that cannot support a positive price post zero, and firms’ abatement respects physical bounds. Firm responses solve box-constrained quadratic programs. The community pricing game is solved as a complementarity problem in which each community either posts a positive price satisfying its first-order condition or posts zero, with the carbon market re-clearing at each step. Because the diagonal-dominance condition of Lemma 1 fails under dense atmospheric transport, uniqueness is verified computationally. Best-response iteration from random starting points converges to the same equilibrium at every community definition (Table G2). Appendix F gives details.

5.2 Aggregate Welfare Results

Table 3 reports the welfare loss of each regime relative to the constrained first best. These losses are small relative to the welfare magnitudes of the carbon market itself, but the relevant comparison (developed at the end of this section) is to the \$32 million of damages from *local* co-pollutants. The first column is the certainty benchmark, in which the regulator has perfect information, so the differentiated instruments and the price-taking mechanism all reach the first best, while the uniform instruments entail a \$2.6 million cost of ignoring the spatial incidence of co-pollutants. When communities behave as monopolists and there is no sub-community partition, our mechanism generates a \$15.9 million loss, roughly six times the welfare cost of the undifferentiated policy that motivates spatial differentiation, and a net loss of \$12.0 million relative to running the carbon market alone. However, implementing the within-community partition n_j^* (which is feasible with perfect information) restores the constrained first-best outcome. The uniform partitions both perform similarly to (albeit slightly worse than) undifferentiated carbon pricing in the presence of perfect information.

Table 3. Welfare losses by regime (communities are counties)

Welfare loss relative to the constrained first best (\$M per year; larger = worse)		
Regime	Perfect information	Damage-estimation error (sd = 1 d_j)
I. Undifferentiated cap-and-trade	2.64	2.88
II. Differentiated cap-and-trade	0	2.94
III. Undifferentiated tax	2.64	2.88
IV. Differentiated tax	0	2.94
V. Community (price-taking)	0	0
V. Community (monopoly, $n_j=1$)	15.93	15.93
V. Community (Cournot, $\hat{n} = n^*(\hat{d})$)	0.00	1.84
V. Community (Cournot, uniform $\bar{n}=20$)	2.64	2.64
V. Community (Cournot, uniform $\bar{n}=50$)	2.65	2.65

Notes: Communities are counties ($J = 58$). Every entry is a welfare *loss*; zero is the best attainable outcome given the carbon cap. Left column: the regulator knows community damages exactly. Right column: the regulator acts on noisy estimates ($\hat{d}_j = d_j + \varepsilon_j$, independent, sd = 1 d_j); communities always know their own damages, so Regime V's price-taking limit is unaffected. Regimes I and III (and II and IV) coincide under known abatement costs. The monopoly $n_j=1$ row and the two uniform-partition rows are identical across columns, because market-power losses and the uniform partition do not depend on the regulator's estimates.

The second column introduces estimation error at a level at which the differentiated instruments lose their advantage over the uniform ones, the reversal documented by [Fowle and Muller \(2019\)](#). With the standard deviation of each community's error equal to its own damages, the differentiated instruments' expected loss of \$2.94 million exceeds the uniform instruments' \$2.88 million. In this setting, the best traditional policy is now to use undifferentiated taxes or cap and trade, as Proposition 6 predicts. The price-taking, monopoly, and uniform-partition rows are unchanged across the two columns, because none of them depend on community-level damage

estimates. At this level of damage estimation error, the welfare losses driven by the market distortions from our proposed mechanism exceed the losses due to estimation error associated with traditional differentiated policy instruments. The within-community partition at $\hat{n}_j$ mostly undoes these losses, since this version of the mechanism loses \$1.84 million in expectation relative to the first best, roughly a third less than the best traditional alternative.¹⁴

The two uniform-partition rows illustrate how the mechanism performs under the simplest implementation that maintains the need for no damage information. Giving every county the same number of seats leaves a loss of \$2.64 million at $\bar{n} = 20$ and \$2.65 million at $\bar{n} = 50$, in both columns of Table 3. Assigning twenty seats therefore removes 83 percent of the monopoly loss. Both uniform partitions outperform all four traditional instruments at this level of damage estimation error. Appendix H examines how the number of seats a uniform partition requires (and the associated loss) change with the scale at which communities are defined.

In the sections that follow, we characterize the distributional implications of each regime and show how the comparison changes in response to different levels of estimation error and different community definitions, and we report how the results change with the carbon-price anchor, the cleanup-cost anchor, the damage slope and value of statistical life, and the cost of carbon capture.

5.3 Distributional Implications


Figure 2 maps the welfare incidence of each regime for individual counties, combining avoided damages and permit revenue, following 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.¹⁵ 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 standard instruments, community gains consist of avoided damages alone and total \$3.4 to \$3.6 million per year statewide. Under the community permit mechanism, communities also collect the scarcity rents their permits command, and statewide community-side incidence rises by an order of magnitude, to roughly \$32 million per year under price-taking. These are community-side gains rather than social welfare. Because permit revenue is a transfer from firms, the mechanism’s large community gains result in the modest aggregate numbers of Table 3, which net all transfers out. Firms’ permit payments and abatement costs, which are ultimately borne by

¹⁴How far each community pulls its price back depends on $\theta_j = 1/(\hat{n}_j + 1)$, and the damping is slight at the median community, where θ_j is 0.009 at counties and 0.002 at places, so the typical community passes the regulator’s error through almost one for one. The aggregate behaves differently, however, because the communities that account for most of the welfare loss are not typical. Weighting each community by its contribution to the differentiated instruments’ expected loss, the mean θ_j is 0.21 at counties and 0.27 at places, and at counties $(1 - \bar{\theta})^2 = 0.618$ comes within a percentage point of the realized ratio of expected losses in Table 3, 0.625. The damping is largest in communities with a small $\hat{n}_j$, which arises when a community’s exposure is large relative to its residual demand, and that is the same condition that makes a pricing error in that community expensive.

¹⁵Even the undifferentiated regimes depict gains relative to this baseline because they embed information about damages from local co-pollutants, which the status-quo regime does not.

Figure 2. Community welfare incidence by county




Notes: Each panel shades county-level welfare incidence relative to the status quo, in which the carbon market operates with no local instrument; the dollar figure in each panel label is the statewide total. Darker shading indicates larger gains, hatching indicates losses, and the V-regime panels include permit revenue (a transfer from firms). Shading is log-scaled with a \$10,000 floor. Dots mark the 287 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 CalEnviro-Screen 4.0 tract geography.

owners and consumers, are not depicted. The patterns in Figure 2 are consistent with exposure to local co-pollutants, since gains concentrate in the South Coast, the Bay Area, and the southern Central Valley, where the covered facilities cluster.

Figure 3 splits the incidence between disadvantaged communities and other communities. The community permit mechanism generates substantially larger gains to disadvantaged communities than traditional policies, and these gains consist mostly of permit revenue rather than avoided damages. Put differently, the scarcity rents that other designs allocate to firms or to the state accrue here to the communities bearing the exposure (Fullerton, 2011; Fullerton and Metcalf, 2001). Strategic pricing strengthens rather than weakens this pattern, because markups are transfers to the selling communities. Monopoly pricing raises statewide community receipts to roughly \$104 million per year even as it destroys aggregate surplus. The uniform partition at $\bar{n} = 20$ sits between the two in aggregate, with community-side receipts of roughly \$33 million per year, but it divides them almost evenly between DAC and other counties, at about \$0.6 million per county in each group. The counties it overshoots are the high-damage counties in which price-taking concentrates its gains, while the markups that survive accrue to counties with low damages relative to demand, which is the pattern Corollary 2(a) implies. Hence market power in this mechanism reduces aggregate efficiency without necessarily diminishing the share of gains received by disadvantaged communities. Even though the largest gains to communities are in the form of revenue, the avoided-damage gains in disadvantaged communities (DACs) are also roughly twice as large under Regime V as under the standard instruments.

Figure 3. Incidence in DAC vs. non-DAC counties, by regime



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.

Figure 3 speaks to a tradeoff between efficiency and equity that has received increasing attention in recent years. Environmental markets might be designed in ways that cause efficiency and equity to compete (Kroetz et al., 2015), or in ways that address equity only after the fact, through the redistribution of revenues. In our mechanism the two objectives are aligned by construction, because the equity instrument also reveals efficiency-enhancing information. The features that deliver the first best are also the ones that concentrate gains in disadvantaged communities. An important caveat is that our incidence accounting is partial equilibrium, but general-equilibrium channels can shift burdens across factors and products (Fullerton and Heutel, 2007, 2010).


5.4 Damage Uncertainty and Community Scale

The results in the preceding two subsections are specific to the chosen level of estimation error (1 s.d. of d_j) and the use of counties to define communities, but Propositions 6 and 8 imply that the ranking of instruments depends on both the size of the estimation error and the scale at which communities are defined. Accordingly, we now demonstrate how varying each margin affects the relative and absolute performance of each policy regime. Figure 4 maintains the county definition of communities and varies the error from zero to 1.5 times damages, while Figure 5 fixes the error at 1 standard deviation of damages and compares tract, place, and county community definitions.

Figure 4 plots expected welfare losses as damage-estimation error grows, using the closed forms of Propositions 6 and 8 with independent errors scaled by the horizontal axis. Three of the six curves do not move as estimation error grows, beginning with the perfectly competitive version of the community permit system, which achieves the first best at any error level because it does not require the regulator to estimate damages in the first place. At the other extreme, monopoly pricing with no adjustment results in welfare losses of \$15.9 million per year. As with the perfectly competitive version of community pricing, this equilibrium does not depend on the regulator’s damage estimates. Between these two extremes lies the uniform partition at $\bar{n} = 20$, whose loss of \$2.64 million is likewise independent of the estimates.

The upward-sloping curves rank the information-dependent designs, starting with the differentiated instruments, which pass through estimation error one for one, so their expected losses rise quadratically and cross the nearly flat uniform line at $CV^* \approx 1.0$. This identifies the exact point at which the Fowlie and Muller (2019) reversal occurs in our calibration. The within-community partition $\hat{n} = n^*(\hat{d})$ is also subject to estimation error, but those errors are dampened by the fact that communities, not the regulator, set their prices. Because of this, the $\hat{n}$ losses rise more slowly and stay below the differentiated instruments’ at every error level in the figure. This means that a regulator confident enough in its estimates to differentiate taxes would do strictly better using those estimates to implement the community permit regime with the $\hat{n}$ partition, regardless of estimation error. The comparison between our mechanism and the uniform instrument, however, depends on the size of the estimation error, because beyond an error level of roughly 1.3 the uniform tax or cap-and-trade performs better than every estimate-dependent design. This follows the same logic underlying the Fowlie and Muller (2019) results, which show that when estimates

Figure 4. Instrument performance as a function of estimation error




Notes: Each curve is the closed-form expected welfare loss relative to the constrained first best (\$M per year; communities are counties) when the regulator acts on estimates $\hat{d}_j = d_j + \varepsilon_j$, with independent mean-zero errors whose standard deviation equals the horizontal-axis share of d_j . The uniform partition at $\bar{n} = 20$ and the monopoly outcome are flat because neither depends on the regulator's estimates.

are bad enough, the best policy uses them as little as possible.

Figure 5 fixes the estimation error at the level from Table 3 and varies how we define a community. The finest definition we consider is the census tract, the geography at which CalEnviroScreen designates disadvantaged communities, and we take Census places (i.e. cities or towns) as an intermediate choice. The undifferentiated instruments are insensitive to this choice because a single price averages over the errors no matter how finely they are distributed. The expected losses from these instruments are roughly \$2.9 million per year whether communities are counties, places, or tracts.

With differentiated instruments, the choice of how to define a community makes an important difference. Expected losses from differentiated taxes or cap-and-trade fall from \$2.94 million per year at the county definition to \$0.8 million with places and \$0.03 million with tracts. These instruments perform better with a finer community definition because we have assumed that estimation errors are uncorrelated across communities, so those errors diversify as the geography grows finer. We do not examine correlated errors, which would diversify less and narrow this advantage. At the same time, the errors themselves are being applied to smaller and smaller populations, so the welfare implications of getting a single community's estimate wrong are dampened. The reversal depicted at the county level in Figure 4 therefore implies that the Fowlie and Muller (2019) caution against differentiating on noisy estimates is specific to the scale at which communities are defined. The caution applies when communities are coarse enough that one bad estimate distorts the price faced by a large population, and it does not arise at the tract geography where disadvantaged communities are actually designated in practice.

Figure 5. Expected welfare losses at different community definitions



Notes: Markers report closed-form expected welfare losses relative to the constrained first best (log scale) at estimation error $sd = 1.0 d_j$, independent across communities, for three community definitions. Regime V price-taking is exactly zero at every definition and therefore does not appear on the log axis.

The community-permit regimes respond to scale differently because community size affects

both estimation error and market structure. Once again, the competitive version of the mechanism achieves the first best and is therefore invariant to community scale. The welfare implications of Cournot pricing with $\hat{n}$ sub-entities are non-monotone in community size. For instance, expected losses are \$1.84 million per year with county communities, roughly \$0.5 million with places, and roughly \$30 million with tracts. Fine geography helps the mechanism on the information margin just as it helps the differentiated instruments, but it also multiplies the number of markets in which market power operates, so that at the tract scale the market-power cost dominates. This suggests that optimal community scale varies according to the policy regime. At the tract level, expected losses under the differentiated instruments are well below those under the community permit mechanism, while at the county scale the ranking reverses and the differences across regimes narrow.

The performance of the uniform partition also depends on community scale, which we explore in Appendix H. Finer communities give each seller more market power, so the best uniform $\bar{n}$ rises from 20 to 777 seats as the number of communities grows (implying ever-smaller decision-making bodies). The loss that remains after partitioning is largely unaffected, staying between \$1.3 and \$1.6 million at every granularity.

These comparisons all depend the simplifying assumption that the researcher knows what damages actually are, and our calibrated d_j could be wrong in the same ways the regulator's estimate could. Damage estimates play two roles in the analysis, entering first as policy inputs, where Figures 4 and 5 model their error explicitly, and second as the benchmark against which the welfare losses themselves are measured. Table 6 varies the value of statistical life, which scales every d_j proportionally, and shows that misestimation on our part changes the reported magnitudes for every regime but not symmetrically, because the community permit system does not depend on damage estimates at all, in either its perfectly competitive or its monopoly form. Hence the welfare *levels* reported in this section are conditional on the damage model, whereas the *ranking* of estimate-dependent versus estimate-free designs is not. We also note that other papers studying differentiated policy suffer from the same issue.

5.5 Robustness and Discussion

The results above depend on a calibrated level for the carbon-abatement cost curve, which we set so that the model reproduces the 2019 mean allowance price. Table 4 reports the results of re-anchoring the curve to both the floor and the high end of the 2019 auction band to assess the sensitivity of our results to this choice. The key outcomes of interest are largely unaffected as we vary the carbon price. The first-best gain from adding community permits ranges from \$3.95 to \$3.99 million per year, the monopoly wedge stays within two percent of \$15.9 million, the uniform-partition loss at $\bar{n} = 20$ stays at \$2.64 to \$2.65 million, and the mechanism's expected-loss ratio relative to the differentiated instruments is 0.62–0.63 in every column. There are two features of the setting that account for this insensitivity. The first is that the auction band itself is narrow enough to generate only about \$2 of variation in the carbon price. The second is that this anchor

sets the level of the aggregate allowance price, but the comparison across regimes is driven by the ability of each regime to price local co-pollutants, which does not depend on the carbon price.

Table 4. Sensitivity to the carbon-cost anchor

Carbon-price anchor (\$/tCO ₂ e)	Reserve floor	2019 mean	2019 high
	15.62	16.84	17.45
τ^A at status quo	15.62	16.84	17.45
First-best gain vs. SQ (\$M/yr)	3.99	3.96	3.95
Monopoly wedge, $n_j=1$ (\$M/yr)	15.74	15.93	16.01
Uniform-partition loss, $\bar{n}=20$ (\$M/yr)	2.65	2.64	2.64
$\hat{n}$ dominance ratio vs. II/IV	0.624	0.625	0.626

Notes: County partition. Each column rescales the carbon-abatement cost level so that the status-quo carbon price equals the stated 2019 auction value. The first-best gain is the welfare gain from adding community permits; the other rows are welfare losses relative to the constrained first best, with the dominance ratio equal to the $\hat{n}$ mechanism's expected loss divided by the differentiated instruments' at estimation error $sd = d_j$.

The intercepts of the of the cleanup-cost curves are the other calibrated level. Table 5 reports the results of re-solving the model at the factors implied by the 2018 and 2020 credit vintages traded in 2019, 8.2 and 44.4, in place of the 19.7 implied by the 2019 vintage. The first-best gain falls from \$7.0 million at 8.2 to \$2.0 million at 44.4, because steeper cleanup costs makes local abatement more expensive. The monopoly wedge rises from \$5.7 to \$35.7 million and the mechanism's expected-loss ratio relative to the differentiated instruments rises from 0.42 to 0.77. The mechanism's advantage under estimation error is therefore largest when cleanup is cheap. The point at which undifferentiated instruments dominate differentiated ones moves only from 0.97 to 1.02 standard deviations across the range we consider. The uniform-partition loss at $\bar{n} = 20$ is lowest at the 2019 vintage because the best uniform size rises with the cost level (Appendix H).

Table 5. Sensitivity to the cleanup-cost anchor

	2018 vintage	2019 vintage	2020 vintage
RECLAIM credit price (\$/ton NO _x)	2,261	5,410	12,190
Cleanup-cost scale s	8.2	19.7	44.4
τ^A at status quo	16.84	16.84	16.84
First-best gain vs. SQ (\$M/yr)	6.99	3.96	1.99
Monopoly wedge, $n_j=1$ (\$M/yr)	5.71	15.93	35.74
Uniform-partition loss, $\bar{n}=20$ (\$M/yr)	3.09	2.64	4.27
$\hat{n}$ dominance ratio vs. II/IV	0.418	0.625	0.770
CV* (Fowlie–Muller reversal)	0.97	0.99	1.02

Notes: County partition. Each column rescales the cleanup-cost slope by the factor s implied by the stated RECLAIM credit price (the 2018, 2019, and 2020 vintages traded in 2019; the 2019 column is the paper's calibration), with the carbon-cost level re-solved so the status quo clears at \$16.84 in every column. Rows are defined as in Table 4. CV* is the estimation error at which the undifferentiated and differentiated instruments have equal expected loss.

Finally, in Table 6 we report the sensitivity of the assumptions used to inform the damage side

of the calibration. Raising the value of statistical life from \$7 million to the \$11.6 million central value and to \$14 million scales every community's damage slope by 1.7 and by 2. The losses from the traditional instruments scale with the square of damages, so the undifferentiated instruments' perfect-information loss rises from \$2.6 million to \$7.2 and \$10.5 million. In contrast, the the monopoly wedge falls slightly from \$15.9 to \$12.5 million, because higher damages raise each community's own valuation of exposure along with its markup. The mechanism's expected-loss ratio relative to the differentiated instruments falls from 0.63 to 0.48 and 0.43, so the mechanism's advantage under estimation error widens as damages rise, and the reversal point does not move, because both traditional losses scale by the same factor.

Table 6. Sensitivity to the damage slope and the value of statistical life

Value of statistical life (\$M)	7.0	11.6	14.0
<i>Welfare loss relative to the constrained first best (\$M/yr; larger = worse)</i>			
I/III, perfect information (hotspot)	2.64	7.24	10.54
I/III, estimation error ($sd = d_j$)	2.88	7.90	11.51
II/IV, estimation error ($sd = d_j$)	2.94	8.06	11.75
V, monopoly, $n_j=1$	15.93	13.45	12.50
V, Cournot at $\hat{n}$, estimation error	1.84	3.84	5.02
V, Cournot, uniform $\bar{n}=20$	2.64	4.77	6.22
$\hat{n}$ dominance ratio vs. II/IV	0.625	0.476	0.428
CV* (Fowlie–Muller reversal)	0.989	0.989	0.989
First-best gain vs. status quo (\$M/yr)	3.96	9.69	13.11
Median n_j^*	112.4	62.3	47.6
Counties with $n_j^* < 1$ (of 58)	0	1	2

Notes: County partition. Each column rebuilds d_j with the [Krewski et al. \(2009\)](#) concentration–response slope and the stated value of a statistical life; the first column is the paper's baseline and reproduces Table 3. The carbon-cost level is unchanged across columns because the status-quo anchor involves no community damages. Rows are defined as in Tables 4 and 5. Counties with $n_j^* < 1$ are those where marginal damage exceeds the residual-demand intercept, so the efficient permit supply is zero and the solver sets $\hat{n}_j = 1$.

Figure G2 in Appendix G reports the distribution of n_j^* across counties. The median county requires 112 sellers, the range runs from a handful in rural counties to roughly one hundred thousand in the largest markets, and every county's optimum is greater than $n = 1$. These are resident-scale bodies, so the within-community partition of Section 4 could plausibly be implemented through existing layers of governance such as school districts or neighborhood associations.

Figure G3 in the same appendix explores how cheap carbon capture would affect the case for the local instrument. In the 2019 baseline no covered facility had installed such technology, so the joint abatement channel was the only one available to firms. Hence, in the calibration, *any* carbon abatement also reduces local co-pollutants as a byproduct. In the future, improvements in the cost of carbon capture could make the carbon-only abatement channel feasible, and Figure G3 explores the consequences. As capture costs fall from four times the engineering estimates of [National Energy Technology Laboratory \(2022\)](#) to one quarter of them, the co-benefit falls by more than eighty percent, because firms shift toward capturing CO₂ and engage in less abatement that

affect co-pollutants (changes in output or fuel). The welfare gain from the community-permit layer moves in the opposite direction, ending about four percent above its baseline value. The divergence follows from Lemma 2(d), because cheaper capture weakens the cost complementarity $|\gamma_i|$ that links the two layers. Carbon policy then does less incidental local cleanup, and more of the co-pollutant problem is left for the local instrument to solve. This comparative static formalizes the environmental-justice objection that carbon capture perpetuates local pollution burdens (Batres et al., 2021). The upshot is that a technology that decouples carbon from its co-pollutants raises the value of a dedicated local instrument.

Whether these magnitudes are economically meaningful depends on the point of comparison. The value of required carbon abatement alone is roughly \$480 million per year in California, making the welfare numbers in this section small by comparison. However, our mechanism is geared toward reducing the damages associated with local co-pollutants, not climate change. The relevant comparison is therefore to total baseline NO_x damages from all 287 covered facilities, which are roughly \$32 million per year. The competitive equilibrium’s \$4.0 million gain over the status quo captures about an eighth of the entire local damage baseline in pure efficiency terms, the \$15.9 million monopoly wedge (measured relative to the first best) can dissipate about half of it, and the \$32 million of community-side receipts under price-taking amounts to a transfer roughly the size of the entire baseline welfare cost. This asymmetry follows from the familiar point that efficiency losses from mispricing are Harberger triangles that are second order in the pricing error, while the permit payments that change hands are rectangles that are first order in the quantities priced (Harberger, 1964; Hines, 1999). Transfers should therefore dwarf deadweight loss in any calibration of this kind, which is why the distributional results may be of greater significance for policymaking than the efficiency comparison.

Our magnitudes are also consistent with previous literature that includes similar calibration exercises. The gains from spatial differentiation scale with the *dispersion* of damages rather than their level, which is why they are modest both in our setting and in Fowlie and Muller (2019)’s regional NO_x application, but become large at national scope where damage heterogeneity spans orders of magnitude (Muller and Mendelsohn, 2009; Hollingsworth et al., 2026). We also focus on a single co-pollutant and a single mortality channel, and adding others would raise the welfare magnitudes reported here.

6 Conclusion

This paper analyzes an environmental market design in which the communities that bear local damages own and price the permits governing their exposure. Firms operating under an aggregate cap must also hold community-specific exposure permits, communities receive endowments that exceed baseline exposure, and posted prices clear each local market. The design’s key advantage is informational, in that the benchmark competitive equilibrium of the stacked market attains the constrained first best without the regulator estimating any community’s damages. Endowing

communities with permits also creates the scope for market power, because when communities price strategically, monopoly markups and the stacking of complementary permits (Heller, 1998; Buchanan and Yoon, 2000) erode the informational gains and can result in a net welfare loss, so whether those distortions outweigh the informational advantage is an empirical question.

We calibrate the model California’s carbon market. Competitive community pricing would improve on carbon-only regulation by \$4.0 million per year relative to \$32 million in damage under carbon-only policy, while monopoly pricing would result in a net loss of roughly \$12 million per year. Relative to the constrained first best, the monopoly loss is \$15.9 million per year, roughly six times the cost of the undifferentiated policy that motivates spatial differentiation in the first place. Partitioning each community into n_j^* independent selling bodies that receive a common price introduces within-community free-riding that exactly offsets the between-community markups and restores the first best under certainty. Implementing this scheme in practice also requires damage estimates, yielding an expected loss that is roughly one third below that of differentiated taxes or cap-and-trade. Alternatively, the regulator could give every community the same number of seats, and in the calibration such a uniform partition removes 83 percent of the market-power loss.

Disadvantaged communities capture the largest share of welfare gains under every variant of the mechanism we analyze. Although the avoided damages are roughly twice as large under community permits as under standard policies, most of disadvantaged communities’ gains are in the form of permit revenue. The scarcity rents that existing designs allocate to firms or the state (Fullerton, 2011; Fullerton and Metcalf, 2001) accrue instead to the communities bearing the exposure. This alignment does not require restricting trade, an approach whose efficiency costs are documented by Kroetz et al. (2015), and it avoids the outcome Kokorsch and Benediktsson (2018) describe in Icelandic fisheries, where a market designed for aggregate efficiency shifted quota holdings away from the communities that depended on them. Ownership also alters the standing of communities in environmental markets, in that communities set prices rather than passively receive compensation, so the mechanism is consistent with procedural justice notions, which the environmental justice literature identifies as a demand distinct from, and often prior to, distributive outcomes (Schlosberg, 2007; Banzhaf et al., 2019).

Mechanism design offers a natural alternative, but existing schemes focus on the cost side of the problem rather than the damage side. When damages are known, elicitation mechanisms recover firms’ private abatement costs (Kwerel, 1977; Dasgupta et al., 1980; Duggan and Roberts, 2002; Montero, 2008), and unstructured communication can transmit some of that information (Farrell, 1987). In our setting the missing information held privately by communities is damages, and Myerson and Satterthwaite (1983) show that no budget-balanced mechanism can efficiently elicit private valuations from the parties who hold them. Assigning ownership avoids elicitation altogether, because a community that owns the permits reveals its damages through the price it sets for them.

There is some empirical precedent for the mechanism we examine in this paper, since community ownership of tradable environmental rights has previously been implemented in the rela-

tively limited context of rights-based fisheries management in Alaska. The Community Development Quota program has allocated a share of Bering Sea fishing quota to sixty-five western Alaska villages since 1992 ([Ginter, 1995](#)). CDQ groups moved over time from harvesting their allocations toward leasing them for royalties ([Haynie, 2014](#)), the analog of communities collecting permit revenue rather than withholding permits to effectively increase abatement. The program's evaluation record is largely descriptive ([National Research Council, 1999](#); [Szymkowiak and Himes-Cornell, 2018](#)), and benefits have been distributed unevenly across member communities ([Lavoie et al., 2019](#)). To our knowledge, the approach has not been seriously considered for the co-pollutants of CO₂ regulation or for any other environmental market.

The mechanism we examine here has a specific structure, but it is not unique to air pollution. Whenever regulators seek to address an aggregate externality that has heterogeneous impacts across space, community permit markets can target the spatial incidence of compliance with aggregate targets. The key requirements are a measurable mapping from regulated entities to affected places, and communities identifiable enough to hold and sell permits. Water-quality trading has these elements, as nutrient credits apply at a watershed scale, nutrient loadings damage particular intakes and shellfish beds in specific districts, and hydrologic routing can supply a transport mapping ([Hung and Shaw, 2005](#)). Groundwater management does as well, since it entails basin-wide extraction cap, drawdown at neighboring wells causes local externalities, the aquifer's response function provides a mapping, and overlying well owners stand in for the communities ([Ayres et al., 2021](#)). Habitat and wetland banking fit less exactly, because the local services lost at a particular site are harder to meter than concentrations ([Salzman and Ruhl, 2000](#)), although offset prices differentiated by site-specific ecological value have been proposed for wetlands ([Aronoff and Rafey, 2026](#)). While these settings need other forms for specifying how emissions map into damages, the idea of using community ownership to reveal information about damages remains the same.

References

- Agee, M. D., S. E. Atkinson, T. D. Crocker, and J. W. Williams (2014). Non-separable pollution control: Implications for a CO₂ 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.
- Berman, M. (2024). Long-term effects of group rights to fisheries: Evaluating the Western Alaska community development quota program. *PLOS ONE* 19(12), e0312682.
- 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.
- 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 Air Resources Board (2022). 2022 scoping plan update: Ab 32 ghg emissions modeling (e3 pathways). 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.
- Farrell, J. (1987). Information and the coase theorem. *Journal of Economic Perspectives* 1(2), 113–129.
- 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., M. Reguant, and S. P. Ryan (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy* 124(1), 249–302.
- 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. (2001). A framework to compare environmental policies. *Southern Economic Journal* 68(2), 224–248.
- Fullerton, D. (2011). Six distributional effects of environmental policy. *Risk Analysis* 31(6), 923–929.
- Fullerton, D. and G. Heutel (2007). The general equilibrium incidence of environmental taxes. *Journal of Public Economics* 91(3–4), 571–591.
- Fullerton, D. and G. Heutel (2010). The general equilibrium incidence of environmental mandates. *American Economic Journal: Economic Policy* 2(3), 64–89.
- 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.
- Gillingham, K. and J. H. Stock (2018). The cost of reducing greenhouse gas emissions. *Journal of Economic Perspectives* 32(4), 53–72.
- 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.
- 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.

- Lavoie, A., K. Sparks, S. Kasperski, A. Himes-Cornell, K. Hoelting, and C. Maguire (2019). Groundfish fishery and community development quota program: Alaska's community development quota program: A complex institution affecting rural communities in disparate ways. *Marine Policy* 108, 103560.
- 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 CO₂ from industrial sources. Technical report, U.S. Department of Energy, NETL.
- National Research Council (1999). *The Community Development Quota Program in Alaska*. Washington, DC: National Academies Press.
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.
- 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 CO₂. *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 (2021a). Annual reclaim audit report for the 2019 compliance year. Technical report, SCAQMD.
- South Coast Air Quality Management District (2021b). Best available retrofit control technology (BARCT) assessments for reclaim facilities, staff reports. Technical report, SCAQMD.
- Szymkowiak, M. and A. Himes-Cornell (2018). Fisheries allocations for socioeconomic development: Lessons learned from the Western Alaska community development quota (CDQ) program. *Ocean & Coastal Management* 155, 40–49.
- 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.
- Wilen, J. E., J. Cancino, and H. Uchida (2012). The economics of territorial use rights fisheries, or TURFs. *Review of Environmental Economics and Policy* 6(2), 237–257.
- 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,

$$\frac{\partial(a_i^A, a_i^L)}{\partial(\tau^A, \tau_i^L)} = (\nabla^2 A_i)^{-1} \equiv \begin{pmatrix} \psi_i & p_i \\ p_i & \tilde{\beta}_i \end{pmatrix}, \quad \psi_i, \tilde{\beta}_i > 0, \quad p_i = \frac{-\gamma_i}{\det \nabla^2 A_i} \geq 0, \quad (23)$$

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,

$$a_i^A = \psi_i \tau^A + p_i \tau_i^L, \quad a_i^L = \tilde{\beta}_i \tau_i^L + p_i \tau^A, \quad \tau^A(\sigma) = \frac{R - \sum_i p_i \sum_j \sigma_j h_{ji}}{\sum_i \psi_i}. \quad (24)$$

All further symbols are defined where they are introduced within each proof. We maintain the following regulatory conditions throughout: 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

$$\sigma_j^* = D_j'(c_j^*), \quad \tau^A = \lambda, \quad \tau_i^L = \sum_j \sigma_j^* h_{ji} = \text{MD}_i^*. \quad (25)$$

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 (25) 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

$$\tau_i^L = \sum_j \sigma_j h_{ji} = \sum_j D'_j(c_j) h_{ji}. \quad (26)$$

Substituting Equation (26) 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). $\square$

A.2 Proof of Proposition 2

Proof. (a) We establish existence as a fixed point of the best-response map. Fix (σ_{-j}, τ^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 (23) 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 so is the objective in Equation (8). The best response is therefore unique and continuous in (σ_{-j}, τ^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 $\tau^A(\sigma)$ therefore exists and is continuous in σ .

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,

$$\sigma_j + \sigma'_j(s_j) s_j = D'_j(c_j), \quad (27)$$

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 $\Pi(\sigma) = \sum_k [\sigma_k s_k(\sigma) - D_k(s_k(\sigma))]$ denote joint community profit. Differentiating with respect to σ_j gives

$$\frac{\partial \Pi}{\partial \sigma_j} = \underbrace{s_j + \mu_j \frac{\partial s_j}{\partial \sigma_j}}_{\text{own marginal profit}} + \sum_{k \neq j} \mu_k \frac{\partial s_k}{\partial \sigma_j}. \quad (28)$$

A coordinated seller sets Equation (28) 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 (33),

$$\frac{\partial s_k}{\partial \sigma_j} = - \sum_i h_{ki} \left[\tilde{\beta}_i h_{ji} + p_i \frac{\partial \tau^A}{\partial \sigma_j} \right] = - \left(\sum_i \tilde{\beta}_i h_{ji} h_{ki} - \frac{(\sum_i p_i h_{ji})(\sum_i p_i h_{ki})}{\sum_i \psi_i} \right). \quad (29)$$

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 (29) shows that the equilibrium own slope exceeds the fixed- τ^A slope by $(\sum_i p_i h_{ji})^2 / \sum_i \psi_i$, so community j 's own marginal profit at the Nash price is $\mu_j (\sum_i p_i h_{ji})^2 / \sum_i \psi_i \geq 0$ rather than exactly zero. Substituting into Equation (28) yields

$$\left. \frac{\partial \Pi}{\partial \sigma_j} \right|_{\text{Nash}} = \mu_j \frac{(\sum_i p_i h_{ji})^2}{\sum_i \psi_i} - \sum_{k \neq j} \mu_k \left(\sum_i \tilde{\beta}_i h_{ji} h_{ki} - \frac{(\sum_i p_i h_{ji})(\sum_i p_i h_{ki})}{\sum_i \psi_i} \right). \quad (30)$$

When the net cross-demand externalities are positive and outweigh the own feedback term, Equation (30) 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

$$\frac{\partial A_i}{\partial a_i^L} = \sum_j \sigma_j^C h_{ji} > \sum_j D'_j(c_j^*) h_{ji} = \text{MD}_i^*, \quad (31)$$

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 $\mathcal{L}^C > 0$. $\square$

A.3 Proof of Proposition 3

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

$$\Phi(\tau^A, \sigma) \equiv \sum_i a_i^A \left(\tau^A, \sum_j \sigma_j h_{ji} \right) - R = 0. \quad (32)$$

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 (32) yields

$$\frac{\partial\tau^A}{\partial\sigma_j} = - \frac{\sum_i p_i h_{ji}}{\sum_i \psi_i} \leq 0 \quad (33)$$

at every price vector. Integrating Equation (33) along the segment from $\mathbf{0}$ to σ gives $\tau^A(\sigma) \leq \tau^A(\mathbf{0})$. 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 (33) is zero, so the carbon price does not change. Under Assumption 2 the responses are linear and aggregate to Equation (24). $\square$

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*

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

and the equilibrium solves a linear system $\mathbf{G}\sigma^C = \mathbf{g}$. Under the diagonal-dominance condition in Equation (41), 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 (24) with $\tau_i^L = \sum_k \sigma_k h_{ki}$ yields

$$s_j = c_j^0 - \sigma_j \sum_i \tilde{\beta}_i h_{ji}^2 - \sum_{k \neq j} \sigma_k \sum_i \tilde{\beta}_i h_{ji} h_{ki} - \tau^A \sum_i p_i h_{ji}. \quad (34)$$

Solving Equation (34) for σ_j gives the linear residual inverse demand

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

We next derive the best response. Community j chooses s to maximize $\sigma_j(s) s - d_j s$, taking (σ_{-j}, τ^A) , and hence the intercept of Equation (35), as given. The objective is strictly concave, with

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

$$s_j^C = \frac{1}{2} \left[\left(c_j^0 - \sum_{k \neq j} \left(\sum_i \tilde{\beta}_i h_{ji} h_{ki} \right) \sigma_k - \left(\sum_i p_i h_{ji} \right) \tau^A \right) - \left(\sum_i \tilde{\beta}_i h_{ji}^2 \right) d_j \right]. \quad (36)$$

Substituting Equation (36) into Equation (35) 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 (24). We substitute it into the intercept of Equation (35) and separate the $k = j$ term of $\sum_i p_i \sum_k \sigma_k h_{ki} = \sum_k \left(\sum_i p_i h_{ki} \right) \sigma_k$, which yields

$$\text{intercept} = \underbrace{c_j^0 - \frac{(\sum_i p_i h_{ji})R}{\sum_i \psi_i}}_{\text{constant in } \sigma} - \sum_{k \neq j} \left(\sum_i \tilde{\beta}_i h_{ji} h_{ki} - \frac{(\sum_i p_i h_{ji})(\sum_i p_i h_{ki})}{\sum_i \psi_i} \right) \sigma_k + \frac{(\sum_i p_i h_{ji})^2}{\sum_i \psi_i} \sigma_j. \quad (37)$$

The own-price term in Equation (37) 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

$$2 \left(\sum_i \tilde{\beta}_i h_{ji}^2 \right) \sigma_j^C = \left(\sum_i \tilde{\beta}_i h_{ji}^2 \right) d_j + \text{intercept}. \quad (38)$$

Substituting Equation (37) into Equation (38) and collecting price terms gives $\mathbf{G}\boldsymbol{\sigma}^C = \mathbf{g}$ with

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

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

Finally, we establish uniqueness. Let

$$x_j \equiv \frac{(\sum_i p_i h_{ji})^2}{(\sum_i \tilde{\beta}_i h_{ji}^2)(\sum_i \psi_i)}, \quad (39)$$

so that $G_{jj} = 2 - x_j$. Positive definiteness of Equation (23) gives $p_i^2 < \psi_i \tilde{\beta}_i$. The Cauchy–Schwarz inequality then yields

$$\left(\sum_i p_i h_{ji} \right)^2 = \left(\sum_i \sqrt{\psi_i} \cdot \frac{p_i h_{ji}}{\sqrt{\psi_i}} \right)^2 \leq \left(\sum_i \psi_i \right) \left(\sum_i \frac{p_i^2 h_{ji}^2}{\psi_i} \right) < \left(\sum_i \psi_i \right) \left(\sum_i \tilde{\beta}_i h_{ji}^2 \right) \quad (40)$$

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

$$\frac{1}{\sum_i \tilde{\beta}_i h_{ji}^2} \left[\sum_{k \neq j} \left| \sum_i \tilde{\beta}_i h_{ji} h_{ki} - \frac{(\sum_i p_i h_{ji})(\sum_i p_i h_{ki})}{\sum_i \psi_i} \right| + \frac{(\sum_i p_i h_{ji})^2}{\sum_i \psi_i} \right] < 1 \quad \text{for all } j. \quad (41)$$

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 (41) 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. In words, each community's price depends more on its own demand than on all rivals' prices combined, which holds when receptor fields overlap only partially. A strictly diagonally dominant matrix is invertible, so the equilibrium $\sigma^C = \mathbf{G}^{-1} \mathbf{g}$ exists and is unique.

Equation (41) is sufficient, not necessary, and in the calibration it fails at every community definition, because atmospheric transport makes every receptor field overlap with every other. Appendix F describes how uniqueness is verified computationally in that case. $\square$

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

$$\sum_j d_j c_j = \sum_j d_j \sum_i h_{ji} (\bar{e}_i^L - a_i^L) = \underbrace{\sum_j d_j \sum_i h_{ji} \bar{e}_i^L}_{\text{constant}} - \sum_i a_i^L \text{MD}_i. \quad (42)$$

Up to a constant, welfare is therefore $W = \sum_i [\text{MD}_i a_i^L - A_i(a_i^A, a_i^L)]$ 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

$$W(a) = W(a^*) + \nabla W(a^*)' \Delta a - \frac{1}{2} \sum_i \Delta a_i' (\nabla^2 A_i) \Delta a_i, \quad (43)$$

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 (43) vanishes.

We now express the deviations in terms of price wedges. Differencing the firm conditions $(\nabla^2 A_i) a_i = (\tau^A, \tau_i^L)'$ and the planner conditions $(\nabla^2 A_i) a_i^* = (\lambda, \text{MD}_i)'$ gives $\Delta a_i = (\nabla^2 A_i)^{-1} (\Delta^A, \Delta_i^L)'$. Substituting into Equation (43) yields

$$\mathcal{L} = W(a^*) - W(a) = \frac{1}{2} \sum_i \Delta a_i' (\nabla^2 A_i) \Delta a_i = \frac{1}{2} \sum_i \begin{pmatrix} \Delta^A \\ \Delta_i^L \end{pmatrix}' (\nabla^2 A_i)^{-1} \begin{pmatrix} \Delta^A \\ \Delta_i^L \end{pmatrix}. \quad (44)$$

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

$$\mathcal{L} = \sum_i \left[\frac{\psi_i}{2} (\Delta^A)^2 + p_i \Delta^A \Delta_i^L + \frac{\tilde{\beta}_i}{2} (\Delta_i^L)^2 \right], \quad \Delta^A = - \frac{\sum_i p_i \Delta_i^L}{\sum_i \psi_i}. \quad (45)$$

The second equation in Equation (45) follows from differencing the two market-clearing conditions $(\sum_i \psi_i) \tau^A = R - \sum_i p_i \tau_i^L$ and $(\sum_i \psi_i) \lambda = R - \sum_i p_i \text{MD}_i$, both obtained by aggregating Equation (24). Substituting it into the first equation, the $(\Delta^A)^2$ terms contribute $(\sum_i p_i \Delta_i^L)^2 / (2 \sum_i \psi_i)$ and the cross terms contribute $-(\sum_i p_i \Delta_i^L)^2 / \sum_i \psi_i$. Their sum is the negative correction in Equation (10).

Finally, we sign the loss. Each firm's term in Equation (44), $\frac{1}{2} v_i' (\nabla^2 A_i)^{-1} v_i$ 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 (45) forces $\Delta^A = 0$. The loss is therefore zero exactly when $\Delta_i^L = 0$ for all i . $\square$

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 $\Delta_i^L = \sum_j d_j h_{ji} - \text{MD}_i = 0$ for all i . The second equation of Equation (45) 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 $\Delta_i^L = t^*(\hat{\mathbf{d}}) - \text{MD}_i$, 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 $\tau_i^L = \widehat{\text{MD}}_i$, so $\Delta_i^L = \widehat{\text{MD}}_i - \text{MD}_i = b_i$. The bias enters the local wedge one for one. Regimes I/III post $t^*(\hat{\mathbf{d}})$, which is linear in $\{\widehat{\text{MD}}_i\}$ and hence in $\hat{\mathbf{d}}$ (Appendix D). The bias therefore enters the wedge profile $\{t^*(\hat{\mathbf{d}}) - \text{MD}_i\}$ linearly. In either case, when $\gamma_i < 0$ for some i , the induced carbon wedge

$$\Delta^A = - \frac{\sum_i p_i \Delta_i^L}{\sum_i \psi_i} \quad (46)$$

is nonzero whenever $\sum_i p_i b_i \neq 0$. The cross terms $p_i \Delta^A \Delta_i^L$ in Equation (45) capture 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 $\mathcal{L}(\{b_i\}) \rightarrow \infty$ as $\zeta \rightarrow \infty$, while $\mathcal{L}(\{\Delta_i\})$ is unchanged. The comparison is therefore quantitative, and the calibration of Section 5 provides an illustrative example. $\square$

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

$$\mathcal{L} = \frac{1}{2}(\mathbf{\Delta}^L)^T N \mathbf{\Delta}^L, \quad N \equiv \text{diag}(\tilde{\beta}_i) - \frac{\mathbf{p}\mathbf{p}^T}{\sum_i \psi_i}, \quad (47)$$

where the diagonal of N captures 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 $\widehat{\text{MD}}_i$, 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 (47) 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, $\mathbb{E}[\boldsymbol{\varepsilon}^T M \boldsymbol{\varepsilon}] = \text{tr}(M\Sigma)$. Because $\mathbf{b}$ has mean zero, this equals

$$\frac{1}{2}\text{tr}(M\Sigma) = \frac{1}{2} \sum_i \tilde{\beta}_i \text{Var}(b_i) - \frac{\text{Var}(\sum_i p_i b_i)}{2 \sum_i \psi_i}. \quad (48)$$

(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

$$s = \frac{\mathbf{1}^T N \mathbf{b}}{\mathbf{1}^T N \mathbf{1}} = \frac{(\mathbf{H}N\mathbf{1})^T \boldsymbol{\varepsilon}}{\Lambda} = \mathbf{w}^T \boldsymbol{\varepsilon}. \quad (49)$$

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

$$\mathcal{L}(t) = \frac{1}{2}(t\mathbf{1} - \mathbf{y})^T N (t\mathbf{1} - \mathbf{y}), \quad y_i = \text{MD}_i. \quad (50)$$

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 $\mathbf{w}^T \Sigma \mathbf{w} = \text{Var}(t^*(\hat{\mathbf{d}}))$,

$$\mathcal{L}(t^*(\hat{\mathbf{d}})) = \mathcal{L}_0^{\text{undiff}} + \frac{1}{2} \Lambda s^2, \quad \mathbb{E}[\mathcal{L}] = \mathcal{L}_0^{\text{undiff}} + \frac{1}{2} \Lambda \text{Var}(t^*(\hat{\mathbf{d}})), \quad (51)$$

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, with equality only when $N^{1/2}\mathbf{b}$ is proportional to $N^{1/2}\mathbf{1}$, that is, when the bias is a uniform shift across sources up to directions that carry no loss. Taking expectations, $\text{tr}(M\Sigma) \geq \Lambda \mathbf{w}^T \Sigma \mathbf{w}$ for every Σ .

Finally, we scale the error. Under $\Sigma = \text{CV}^2 \Sigma_0$, 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 ,

$$\mathbb{E}[\mathcal{L}^{\text{undiff}}] - \mathbb{E}[\mathcal{L}^{\text{diff}}] = \mathcal{L}_0^{\text{undiff}} - \frac{1}{2} \text{CV}^2 [\text{tr}(M\Sigma_0) - \Lambda \mathbf{w}^T \Sigma_0 \mathbf{w}]. \quad (52)$$

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 $\text{CV} = 0$ whenever MD_i varies across sources. It therefore crosses zero exactly once, at

$$\text{CV}^{*2} = \frac{2 \mathcal{L}_0^{\text{undiff}}}{\text{tr}(M\Sigma_0) - \Lambda \mathbf{w}^T \Sigma_0 \mathbf{w}}, \quad (53)$$

which is the reversal documented by [Fowle and Muller \(2019\)](#).

(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 . $\square$

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, $\lambda = \text{SCC}$. From the planner's clearing condition $(\sum_i \psi_i) \lambda = R - \sum_i p_i \text{MD}_i^*$, the cap that delivers this is

$$R^* = \left(\sum_i \psi_i \right) \text{SCC} + \sum_i p_i \text{MD}_i^*. \quad (54)$$

Setting Equation (54) requires a forecast of the co-abatement term $\sum_i p_i \text{MD}_i^*$, which is a function of $\{d_j\}$.

We next evaluate the cap set under beliefs. The regulator sets $\hat{R} = (\sum_i \psi_i) \text{SCC} + \sum_i p_i \widehat{\text{MD}}_i$,

and the realized shadow price is

$$\lambda(\hat{R}) = \frac{\hat{R} - \sum_i p_i \text{MD}_i}{\sum_i \psi_i} = \text{SCC} + \frac{\sum_i p_i b_i}{\sum_i \psi_i}. \quad (55)$$

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 $\text{SCC} - \lambda(R)$, with $\lambda(R) = (R - \sum_i p_i \text{MD}_i) / \sum_i \psi_i$. Integrating from R^* to $\hat{R}$ yields

$$\mathcal{L}^{\text{cap}} = \int_{R^*}^{\hat{R}} (\lambda(R) - \text{SCC}) dR = \int_{R^*}^{\hat{R}} \frac{R - R^*}{\sum_i \psi_i} dR = \frac{(\hat{R} - R^*)^2}{2 \sum_i \psi_i} = \frac{(\sum_i p_i b_i)^2}{2 \sum_i \psi_i}, \quad (56)$$

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

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-entities. Each sub-entity ℓ chooses its own supply $s_{j\ell} \geq 0$, taking the supplies of the other sub-entities $S_{j,-\ell}$, the prices of other communities σ_{-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

$$\max_{s_{j\ell} \geq 0} \sigma_j (s_{j\ell} + S_{j,-\ell}) s_{j\ell} - \frac{d_j}{n_j} (s_{j\ell} + S_{j,-\ell}), \quad (57)$$

where $\sigma_j(\cdot)$ is the linear residual inverse demand expressed in Equation (35), with slope $-1/\sum_i \tilde{\beta}_i h_{ji}^2$. The objective in Equation (57) is strictly concave in $s_{j\ell}$, and its interior first-order condition is

$$\sigma_j + \sigma'_j(s_j) s_{j\ell} = \frac{d_j}{n_j} \iff \sigma_j - \frac{s_{j\ell}}{\sum_i \tilde{\beta}_i h_{ji}^2} = \frac{d_j}{n_j}. \quad (58)$$

Every sub-entity's condition in Equation (58) determines the same offer, so any interior equilibrium is symmetric. Summing Equation (58) over ℓ gives $n_j \sigma_j - s_j / \sum_i \tilde{\beta}_i h_{ji}^2 = d_j$. Substituting demand from Equation (35) yields the price in the proposition. Setting $n_j = 1$ recovers the unitary best response of Lemma 1.

The partition n_j^* is the number of seats 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, $\sigma_{-j} = \mathbf{d}_{-j}$, and the carbon price to its first-best value λ .

We now verify that $\sigma = \mathbf{d}$ is the joint equilibrium at $\mathbf{n}^*$. At that profile, the carbon price from

Equation (24) is

$$\tau^A(\mathbf{d}) = \frac{R - \sum_i p_i \sum_j d_j h_{ji}}{\sum_i \psi_i} = \frac{R - \sum_i p_i MD_i}{\sum_i \psi_i} = \lambda, \quad (59)$$

where the middle step exchanges the order of summation. By Equation (59), 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^* . Hence $\sigma = \mathbf{d}$ is a fixed point of the joint system.

Finally, we establish uniqueness. Under Equation (41) 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 (39), 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. $\square$

A.10 Proof of Corollary 2

Proof. The corollary applies Proposition 7 to a partition that gives every community the same number of seats $\bar{n}$. Throughout, we write

$$I_j \equiv c_j^0 - \left(\sum_i p_i h_{ji} \right) \tau^A, \quad B_j \equiv \sum_i \tilde{\beta}_i h_{ji}^2, \quad (60)$$

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 partition.

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

$$\sigma_j(\bar{n}) = \frac{d_j + I_j/B_j}{\bar{n} + 1}. \quad (61)$$

By the definition of $n_j^*(\bar{n})$ in the statement, which is Equation (17) with the cross terms dropped and the carbon price at τ^A , we have $I_j/B_j = n_j^*(\bar{n}) d_j$. Substituting into Equation (61) gives Equation (21). The markup is

$$\mu_j(\bar{n}) = \sigma_j(\bar{n}) - d_j = d_j \frac{n_j^*(\bar{n}) - \bar{n}}{\bar{n} + 1}, \quad (62)$$

whose magnitude falls at rate $1/(\bar{n} + 1)$ and whose sign matches the sign of $n_j^*(\bar{n}) - \bar{n}$. Community j therefore prices above marginal damage when it has fewer seats than $n_j^*(\bar{n})$ 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 $\bar{n} = n_j^*(\bar{n})$ 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 partition 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 (62) gives the unitary markup $d_j (n_j^*(\bar{n}) - 1)/2$. For a community with $n_j^*(\bar{n}) \geq \bar{n}$, the ratio of the markup at $\bar{n}$ to the unitary markup is

$$\frac{\mu_j(\bar{n})}{d_j (n_j^*(\bar{n}) - 1)/2} = \frac{2}{\bar{n} + 1} \cdot \frac{n_j^*(\bar{n}) - \bar{n}}{n_j^*(\bar{n}) - 1} \leq \frac{2}{\bar{n} + 1}, \quad (63)$$

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 $\mathbf{H}$, $\{\psi_i, p_i, \tilde{\beta}_i\}$, $\{d_j\}$, R , and the integer $\bar{n}$. 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 (61) is the unitary best response of Lemma 1. As $\bar{n} \rightarrow \infty$, Equation (21) gives $\sigma_j(\bar{n}) \rightarrow 0$ 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 (62). When the partition overshoots, $\bar{n} \geq n_j^*(\bar{n})$, the price is below marginal damage and $|\mu_j(\bar{n})| = d_j (\bar{n} - n_j^*(\bar{n})) / (\bar{n} + 1) < d_j$. When it undershoots, $\bar{n} < n_j^*(\bar{n})$, the price is above marginal damage, the markup is decreasing in $\bar{n}$, and its largest value is the unitary markup at $\bar{n} = 1$. $\square$

A.11 Proof of Proposition 8

Proof. The regulator chooses the partition $\hat{n}$ 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

$$B_j \equiv \sum_i \tilde{\beta}_i h_{ji}^2, \quad P_j \equiv \sum_i p_i h_{ji}, \quad \Psi \equiv \sum_i \psi_i. \quad (64)$$

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

$$(\hat{n}_j + 1) \sigma_j = d_j + \frac{\text{intercept}_j}{B_j}, \quad \text{intercept}_j = c_j^0 - \sum_{k \neq j} \left(\sum_i \tilde{\beta}_i h_{ji} h_{ki} \right) \sigma_k - P_j \tau^A. \quad (65)$$

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 (24) and collecting all price terms on the left, as in the proof of Lemma 1, turns Equation (65) into the linear system

$$(\mathcal{N} + F) \boldsymbol{\sigma} = \mathbf{d} + \boldsymbol{\kappa}, \quad \mathcal{N} \equiv \text{diag}(\hat{n}_j + 1), \quad \kappa_j = \frac{c_j^0 - P_j R / \Psi}{B_j}, \quad (66)$$

where the matrix F collects the spillovers,

$$F_{jj} = -\frac{P_j^2}{B_j\Psi}, \quad F_{jk} = \frac{\sum_i \tilde{\beta}_i h_{ji} h_{ki} - P_j P_k / \Psi}{B_j} \quad (k \neq j). \quad (67)$$

Under Equation (41), 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 $\sigma(\mathbf{d}) = (\mathcal{N} + F)^{-1}(\mathbf{d} + \kappa)$, an affine function of the realized damages.

We now take the limit in which spillovers are negligible, so that $\sum_i \tilde{\beta}_i h_{ji} h_{ki} \rightarrow 0$ for $j \neq k$ and $P_j^2/\Psi \rightarrow 0$. Then $F \rightarrow 0$, and the intercept in Equation (65) no longer depends on any community's price. Its value is therefore the same whether damages are $\mathbf{d}$ or $\hat{\mathbf{d}}$. The rule in Equation (18) chose $\hat{n}_j$ as the number of seats at which the price would equal $\hat{d}_j$ if damages were $\hat{\mathbf{d}}$, which by Equation (65) means $\text{intercept}_j/B_j = \hat{n}_j \hat{d}_j$. Substituting into Equation (65) gives

$$\sigma_j = \frac{d_j + \hat{n}_j \hat{d}_j}{\hat{n}_j + 1} = \theta_j d_j + (1 - \theta_j) \hat{d}_j, \quad \theta_j = \frac{1}{\hat{n}_j + 1}, \quad (68)$$

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 partition 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 (68) when spillovers are not negligible. Let $\sigma^\theta \equiv \mathcal{N}^{-1}(\mathbf{d} + \kappa - F\hat{\mathbf{d}})$, whose j -th entry is Equation (68) 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 (66) gives

$$\sigma(\mathbf{d}) - \sigma^\theta = (\mathcal{N} + F)^{-1} F (\hat{\mathbf{d}} - \sigma^\theta). \quad (69)$$

(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

$$\mathcal{L}(\mathbf{v}) = \frac{1}{2} \mathbf{v}^T M \mathbf{v}, \quad M = \mathbf{H} \text{diag}(\tilde{\beta}_i) \mathbf{H}^T - \frac{1}{\Psi} \mathbf{P} \mathbf{P}^T, \quad (70)$$

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

$\text{tr}(Q\Sigma)$ applied to Equation (70) gives

$$\mathbb{E}[\mathcal{L}^{\text{diff}}] = \frac{1}{2}\text{tr}(M\Sigma), \quad \mathbb{E}[\mathcal{L}^{\text{mech}}] = \frac{1}{2}\text{tr}(\bar{W}M\bar{W}\Sigma). \quad (71)$$

The source wedge produced by $\bar{W}\varepsilon$ is the damped bias b_i^θ of the statement, so the second expression in Equation (71) is Equation (20).

(c) Subtracting the two expressions in Equation (71),

$$\mathbb{E}[\mathcal{L}^{\text{diff}}] - \mathbb{E}[\mathcal{L}^{\text{mech}}] = \frac{1}{2}\text{tr}((M - \bar{W}M\bar{W})\Sigma), \quad (72)$$

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 (72) is $\sum_j \varsigma_j^2 M_{jj} \theta_j(2 - \theta_j)$. Every term is nonnegative, because $M_{jj} > 0$ and $\theta_j \in (0, 1)$, and the sum is strictly positive whenever some $\varsigma_j^2 > 0$. 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 $\mathbb{E}[\mathcal{L}^{\text{mech}}] = (1 - \bar{\theta})^2 \mathbb{E}[\mathcal{L}^{\text{diff}}]$ for any Σ .

(iii) With heterogeneous θ and correlated errors, dominance can fail, because Ξ need not be positive semidefinite. At $\theta = (0.1, 0.9)$, $\Xi = \begin{pmatrix} 0.19 & 0.91 \\ 0.91 & 0.99 \end{pmatrix}$ has determinant $0.19 \times 0.99 - 0.91^2 = -0.64 < 0$. For two communities, Equation (72) is

$$\text{tr}((M - \bar{W}M\bar{W})\Sigma) = \varsigma_1^2 M_{11}\Xi_{11} + \varsigma_2^2 M_{22}\Xi_{22} + 2\rho\varsigma_1\varsigma_2 M_{12}\Xi_{12}, \quad (73)$$

which is negative when $M_{12} > 0$, Ξ_{12} is large relative to Ξ_{11} , and the correlation ρ is sufficiently negative. The calibration evaluates this case numerically. $\square$

B Microfoundation: Three Abatement Margins

This appendix grounds Assumption 1 and discusses how it is applied in the calibration. The margin-level notation introduced here, $(\phi_i, c_i^g, c_i^c, c_i^g)$ and the margin quantities (a_i^g, a_i^c, a_i^g) , is used only in this appendix and in the calibration discussion of Section 5 and Appendix F. Say firm i has a *three-margin technology* if it abates through three underlying margins, each with quadratic cost: activity reduction a_i^g (cost $\frac{c_i^g}{2}(a_i^g)^2$), cutting aggregate emissions one-for-one and local emissions at the co-emission ratio $\phi_i \geq 0$; local cleanup a_i^c (cost $\frac{c_i^c}{2}(a_i^c)^2$), cutting local emissions only; and capture a_i^g (cost $\frac{c_i^g}{2}(a_i^g)^2$), cutting 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*

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

(b) $\gamma_i = 0$ if and only if $\phi_i = 0$ (separability), and shutting down cleanup and capture ($c_i^c, c_i^g \rightarrow \infty$) gives fixed proportions $a_i^L = \phi_i a_i^A$, the one-pollutant case, in which the composite marginal damage of activity is $\text{SCC} + \phi_i \sum_j d_j h_{ji}$ (with SCC the social cost of the aggregate pollutant, as in Corollary 1) and carbon-only regulation is exactly the undifferentiated instrument of Fowlie and Muller (2019); Appendix E develops this mapping in full. (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

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

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 $\partial q^* / \partial a^L = \phi c^c / (c^g + c^q + \phi^2 c^c)$,

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

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, $dx = D^{-1}B^T d\vartheta$ with $D = \text{diag}(G'', Q'', K'')$ evaluated at the optimum and $B = \begin{pmatrix} 1 & 1 & 0 \\ 0 & \phi & 1 \end{pmatrix}$, 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

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

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. $\square$

The margin parameters are not identified by $(\alpha_i^A, \alpha_i^L, \gamma_i)$ (three equations, four unknowns), but the calibration works in the opposite direction, measuring $(\phi_i, c_i^q, c_i^c, c_i^g)$ from emissions and engineering data and computing the cost parameters. Part (c) also has direct policy content, because as capture gets cheap, carbon policy stops delivering local co-benefits, which raises rather than lowers the value of a separate local instrument. This is the formal content of the environmental-justice objection to carbon capture (Batres et al., 2021), and we examine it quantitatively in Section 5.5.

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 (neither is built from cost data), and with $\tau_i^L = \text{MD}_i$ 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, which makes both regimes 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 connects this to the corresponding Fowle and Muller (2019) results.

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

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

Proposition 9 (Optimal partition, general form). *Maintain Assumption 1, weakly concave residual inverse demand, and interior planner supplies $s_j^* = c_j^* > 0$. Evaluate each community's residual demand at rivals' efficient prices $\sigma_k = D'_k(c_k^*)$ and the efficient carbon price λ , and allow n_j real-valued. Then, with*

$$n_j^* = 1 + \frac{(-\sigma'_j(s_j^*)) s_j^*}{D'_j(c_j^*)} \in [1, \infty),$$

(a) *within each community, the symmetric profile $s_{j\ell} = s_j^*/n_j^*$ is the unique within-community equilibrium of Equation (15) given the efficient environment, so the planner's allocation, priced at $\sigma_j = D'_j(c_j^*)$ for every j , is an equilibrium of the joint partitioned game and both strategic distortions vanish; (b) the within-community equilibrium price is strictly decreasing in n_j , so n_j^* is the unique partition with this property (monopoly pricing above it, commons pricing below 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 seats. 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-entity'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

$$\sigma(s) + \sigma'(s) s_{j\ell} = \frac{D'_j(s)}{n}. \quad (74)$$

Given the aggregate s , Equation (74) determines the same offer for every ℓ , since $\sigma'(s) \neq 0$, so any interior equilibrium is symmetric. Summing Equation (74) over ℓ gives the aggregate equilibrium condition

$$\Phi(s) \equiv \sigma(s) + \frac{\sigma'(s) s}{n} - \frac{D'_j(s)}{n} = 0, \quad (75)$$

and

$$\Phi'(s) = \sigma'(s) + \frac{\sigma'(s) + \sigma''(s) s}{n} - \frac{D''_j(s)}{n} < 0 \quad (76)$$

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 (75) at $n = n_j^*$. At rivals' efficient prices and $\tau^A = \lambda$, the price-taking equilibrium of Proposition 1 clears community j 's market at $(s_j^*, D'_j(c_j^*))$, so the residual demand passes through that point and $\sigma(s_j^*) = D'_j(c_j^*)$. Substituting into $\Phi(s_j^*) = 0$ and solving for n ,

$$D'_j\left(1 - \frac{1}{n}\right) = \frac{-\sigma'(s_j^*) s_j^*}{n} \iff n = 1 + \frac{(-\sigma'(s_j^*)) s_j^*}{D'_j(c_j^*)} = n_j^*. \quad (77)$$

At n_j^* the unique within-community equilibrium is therefore s_j^* at price $D'_j(c_j^*)$. 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 (75) implicitly in n . Since $\partial\Phi/\partial n = (D'_j(s) - \sigma'(s)s)/n^2 > 0$ and $\Phi'(s) < 0$ by Equation (76), $ds/dn = -(\partial\Phi/\partial n)/\Phi'(s) > 0$. 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 $D'_j = d_j$, and the efficient supply is $s_j^* = \tilde{Q}_j - d_j \sum_i \tilde{\beta}_i h_{ji}^2$, where $\tilde{Q}_j$ is the intercept of Equation (35) at the efficient environment. Substituting into Equation (77) gives $n_j^* = 1 + s_j^*/(d_j \sum_i \tilde{\beta}_i h_{ji}^2) = \tilde{Q}_j/(d_j \sum_i \tilde{\beta}_i h_{ji}^2)$, the closed form of Proposition 7. $\square$

The numerator of the fraction in n_j^* 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 seats 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 seats. Proposition 8 does not extend to this setting, because its argument relies on prices being linear in damages.

D Regime Derivations for Table 1

All entries evaluate Proposition 4 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 $\Delta_i^L = t - \text{MD}_i$. 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

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

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

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

is linear in t and solves to

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

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 $t^* = \sum_i \tilde{\beta}_i \text{MD}_i / \sum_i \tilde{\beta}_i$, the [Diamond \(1973\)](#) weighted average recovered in [Appendix E](#). Under beliefs, $t^*(\hat{\mathbf{d}})$ replaces MD_i with $\widehat{\text{MD}}_i$ (the formula is linear). At $b = 0$ the residual loss is the dispersion term $\frac{1}{2} \sum_i \tilde{\beta}_i (t^* - \text{MD}_i)^2$ net of the rebate correction, positive whenever MD_i varies across sources.

Regimes II/IV. Differentiated instruments computed from beliefs support $\tau_i^L = \widehat{\text{MD}}_i$, 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 $\Delta^A = -(\sum_i p_i \Delta_i) / \sum_i \psi_i \leq 0$, 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 makes the relation to [Fowlie and Muller \(2019\)](#) (henceforth FM) exact. 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 $i = 1, \dots, I$ 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: 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), under which source i must hold r_i permits per unit of emissions, with r_i proportional to δ_i . Their main analytical findings are: (FM-R1) only the differentiated tax requires no cost information. The remaining four are that (FM-R2, FM-R3, their eq. 11) the gains from differentiation under certainty grow with the dispersion of marginal damages and shrink with the slope of marginal abatement costs; (their eq. 8b) the optimal undifferentiated tax is the [Diamond \(1973\)](#) weighted average of marginal damages, weighted by the inverses of the marginal-cost slopes; (their eq. 14) 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; and (FM-R5, their eq. 15 and §2.3.3) under cost uncertainty the differentiated *tax* still weakly dominates the undifferentiated tax, while differentiated *trading* can be dominated by undifferentiated trading when the

cap was set on mistaken cost estimates, because the compliance ratios fix the relative allocation that was optimal only for the intended aggregate.

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 characterize how these difference 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

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

and with our $\mathbf{H}$ supply providing microstructure of their δ_i . (b) Fixed-proportions limit. When cleanup and capture shut down ($c_i^c, c_i^g \rightarrow \infty$ in the technology of 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 $\text{SCC} + \phi_i \text{MD}_i$, so FM's non-uniform marginal damage 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 $\phi_i \text{MD}_i$, 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 (24) 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, since quadratic local abatement cost $\frac{\alpha_i^L}{2}(a_i^L)^2$ has marginal-cost slope α_i^L , matching their $2\beta_i$, and total local damages are $\sum_j d_j c_j = \text{const} - \sum_i \text{MD}_i a_i^L$ (see the proof of Proposition 4), matching their $\sum_i \delta_i e_i$ with $\delta_i = \text{MD}_i$. (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 $\text{SCC} + \phi_i \text{MD}_i$ (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 $\text{SCC} \cdot \bar{E}^A$ to damages in every regime and differences out, leaving $\phi_i \text{MD}_i$ as the operative source-specific marginal damage. $\square$

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 $\tau_i^L = \text{MD}_i$ 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

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

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^* ,

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

which with two sources equals $(\delta_L - \delta_H)^2 / [4(\beta_L + \beta_H)]$, 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 $r_i \propto \text{MD}_i$ and an optimally chosen permit total implements the same source prices as the differentiated tax (in equilibrium marginal abatement costs are equalized per unit of damage, their eq. 7), so our Regime II is their ratio-based design; (f) among Regimes I–IV, only the differentiated tax requires no cost data (FM-R1), as recorded in the information column of Table 1.

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 $\mathcal{L}^{\text{undiff}} = \frac{1}{2} \sum_i (t^* - \text{MD}_i)^2 / \alpha_i^L$ 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

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

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 $\sum_i \text{MD}_i / \alpha_i^L$; these coincide exactly 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 $r_i \tau$, where τ is the single permit price; with $r_i = \text{MD}_i / \bar{\delta}$ for any positive normalization $\bar{\delta}$, the supported prices are $\tau_i^L = \tau \text{MD}_i / \bar{\delta}$, and choosing the permit total so that $\tau = \bar{\delta}$ delivers $\tau_i^L = \text{MD}_i$, the differentiated-tax prices, and hence the same allocation, with marginal costs proportional to damages across sources. (f) The differentiated tax posts $\widehat{\text{MD}}_i$, built from **H** and beliefs alone; the uniform designs require t^* or its supporting cap, which weight damages by cost slopes; and the quantity designs require cost data to choose totals, as the formulas above show. $\square$

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 $\widehat{MD}_i = \sum_j m_j h_{ji}$ under differentiation and $t^*(\mathbf{m})$ under uniformity): (a) for any damage realization $\mathbf{d}$, the realized gain from differentiation is*

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

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 certain damages, the differentiated tax remains ex post optimal for any realized costs (its rates require no cost data), so its gain over the undifferentiated tax is 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 realized-cost optimum conditional on the mis-set cap calls for a different allocation, so differentiated trading can be welfare-dominated by undifferentiated trading, which is their §2.3.3 mechanism and the cap-versus-tax asymmetry of Appendix C.

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

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

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's rates $\widehat{MD}_i$ involve no cost parameters, so for any realized costs the induced allocation solves the planner's problem at prices equal to (believed) marginal damages; with certain damages ($\widehat{MD}_i = MD_i$) it is ex post first-best, and its loss is zero for every cost realization, while the undifferentiated tax's loss is nonnegative (FM-R5). For the quantity designs, the realized-cost market clearing under a cap fixed at the ex ante optimum delivers a price level different from the intended one. Under differentiated trading, the compliance ratios continue to force $\tau_i^L \propto MD_i$ 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, which equates marginal costs across sources, can then be closer to the cap-constrained optimum, which is the case FM document empirically. $\square$

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, each vanishing continuously as $\gamma_i \rightarrow 0$: (a) the carbon-market rebate: every regime’s loss includes the correction $-(\sum_i p_i \Delta_i^L)^2 / (2 \sum_i \psi_i)$ of Equation (10), so uniform-instrument losses are no longer a pure weighted dispersion of marginal damages, and the optimal uniform price t^* is corrected toward the co-abatement-weighted mean (Appendix D); (b) an endogenous aggregate price: local policy moves the carbon price (Proposition 3), a channel with no counterpart in FM’s single market; (c) the cap-level channel: the optimal cap itself 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 $\gamma_i \rightarrow 0$ 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 $(\sum_i p_i b_i)^2 / (2 \sum_i \psi_i)$. Since $p_i = -\gamma_i / (\alpha_i^A \alpha_i^L - \gamma_i^2) \rightarrow 0$ and $\psi_i, \tilde{\beta}_i$ converge to $1/\alpha_i^A, 1/\alpha_i^L$ as $\gamma_i \rightarrow 0$, 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). $\square$

F Calibration Details

Measurement

Transport. We construct **H** from facility-specific InMAP runs (Tessum et al., 2017). For each source we perturb NO_x emissions by 1,000 tons, hold other precursors fixed, and record the change in ambient $\text{PM}_{2.5}$ at each of 9,109 census-tract receptors per ton emitted. We measure GHG quantities in metric tons and criteria pollutants in short tons, following the reporting conventions of the underlying data.

Damages. We estimate tract damages by multiplying 2020 Census population by a concentration-response slope for adult mortality (Krewski et al., 2009) and a \$7 million value of statistical life. Table 6 in Section 5.5 reports results with the central and high values of statistical life. The steeper slope of Di et al. (2017), twelve times the Krewski slope, puts marginal damage above the residual-demand intercept in 50 of 58 counties, so the efficient permit supply there is zero and the partition results of Section 4 do not apply; we do not use it. We apply a floor of one person in the thirty-two tracts with zero 2020 population so that every community’s damage slope is positive.

Co-emission ratios. We measure ϕ_i as the facility’s NO_x -to- CO_2 emissions ratio in CARB’s 2019 mandatory reporting data, winsorized at the 1st and 99th percentiles. We observe the ratio directly for 524 of 581 reporting facilities and assign sector medians to the remainder. Of the 287 facilities in the calibration sample, 265 have a measured ratio, 18 receive a sector median, and 4 are absent from the reporting data. We assign the last group no co-emission ratio, so their abatement costs are separable.

Cleanup costs. We estimate c_i^c in two steps. We first fit a through-origin quadratic to NO_x control-cost data that we assemble from EPA’s Control Strategy Tool and SCAQMD’s BARCT technology assessments (U.S. Environmental Protection Agency, 2020; South Coast Air Quality Management District, 2021b). The control measures in these sources are defined at the sector level and scaled to each facility by its baseline emissions, so the fit yields one cost-per-ton schedule for each of the seven sectors in the sample, and a facility’s slope is its sector’s schedule divided by its own emissions. This step fixes the shape of the schedules but not their level, and engineering menus of this kind understate the costs that firms reveal in permit markets (Fowlie, 2010; Fowlie and Muller, 2019).

We use revealed market prices to adjust the level (intercept) of the abatement cost curves. Fowlie and Muller (2019) do this by iterating a permit price until their engineering cost model reproduces the observed cap, which pins the intercepts of the cost curves based on abatement volumes actually associated with the program. We cannot replicate this exact approach for two reasons. First, their condition requires a baseline from which the observed price induced a known amount of abatement. Our baseline is 2019 emissions, which are the levels facilities chose while already facing the RECLAIM price. Relative to that baseline the abatement RECLAIM induced is built in, so the model contains no observed quantity of abatement that corresponds to the 2019 price. Second, because the schedules are sector constants, any single scale factor allocates abatement across sectors in the same proportions regardless of the price it targets. We therefore use the simpler condition that the 2019 credit price, \$5,410 per ton for NO_x trading credits in Southern California’s RECLAIM market (South Coast Air Quality Management District, 2021a), equals the marginal cost of full abatement in the largest sector, which contains 102 of the 267 facilities with positive NO_x emissions. That condition sets the factor at 19.7. At that level the sample as a whole would abate 43 percent of its NO_x at the 2019 credit price. The same price paired with the 2018 and 2020 vintages traded in 2019, \$2,261 and \$12,190 per ton, gives factors of 8.2 and 44.4, and Table 5 in Section 5.5 re-solves the model at both.

Activity costs. We set c_i^q from sector-specific anchors for the marginal cost of CO_2 abatement at 30 percent abatement, in dollars per ton, which we take from the sources cited in the text (Gillingham and Stock, 2018; Fowlie et al., 2016; California Air Resources Board, 2022). The anchors are 30 for electricity, 45 for cogeneration, 40 for oil and gas, 60 for other combustion, 75 for refining, 120 for cement, and the mean of the others for hydrogen. Because only the ratios of these anchors matter, we rescale their common level by a single factor, $\kappa = 0.350$, chosen so that the status-quo carbon market, solved with every firm’s abatement bounded by its emissions, clears at the 2019 allowance price. Equation (22) is the interior form of that condition. We use the 2019 auction band of \$15.62 to \$17.45 to bound the sensitivity analysis in Table 4.

Capture costs. We set c_i^g prohibitively high 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).

Solution and conventions

We compute every equilibrium in three nested steps. First, given a carbon price and a vector of community prices, each firm chooses abatement on each margin to minimize its cost plus permit payments, subject to the constraint that abatement on either margin lies between zero and its business-as-usual emissions. With quadratic costs this is a small quadratic program with bounds, and we solve it exactly. Second, given the community prices, we find the carbon price at which aggregate carbon abatement equals the cap, by successively halving a price bracket until abatement equals the cap. Third, in the strategic-pricing regimes, each community chooses its price given the other communities' prices and the carbon price. Its first-order condition on the linear residual demand of Lemma 1 has a unique solution, and if that solution is negative the community posts a price of zero instead. We iterate the third step across communities, re-solving the first two steps after each round, until no price changes.

The diagonal-dominance condition in Equation (41) fails at every community definition, because atmospheric transport makes every receptor field overlap with every other, so the lemma does not guarantee that the strategic pricing equilibrium is unique. We check uniqueness numerically by starting the iteration from 50 random price vectors drawn uniformly on $[0, 2 \max_j d_j]$ at the county and place definitions, and from 20 at the tract definition, for both the monopoly game and the game at $\hat{n}$. Every run converges, and the converged price vectors agree to within one part in 10^{11} of the largest price (Table G2 in Appendix G).


Three conventions apply throughout. We set the aggregate cap 30 percent below business-as-usual emissions and verify that it binds in every regime. We retain seven facilities that report zero GHG emissions, with a one-ton floor so that they match the rows of **H**. They represent a negligible share of carbon and are the only sources at their abatement bounds in the competitive equilibrium. In the uncertainty experiments, we draw each community's estimation error independently, with standard deviation equal to its damage slope d_j times the coefficient of variation reported with each exhibit.

G Additional Calibration Exhibits

Figure G1 supports the co-abatement assumption of Section 5, and Figures G2 and G3 support the robustness discussion in Section 5.5; the latter two are computed on the county calibration. Table G1 reports numerical checks of the closed forms in Appendices A, B, and E.


Each row of Table G1 recomputes one closed form or identity from the appendices by a method that does not use the closed form, on a random test economy, and reports the largest error observed. Table G2 reports the random-restart check of uniqueness on the calibrated economies.

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.

Figure G2. The distribution of n_j^* across counties



Notes: Histogram of n_j^* from Proposition 7 across the 58 county communities, log₁₀ scale. The dotted line marks the unitary corner $n_j^* = 1$; the dashed line marks the median of 112.

Figure G3. The capture counterfactual



Notes: Both series are indexed to the 2019 base, in which capture is unavailable. Moving right, CO₂ capture cost falls from four times the NETL engineering estimates ([National Energy Technology Laboratory, 2022](#)) to one quarter of them. The co-benefit series is the incidental NO_x reduction from carbon abatement; ΔW is the welfare gain of the price-taking community-permit regime over the carbon-only status quo.

Table G1. Numerical checks of the closed forms

Check	Where in the paper	Test economy	Tolerance	Max. abs. error
<i>Appendix A: propositions and lemmas (Sections 2–4)</i>				
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 (41), holds on the test economy	Lemma A.4	random 12-firm, 4-community	pass/fail	passed
Cournot linear system $\mathbf{G}\boldsymbol{\sigma} = \mathbf{g}$ equals best-response iteration from four starts	Lemma A.4; Prop. 2	random 12-firm, 4-community	10^{-7}	1.8×10^{-15}
Direct Nash solve with carbon feedback satisfies the Cournot first-order conditions, $n \in \{1, 3\}$	Lemma A.4; Prop. 2	random 20-firm, 15-community	10^{-6}	1.2×10^{-10}
Cournot 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}
Cournot welfare loss equals the loss identity at the Cournot 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

continued on next page

Table G1, continued

Check	Where in the paper	Test economy	Tolerance	Max. abs. error
Carbon market clears at the closed-form τ^A (relative clearing residual)	Prop. 3; Equation (45)	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 (45)	Prop. 4; Equation (45)	random 12-firm, 4-community	10^{-8}	8.9×10^{-16}
Partition n^* delivers $\sigma_j = d_j$ at the joint fixed point	Prop. 7	random 12-firm, 4-community	10^{-7}	2.2×10^{-16}
Partition n^* 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}
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-partition price $\sigma_j(\bar{n}) = d_j(n_j^*(\bar{n}) + 1)/(\bar{n} + 1)$, $\bar{n} \in \{2, 5, 20\}$ (relative)	Cor. 2; Equation (21)	block-diagonal 12-firm, 6-community	10^{-10}	3.8×10^{-16}
<i>Appendix B: microfoundation</i>				
Induced cost equals direct minimization over the three margins	Lemma B.1(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 B.1(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 (23)	random 20-firm, 15-community	10^{-12}	1.7×10^{-16}

continued on next page

Table G1, continued

Check	Where in the paper	Test economy	Tolerance	Max. abs. error
Margin-level first-order conditions equal the induced-cost firm response	Equation (24); Lemma B.1	random 12-firm, 4-community	10^{-8}	4.4×10^{-16}
Firm responses satisfy $\nabla A_i = (\tau^A, \tau_i^L)'$	Equation (24)	random 20-firm, 15-community	10^{-9}	1.7×10^{-13}
$ \gamma_i $ strictly increasing in c_i^g (200-point grid, every firm)	Lemma B.1(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 B.1(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 B.1, non-identification remark	random 12-firm, 4-community	10^{-9}	2.2×10^{-15}
<i>Appendix D: regime derivations</i>				
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 E.2(b)	random 20-firm, 15-community	10^{-10}	0
<i>Appendix E: mapping to Fowle and Muller (2019)</i>				
Diamond t^* at $\gamma = 0$; two-source eq. 8b and eq. 11 of FM	Result E.2(b)–(c)	random 12-firm, 4-community	10^{-8}	8.7×10^{-19}
Emissions equivalence; compliance ratios $r_i \propto MD_i$ support the differentiated-tax prices	Result E.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 E.3(a)	random 12-firm, 4-community	10^{-8}	1.7×10^{-16}
$\gamma \rightarrow 0$: every Table 1 entry converges to its FM counterpart	Result E.1; Result E.4	random 12-firm, 4-community	10^{-6}	5.7×10^{-7}
<i>Appendix F: calibration solvers and the status-quo anchor</i>				
Box-constrained price-taking solver reproduces the interior solver on interior economies (relative)	Appendix F, solution and conventions	random 20-firm, 15-community	10^{-6}	3.3×10^{-16}
Box-constrained strategic-pricing solver reproduces the interior solver on interior economies, $n \in \{1, 3\}$ (relative)	Appendix F; Lemma A.4	random 20-firm, 15-community	10^{-6}	6.4×10^{-10}
Carbon-only status quo with box-constrained responses clears the cap (relative residual)	Equation (22)	forced-corner 20-firm, 15-community	10^{-8}	1.1×10^{-15}

continued on next page

Table G1, continued

Check	Where in the paper	Test economy	Tolerance	Max. abs. error
Box-constrained status quo equals the interior closed form $R/\sum_i \psi_i$ on interior economies (relative)	Equation (22)	random 20-firm, 15-community	10^{-6}	2.1×10^{-16}
Carbon-cost level κ solved so the box-constrained status quo clears at the anchor price (relative)	Equation (22)	12-firm synthetic knobs	10^{-8}	1.6×10^{-15}

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 A.4 holds by construction (verified in the second row). Table G2 reports the random-restart uniqueness check on the calibrated economies.

Table G2. Uniqueness of the strategic pricing equilibrium under random restarts

Partition	Pricing game	Converged (of starts)	Max. price deviation		Max. welfare dev. \$/yr	Max. τ^A dev. \$/tCO ₂ e
			\$/unit exposure	rel. to $\max_j \sigma_j$		
County ($J = 58$)	monopoly, $n_j=1$	50 of 50	1.3×10^{-4}	2.4×10^{-14}	7.2×10^{-7}	2.7×10^{-14}
County ($J = 58$)	Cournot at $\hat{n}$	50 of 50	9.9×10^{-5}	3.0×10^{-13}	1.5×10^{-7}	8.5×10^{-14}
Place ($J = 529$)	monopoly, $n_j=1$	50 of 50	3.0×10^{-5}	1.1×10^{-14}	8.9×10^{-7}	1.4×10^{-14}
Place ($J = 529$)	Cournot at $\hat{n}$	50 of 50	6.1×10^{-6}	2.7×10^{-13}	6.0×10^{-8}	3.6×10^{-15}
Tract ($J = 9,109$)	monopoly, $n_j=1$	20 of 20	5.6×10^{-5}	6.8×10^{-14}	1.0×10^{-6}	1.1×10^{-14}
Tract ($J = 9,109$)	Cournot at $\hat{n}$	20 of 20	9.9×10^{-6}	3.6×10^{-12}	8.9×10^{-8}	3.6×10^{-15}


Notes: For each calibrated partition and each pricing game, the box-constrained strategic-pricing solver is started from 50 price vectors (20 at tract) drawn uniformly on $[0, 2 \max_j d_j]$ (twice the largest first-best price) in place of its default start $\sigma^0 = \mathbf{d}$; convergence is a relative sup-norm change below 10^{-8} on (σ, τ^A) within 500 damped sweeps. “Max. price deviation” is the largest sup-norm distance between any two converged price vectors; the welfare and carbon-price columns are the largest spread across converged runs. The $\hat{n}$ game uses $\hat{n}_j = \max\{n_j^*(\mathbf{d}), 1\}$. Every converged run reaches the same equilibrium to solver precision, consistent with the uniqueness statement of Lemma A.4. At the tract definition each game uses 20 starts, because one solve of the $\hat{n}$ game takes about two minutes.

H Community Scale and Uniform Partitions

This appendix examines how the number of seats a uniform partition requires, and the loss that remains after partitioning, change with the scale at which communities are defined. We begin with the county and place definitions of Section 4.3 and then consider a family of synthetic partitions. Figure H1 plots the welfare loss under a uniform partition of size $\bar{n}$ at counties in panel (a) and at places in panel (b), as well as the benchmarks of Table 3. At counties the loss falls from the monopoly level to a plateau that runs from roughly $\bar{n} = 20$ to $\bar{n} = 50$ at \$2.64 million, and then rises toward the status-quo loss as every price falls toward zero, which is the bracket that Corollary 2(d) describes. At places, where the median n_j^* is 535 rather than 112, the same shape

appears with the minimum near $\bar{n} = 300$ at \$1.17 million, so the overshoot occurs at a much larger $\bar{n}$. The uniform partition has lower loss than the undifferentiated instruments at counties for any estimation error above roughly 0.15 standard deviations of damages and than the differentiated instruments above about 0.95. The same comparison at places does not favor the uniform partition over the partition at $\hat{n}_j$ at any error level we compute, because the damping governed by θ_j strengthens as communities get finer.

Figure H1. Welfare loss under a uniform partition at two community definitions



Notes: Each panel plots the welfare loss relative to the constrained first best (\$M per year, log scales) when every community is divided into the same number $\bar{n}$ of selling bodies, alongside the status quo, the undifferentiated and differentiated instruments at estimation error $\text{sd} = 1.0 d_j$, and the partition at $\hat{n}_j$ at the same error. The uniform partition uses no damage information, so its loss is the same under certainty and under estimation error.

Counties and places give two points on the relationship between community size and the seats a uniform partition requires, but the tract partition cannot be solved at arbitrary $\bar{n}$. This is because the corner solver requires a dense system with 9,109 rows, whereas the tract entries in Figure 5 use the closed forms of Proposition 8, and California's counties are so unequal in population that a comparison of counties with places confounds granularity with the shape of the population distribution. We therefore build a family of synthetic partitions by re-aggregating the tract-level primitives into J population-weighted geographic clusters under the same aggregation rules used for counties and places, summing damages and population-weighting the rows of $\mathbf{H}$, which isolates community size from the identity of the jurisdictions. Table H1 reports seven granularities from 29 communities, with a median population near 855,000, to 1,500 communities with a median population near 21,000, and Figure H2(a) plots the loss curves behind it and Figure H2(b) the seat counts.


The market-power loss that a partition corrects grows as communities get finer, from \$23.4

Table H1. Uniform partitions at seven synthetic community granularities

Communities J	29	58	116	250	529	1,000	1,500
Median population (000s)	855	476	272	108	53	30	21
Median n_j^*	13	42	80	310	598	906	1,144
Monopoly loss	23.4	37.2	47.5	68.6	88.1	109.2	120.3
Best uniform $\bar{n}$	20	40	64	163	313	546	777
within 5% of best	18–23	32–45	56–83	140–189	313–371	546–657	777–946
Loss at best $\bar{n}$	1.47	1.56	1.58	1.63	1.56	1.43	1.33
Share of unitary loss remaining	6.3%	4.2%	3.3%	2.4%	1.8%	1.3%	1.1%

Notes: Each column re-aggregates the 9,109 census tracts into J population-weighted geographic clusters and solves the corner strategic-pricing equilibrium at every uniform partition size $\bar{n}$. Losses are in \$M per year relative to the constrained first best. The best uniform $\bar{n}$ minimizes the loss on a grid, and the range beneath it collects every $\bar{n}$ whose loss is within five percent of that minimum. The synthetic clusters have roughly equal populations, so their levels do not match the county and place partitions of Table 3; see the text.

Figure H2. How the uniform partition scales with community granularity



(a) Loss curves by community granularity

(b) Seats required by community granularity

Notes: Panel (a) plots the welfare loss (\$M per year, log scales) under a uniform partition of size $\bar{n}$, one curve per synthetic partition of California into J communities. Panel (b) plots the median n_j^* and the best uniform $\bar{n}$ against J , with the shaded band marking the range of $\bar{n}$ within five percent of the best loss.

million at 29 communities to \$120.3 million at 1,500. Subdividing a community shrinks its damage slope faster than it shrinks the demand it faces, which leaves every seller with more market power over a smaller constituency. The number of seats required rises with it, and rises faster, with the median n_j^* running from 13 to 1,144 across the same range while the best uniform $\bar{n}$ runs from 20 to 777, exceeding the median n_j^* at the coarsest partition and settling between one half and seven tenths of it from 250 communities onward. The set of $\bar{n}$ within five percent of the best loss widens in proportion, so a regulator working at a finer granularity needs a larger number of seats but does not need to choose that number more precisely.

The loss that remains after a well-chosen uniform partition is, by contrast, nearly flat in granularity, lying between \$1.33 million and \$1.63 million across a range over which the monopoly loss grows more than fivefold. The uniform partition accordingly removes 94 percent of the market-power loss at the coarsest partition and 99 percent at the finest. Community size affects the number of seats a uniform partition requires much more than it affects the loss that remains after partitioning. The choice of scale therefore depends on the cost of seating and administering more bodies rather than on efficiency at finer resolution.

The synthetic clusters have roughly equal populations by construction and real jurisdictions do not, so the levels in Table H1 do not line up with the paper’s geographies. At $J = 58$ the clusters give a unitary loss of \$37.2 million and a median n_j^* of 42, compared to \$15.9 million and 112 for California’s actual counties, because the state’s population is concentrated in a small number of large counties. The scale results should therefore be read as a statement about the shape of the relationship within the synthetic family, with counties and places as the real reference points. An intermediate real geography such as ZIP codes or the AB 617 community boundaries would strengthen the comparison, which we leave for future work.