Is Taxing the Externality Enough?*

Christopher Conlon
NYU Stern School of Business
cconlon@stern.nyu.edu

Nirupama Rao
University of Michigan Ross School of Business
nirurao@umich.edu

March 30, 2026

Abstract

A central prescription in public economics is to correct externalities by taxing the externality-generating attribute: ethanol content for alcohol, THC for cannabis, sugar for sweetened beverages, or CO_2 emissions for vehicles. This prescription implicitly assumes independent demands across products, an assumption that fails in any differentiated product market, because the marginal external damage from a product depends not only on its externality content, but also on the external damage generated by its substitutes. We illustrate the conditions under which a simple attribute-based tax is likely to be sufficient and characterize when it fails. We then present case studies for distilled spirits and automobiles that quantify the welfare cost of ignoring within-market substitution.

JEL Classification: H21, H23, L13, L66.
Keywords: Excise tax, Pigouvian tax, externality correction, differentiated products, diversion ratio, sin goods.

*Preliminary draft. Comments welcome.

1 Introduction

Governments routinely levy corrective excise taxes on products that generate negative externalities. The standard policy prescription, dating to Pigou, is straightforward: levy the tax proportional to the externality-generating attribute. This is exactly what policymakers do in many contexts by taxing: ethanol in alcoholic beverages, THC in cannabis, grams of sugar in sweetened drinks, or CO_2 per mile in vehicles. This principle is formalized by the "additivity property" of optimal taxation (Sandmo, 1975), which establishes that the optimal tax on an externality-generating good contains a Ramsey revenue-raising component plus a Pigouvian correction equal to the marginal external damage. Kopczuk (2003) shows this additivity result holds under very general conditions.

However, much of the literature treats each category as a single good with its own demand curve. This requires treating "Alcohol" or "Spirits" as a single good, and "cigarettes" as another. In practice, consumers don't purchase "goods" but rather "products," which can be highly differentiated. We can think about a bottle of whiskey as containing some amount of externality-generating ethanol, but also some degree of quality, branding, or flavor that a planner seeking to address the externality would not necessarily want to tax (unless they were also seeking to raise revenue). Most sin-good markets feature substantial product differentiation, and excise tax policy must confront the question of *which* products to tax and at *what rate* within a given market.

What is already well-understood is that within a category, products may differ significantly in the amount of external damage they generate. Carbonated beverages may have significantly different levels of sugar, cigarettes may differ in the amount of tar or other carcinogens, marijuana products may have significantly different levels of THC, and vehicles may generate very different amounts of CO_2 per mile. The textbook theory says that policymakers should levy corrective taxes that target the different levels of externalities across products, but in practice they often rely on coarser tax instruments. As we show in Table 2, most jurisdictions tax sugar-sweetened beverages by volume rather than sugar content; and while the federal government taxes spirits by ethanol content, most states tax distilled spirits using only volume (at least above some minimum alcohol content).

This paper makes a simple but underappreciated observation. In a differentiated product market, the marginal external damage reduction from raising the price of product j does not reflect only the damage caused by its attribute content e_j, but also the external damage of substitute products e_k, and the degree of substitution between products. Our substitution-adjusted measure, $D^e_{j\to 0}$, accounts for the fact that consumers who leave product j may substitute toward other products within the market rather than exit entirely (which is im-

plicitly assumed when demands are independent). When consumers switch from one spirit to another with similar ethanol content, the externality reduction from the price increase is negligible. This suggests taxing more heavily those products whose consumers are most likely to exit the market entirely.

The conventional prescription to set taxes $\tau_j \propto e_j$ is not a consequence of the additivity principle alone; it is a consequence of the additivity principle *plus* the assumption of independent demands across products within the market. Once we recognize that products are substitutes, the optimal Pigouvian correction depends on the full substitution structure: where consumers go when they leave each product, and what externalities those alternatives carry.

We derive the conditions under which a simple attribute-based tax (proportional to e_j) replicates the optimal product-specific correction, and organize the analysis around a taxonomy of three conditions that jointly determine when and how the naive tax fails. We then present case studies using estimated demand systems for distilled spirits and survey data for automobiles that quantify the welfare cost of ignoring within-market substitution.

2 Framework

Consider a market with J differentiated products and an outside option ($j = 0$). We assume that the externality is *atmospheric*: total external damage depends only on aggregate consumption of the externality-generating attribute, $E = \sum_j e_j q_j$, and is independent of who consumes the product. Each product generates externalities in proportion to an observable attribute (ethanol content, THC, sugar, CO_2), which we denote e_j.

The outside option may also carry externalities ($e_0 \geq 0$): for example, the outside option for legal cannabis includes the illicit market; the outside option for new vehicles includes keeping an older, dirtier car; and the outside options for drinking distilled spirits might include beer or wine.

We make these strong assumptions not because we think they are necessarily an accurate depiction of how externalities work in practice, but rather because we want to demonstrate that even in a highly simplified world, it is important to consider how consumers substitute across products when calculating the marginal external damage and corrective taxes.[1]

In Conlon and Rao (2025), we solve the problem of a social planner who can set product-specific prices in order to maximize consumer surplus net of production costs, subject to a

[1] We are ruling out externalities that depend on who consumes the product or on non-linear interactions. In reality, we may worry that young smokers lose more quality adjusted life years than old smokers, or that heavy drinkers generate more external damage per unit of ethanol than light drinkers because the external damage function is convex rather than linear (see Griffith et al. (2019) for the latter).

revenue requirement and an externality cap:

$$\max_{\mathbf{p}} \; CS(\mathbf{q}) - C(\mathbf{q}) + \lambda_r\big(\mathbf{p}\cdot\mathbf{q} - C(\mathbf{q}) - R\big) - \lambda_e\big(\mathbf{e}\cdot\mathbf{q} - \bar{E}\big), \tag{1}$$

where λ_r and λ_e are the shadow values of revenue and externality reduction.

The first-order conditions of this problem show that the optimal Pigouvian correction for product j is proportional to the **substitution-adjusted marginal externality**:[2]

$$D^e_{j\to 0}(p) = e_j - \sum_{k\neq j} D_{j\to k}(p)\cdot e_k - D_{j\to 0}(p)\cdot e_0, \tag{2}$$

This expression depends on the (incremental) external damage from the focal good e_j, the (incremental) external damage from the substitutes e_k including the outside option e_0, and the *diversion ratio* from j to k: $D_{j\to k}$–the fraction of consumers leaving j who switch to k (as the price rises). The diversion ratio is similar to the cross-elasticity in that it quantifies substitution across products, but has the advantage that it sums to one across all substitutes, so that diversion to the outside option is given by $D_{j\to 0} = 1 - \sum_{k\neq j} D_{j\to k}$.[3]

The expression in (2) captures the net reduction in aggregate external damage when the price of product j rises enough to divert one consumer away from j: some switch to substitutes (generating e_k), some exit the market (generating e_0), and the net effect can be much smaller than e_j. Using the definition of $D_{j\to 0}$ we can decompose (2) as:

$$D^e_{j\to 0} = \underbrace{D_{j\to 0}\cdot(e_j - e_0)}_{\text{extensive margin}} + \underbrace{\sum_{k\neq j} D_{j\to k}\cdot(e_j - e_k)}_{\text{intensive margin}}. \tag{3}$$

The *extensive margin* captures consumers who leave the market entirely, weighted by the externality gap between product j and the outside option. The *intensive margin* captures consumers who switch to other products within the market, weighted by the externality gap between j and the substitute. This decomposition clarifies why a naive attribute tax can fail: it ignores both the possibility that the outside option is "dirty" (reducing the extensive margin) and that within-market substitution may generate little externality reduction (shrinking the intensive margin).

[2] See Conlon and Rao (2025) for a full derivation of the pricing rule, which also includes Ramsey and market-power components.

[3] Mathematically, the diversion ratio is defined as: $D_{j\to k}(p) = (\partial Q_k/\partial p_j)/|\partial Q_j/\partial p_j|$ or using the own and cross-elasticities as: $\frac{\mathcal{E}_{jk}}{|\mathcal{E}_{jj}|}\cdot\frac{p_j}{p_k} = D_{j\to k}$.

2.1 When Does a Simple Attribute Tax Suffice?

A tax proportional to e_j replicates the optimal product-specific correction if and only if $D^e_{j\to 0} \propto e_j$ for all j. Inspecting (3), this proportionality holds when three conditions are jointly satisfied:

(i) **Clean outside option**: $e_0 = 0$, so the extensive margin simplifies to $D_{j\to 0} \cdot e_j$;

(ii) **Homogeneous externality content**: $e_k \approx e$ for all products k, so the intensive margin vanishes (consumers who switch products generate the same externality);

(iii) **Uniform diversion to the outside good**: $D_{j\to 0}$ is constant across products, so the extensive margin is the same fraction of e_j for every product.

These three conditions are properties of the demand system and the externality technology, not of market structure. Whether firms are perfectly competitive, Bertrand oligopolists, or a collusive monopoly, the formula for $D^e_{j\to 0}$ remains the same. At the same time, market structure does affect the *cost* of getting the tax wrong. Firms with market power set markups inversely proportional to demand elasticity, and greater market power means scaling both the production cost and the (corrected) externality cost $D^e_{j\to 0}$ by a larger amount.[4]

When all three hold, $D^e_{j\to 0} = \bar{D} \cdot e_j$ for some constant $\bar{D}$, and taxing proportional to e_j is optimal. When any condition fails, the naive attribute tax generically over- or under-corrects relative to the optimum.[5]

The level of $D^e_{j\to 0}$ relative to e_j is also informative. When the outside option is "dirty" $e_0 > 0$ or there is significant substitution to other dirty goods, the level of the optimal corrective tax may decline because raising the tax on j does little to reduce the external damage but still causes deadweight loss. For example, taxing home heating emissions at \$200/ton of CO_2 might be too high if there is a high-degree of substitution to untaxed wood-burning stoves.

3 Applications

Table 1 summarizes how the three conditions map onto four policy-relevant markets. We discuss each briefly here and present detailed case studies for distilled spirits and vehicles in Sections 3.1 and 3.4.

[4] Conlon and Rao (2025) find that higher quality products have higher markups but lower $D^e_{j\to 0}$– so this works against the externality correction; O'Connell and Smith (2024) find that ignoring market power in SSB taxation reduces welfare gains by 40 percent.

[5] One can quantify the welfare cost of using $\tau_j \propto e_j$ instead of $\tau_j \propto D^e_{j\to 0}$ using the framework of Jacobsen et al. (2016): the R^2 from a quantity-weighted regression of $D^e_{j\to 0}$ on e_j measures the fraction of the optimal welfare gain captured by the naive tax. We compute this statistic in the case studies below.

Table 1: Taxonomy of Condition Failures Across Markets

Market	**Condition (i)** $e_0 = 0$?	**Condition (ii)** $e_k \approx e$?	**Condition (iii)** $D_{j\to 0}$ uniform?	**Primary failure mode**
Distilled spirits†	Partially fails	Mostly holds	Fails strongly	Extensive margin
Cannabis	Fails strongly	Fails strongly	Fails	All three
Sweetened beverages	Partially fails	Moderate	Fails	(i) and (iii)
Vehicles / emissions†	Ambiguous	Fails strongly	Varies	Intensive margin
E-cigarettes / nicotine	Fails strongly	Fails strongly	Fails	(i) and (ii)

† Detailed case study in Sections 3.1 and 3.4.

3.1 Distilled Spirits

The main ingredient in distilled spirits that is associated with both health consequences (cirrhosis, cancer) and safety issues (domestic violence, drunk driving) is the alcohol itself (ethanol). Most distilled spirits are bottled at 40% ABV (80 proof), so condition (ii) approximately holds within the spirits category—though flavored products (20–30% ABV) and overproof products (50%+) introduce some dispersion.

In Conlon and Rao (2025), we leverage extensive data on product-level sales and pricing, including a 2011 change in the excise tax in Connecticut, as well as individual data on purchases by income. We estimate a demand system and recover estimates of $\widehat{D}_{j\to k}$, which we then use to compute $D^e_{j\to 0}$ for each product. We display our estimates in Figure 1.

Condition (i) fails because the outside option for spirits consumers may include beer and wine (where $e_0 > 0$), though this is moderated if the relevant outside option is non-alcoholic beverages – including new non-alcohol spirits substitutes. It is certainly possible to expand the definition of the market and to better measure substitution to beer, wine and non-alcoholic alternatives.[6] To highlight the larger issue with condition (iii), our calculation of $D^e_{j\to 0}$ from the data implicitly assumes that $e_0 = 0$ to shut down this channel.

The primary failure is in condition (iii): $D_{j\to 0}$ varies enormously across products. Conlon and Rao (2025) estimate that budget products (e.g., Dubra Vodka) have diversion ratios to the outside good exceeding 50%, while premium products (e.g., Belvedere, Grey Goose) have much lower diversion, as consumers substitute primarily to other premium spirits. The substitution-adjusted measure $D^e_{j\to 0}$ is highest for the cheapest products, suggesting that externality reduction is most effectively achieved by targeting these products.

In Figure 1, this leads to significant dispersion in $D^e_{j\to 0}$ *conditional on* e_j. Even though Dubra Vodka (the cheapest product) and Bulldog Gin are both 40% Alcohol by Volume (80 Proof), the marginal reduction in ethanol consumption is 60% larger for Dubra Vodka than

[6] Our work in Conlon et al. (2024) suggests that spirits drinkers dominate the heaviest drinking individuals.

Bulldog Gin. This implies that a uniform ethanol tax–even if set correctly on average – would still overtax premium products and undertax budget products relative to the optimum.

Similarly, we find that $D^e_{j\to 0} < 0$ for Malibu Coconut Rum. This means that as we raise the price and consumers switch away, they actually consume more ethanol in net as they switch to products with a greater alcohol content. This happens because Malibu contains a lot of sugar and very little alcohol even compared to other flavored rum products. A planner might respond not only by reducing taxes on Malibu rum, but possibly even by offering a small subsidy to encourage substitution from other products with larger values of e_j.

In Figure 1, we also consider a "tagging" exercise inspired by Jacobsen et al. (2016). We calculate the R^2 from a regression of our marginal external damage measure $D^e_{j\to 0}$ on both ethanol content e_j (which the classic Pigouvian tax would target) and prices, which an *ad valorem* sales tax would target. The bad news is that the sales tax would be almost perfectly uncorrelated with our marginal external damage measure, while the R^2 for the ethanol tax would be a respectable 0.88 (albeit lower if we weighted by product sales).

3.2 Cannabis

Cannabis is the market where the naive attribute tax is most likely to fail, because all three conditions are violated simultaneously.

Condition (i): $e_0 \gg 0$. The outside option for legal cannabis is not "abstinence" but a combination of the illicit market, alcohol, and other substances. The illicit market carries potentially *higher* externalities than the legal market, as unregulated products pose health risks and illicit distribution generates violence and incarceration costs. Miller and Seo (2021) document substantial substitution between legal cannabis and alcohol, finding that cannabis legalization reduces alcohol demand by approximately 15%. Han et al. (2024) show that excessive sin taxes push consumers back toward illicit markets. When e_0 is large, $D^e_{j\to 0}$ can be very small or even *negative*: raising legal cannabis prices may *increase* total harm by diverting consumers to worse alternatives.

Condition (ii): massive e_k dispersion. THC content varies by a factor of five or more across product categories. Low-potency flower contains roughly 15% THC, high-potency flower approximately 30%, and concentrates 60–90% THC. Hansen et al. (2020) study potency-based taxation in Washington State and find that demand for potency is relatively inelastic, so a naive THC tax has limited behavioral effect. The intensive margin term $\sum_k D_{j\to k}(e_j - e_k)$ is large because substitution between flower and concentrates is meaningful.

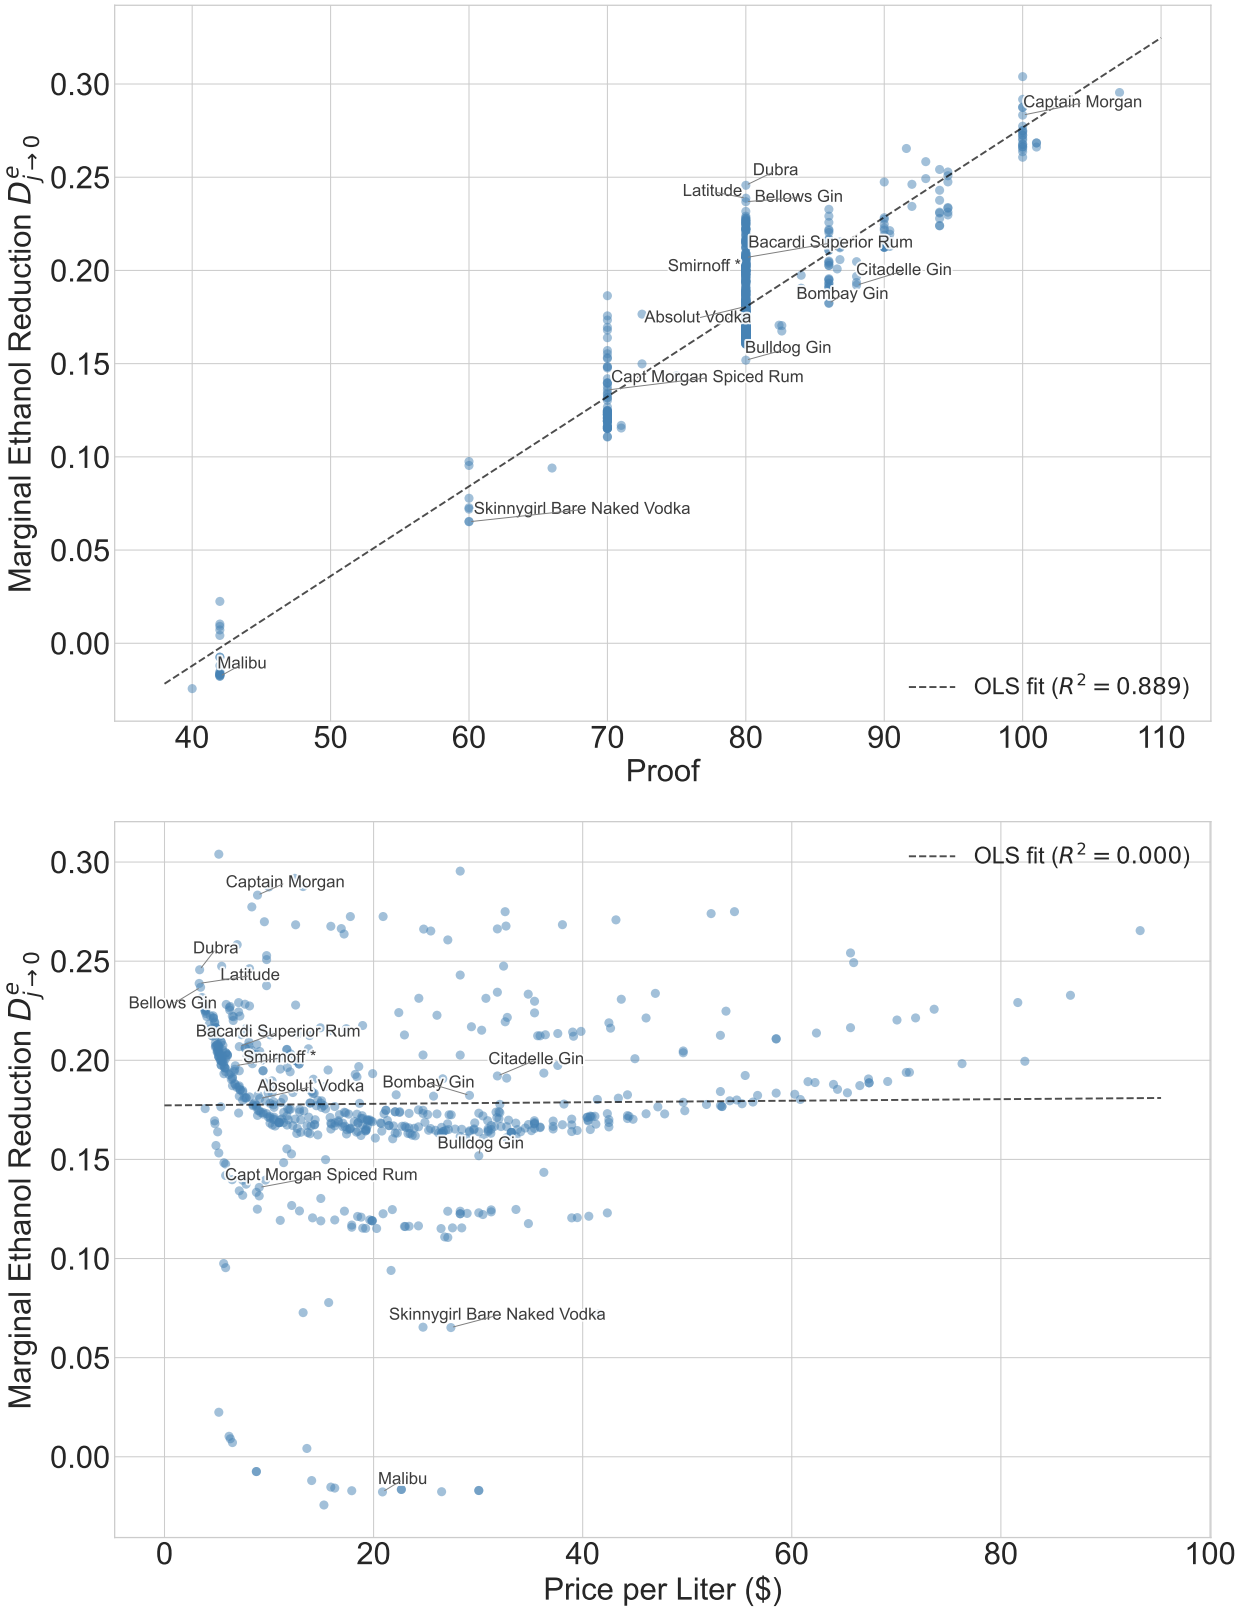


Figure 1: Marginal Ethanol Reduction for Spirits (Conlon and Rao, 2025)

Notes: $D_{j\to 0}$ calculated assuming that outside good ethanol is zero $e_0 = 0$ using the parametric demand estimates of Conlon and Rao (2025).

Condition (iii): heterogeneous $D_{j\to 0}$. Hollenbeck and Uetake (2021) show that entry restrictions in legal cannabis markets create substantial market power, with firms behaving as local monopolists. This generates the same markup-externality misalignment documented for spirits: firms set high markups on premium products and low markups on budget products, uncorrelated with THC content. The illicit market may also serve as a better substitute for certain kinds of products (loose flower) than for others (gummies).

The policy implication is striking: in cannabis, a well-designed tax schedule might involve *lower* taxes on some legal products than a naive THC tax would prescribe, in order to keep consumers in the regulated market and away from illicit alternatives where e_0 is high. It might also involve different tax rates for different kinds of products (with higher taxes on less substitutable baked goods and gummies).

3.3 Sugar-Sweetened Beverages

Here there is little disagreement that sugar content is the main source of the externality, and the case for corrective taxation has been extensively studied (Allcott et al., 2019a,b). A growing number of jurisdictions have implemented SSB taxes, but the designs vary dramatically in both structure and tax base, as summarized in Table 2.

Table 2: SSB Tax Designs Across Jurisdictions

Jurisdiction	Structure	Diet	100% Juice	Milk	Rate
Berkeley, CA (2015)	Flat volumetric	No	No	No	1¢/oz
Philadelphia, PA (2017)	Flat volumetric	Yes	No	No	1.5¢/oz
Cook County, IL (2017)*	Flat volumetric	Yes	No	No	1¢/oz
Boulder, CO (2017)	Flat volumetric	No	No	No	2¢/oz
San Francisco, CA (2018)	Flat volumetric	No	No	No	1¢/oz
Seattle, WA (2018)	Flat volumetric	No	No	No	1.75¢/oz
UK (SDIL, 2018)	Tiered by sugar	No	No	No	19.4–25.9p/L
Mexico (2014, ref. 2026)	Flat volumetric	Reduced	No	No	1.5–3.1 pesos/L
South Africa (2018)	Per gram of sugar	No	No	No	2.21c/g >4g/100mL

* Repealed after four months.

Notes: "Diet" indicates whether artificially sweetened beverages are taxed. "100% Juice" and "Milk" indicate whether these categories are included in the tax base. UK rates as of 2025–26; lower tier applies at 5–8g sugar/100mL, upper tier at 8g+/100mL. Philadelphia and Cook County are the only US jurisdictions to have included diet drinks. All US taxes are levied on distributors.

The most striking feature of existing SSB taxes is how much the tax base varies. Every jurisdiction exempts 100% fruit juice, despite the fact that orange juice contains comparable sugar per ounce to many taxed sodas. Most jurisdictions also exempt milk-based drinks and

sweetened coffee or tea purchased at cafes. These exemptions create precisely the condition (i) failure our framework highlights: the "outside option" for a taxed soda includes untaxed beverages that carry similar sugar loads. Dubois et al. (2020) confirm that this matters empirically: consumers who switch from taxed sodas to untaxed fruit juices generate no externality reduction.

Diet and artificially sweetened drinks pose a different design question. Most jurisdictions exclude them from the tax base, treating them as a "clean" substitute ($e_k \approx 0$). Philadelphia is the notable exception, taxing diet drinks at the same rate as regular soda. This effectively treats the tax base as "beverages" rather than "sugar," and the political backlash that led to the rapid repeal of Cook County's (Chicago,IL) similar tax suggests the choice is consequential. From the perspective of our framework, whether to tax diet drinks depends on their role in the substitution structure: if taxing regular soda diverts consumers to diet soda (a clean substitute), then condition (i) is satisfied for that margin and excluding diet drinks is correct.

The tax *structure* also maps onto our conditions. All US city SSB taxes are flat volumetric: a 12oz Coca-Cola (39g sugar) and a 12oz Gatorade (21g sugar) pay the same tax. This ignores condition (ii) entirely. Grummon et al. (2019) argue that per-gram-of-sugar taxes are better targeted, and South Africa's design comes closest to the textbook prescription $\tau_j \propto e_j$. The UK's tiered Soft Drinks Industry Levy takes an intermediate approach with two sugar thresholds and successfully induced reformulation before the tax even took effect. But even South Africa's "good" design ignores the substitution structure: it does not account for whether consumers leaving a taxed soda switch to untaxed fruit juice ($e_0 > 0$) or to water ($e_0 = 0$).

Condition (iii) fails because budget consumers may have higher price elasticity (and thus higher $D_{j\to 0}$) but may also substitute toward cheap sugary alternatives rather than water. Seiler et al. (2021) study Philadelphia's tax and find a 46% demand reduction in the taxed area, but cross-shopping at stores outside city limits offsets more than half the effect. O'Connell and Smith (2024) embed this substitution structure in an optimal tax framework with market power and find that the concentrated brand ownership of the UK drinks market creates additional distortions that interact with the tax design problem.

3.4 Vehicles and Emissions

There are multiple potential externalities in vehicles. The most obvious one is tailpipe emissions and specifically CO_2. This market features the most dramatic failure of condition (ii): externality content ranges from near-zero for electric vehicles to very high for large trucks and SUVs. The intensive margin $\sum_k D_{j\to k}(e_j - e_k)$ dominates the problem, since

most of the policy-relevant substitution is between vehicle types, not the extensive margin of car ownership.

Condition (i) is ambiguous: the outside option could be keeping an older, dirtier vehicle (e_0 high) or reducing driving ($e_0 = 0$). Jacobsen et al. (2023) show that when the outside option is the used car fleet, the scrappage margin creates important interactions that a naive emissions tax misses.

A fuel tax (roughly proportional to e_j, since CO_2 per gallon is nearly constant across fuels) works reasonably well in this market because it operates on the *intensive margin of driving*, where the proportionality between fuel consumed and emissions is tight. However, policies that operate on the *extensive margin of vehicle choice*—such as CAFE standards or feebate schemes—face the full force of the differentiated products problem. Ito and Sallee (2018) show that attribute-based fuel economy regulations (where compliance depends on vehicle weight) create perverse incentives to distort the secondary attribute, a mechanism analogous to the markup distortions documented for sin goods.

Vehicle weight generates a second externality distinct from emissions: crash incompatibility. Anderson and Auffhammer (2014) estimate that being struck by a vehicle 1,000 pounds heavier raises fatality risk by 40–50%, with total crash-weight externalities of $136 billion annually. This externality is particularly relevant as the vehicle fleet shifts toward heavier battery electric vehicles (BEVs). A Ford F-150 Lightning weighs roughly 35% more than a comparable gasoline F-150, largely due to the battery pack. The gas tax historically served as a rough proxy for weight-based externalities because heavier vehicles consumed more fuel, but this correlation breaks down entirely for BEVs, which are heavy but pay no fuel tax. White (2004) frames this as a "weight arms race": each consumer's decision to buy a heavier vehicle imposes uncompensated risk on occupants of lighter vehicles. The optimal instrument is a weight-varying mileage tax (Anderson and Auffhammer, 2014), but no US jurisdiction currently implements one.

We have data on roughly 50,000 automobile purchasers from MaritzCX (the same data used by Grieco et al. (2024)). In this survey each consumer is asked not only which car they purchased, but also which car they *would have purchased* if they could not have purchased their first-choice car. This nearly provides a direct measure of $D_{j\to k}$ (but without an outside option). We can use estimates in Conlon et al. (2025) to obtain corrected estimates of both $D_{j\to k}$ and $D_{j\to 0}$ for all products and compute $D^e_{j\to 0}$ for both the greenhouse gas externality and the vehicle weight externality.

[This table is not cleared yet]

3.5 E-Cigarettes and Nicotine Products

The nicotine products market may be the setting where the $D^e_{j\to 0}$ framework is most immediately policy-relevant. The externality-generating attribute is typically taken to be nicotine, but the health damage from nicotine products depends overwhelmingly on the *delivery mechanism*: combustion, heating, or vaporization, rather than nicotine content per se. This means condition (ii) fails dramatically: a combustible cigarette and an e-cigarette may deliver similar nicotine but differ by an order of magnitude in health harm.

Condition (i) fails for the same reason it fails in cannabis: the outside option for e-cigarettes is often combustible cigarettes, where $e_0 \gg e_j$. Taxing vaping products aggressively, as many states now do, or banning them outright, risks pushing consumers back toward the far more harmful combustible alternative. The policy parallels the cannabis case, where overtaxing legal products diverts consumers to a worse outside option, but with arguably stronger evidence on the relative harms.

Condition (iii) also fails: budget disposable vapes and premium refillable devices likely have very different diversion patterns to combustible cigarettes (and possibly greater non-nicotine health costs). The implication is that a uniform nicotine tax across all product forms could be counterproductive, raising total health damage by discouraging substitution away from combustion.

Massachusetts, California, New Jersey, New York, Maryland, DC, and 350+ local jurisdictions have some form of ban on flavored e-cigarettes or nicotine products, and San Francisco has banned all e-cigarette sales entirely (but not the sales of combustible cigarettes).

Being able to measure $D^e_{j\to 0}$ from data is crucial to making informed policy decisions.[7] For example, in 2022 the FDA ordered the removal of products from the leading maker of flavored e-cigarettes, JUUL. On one hand, if the main outside good substitution was that teens never started any nicotine product, this would be a public health success. On the other hand if the JUUL users largely diverted to combustible cigarettes or lower quality e-cigarettes, the net effect might be to increase rather than decrease the external damage.

The proposed federal ban on menthol cigarettes illustrates the same logic in reverse. Menthol cigarettes account for roughly a third of US cigarette sales and are disproportionately used by Black smokers. The FDA proposed banning menthol as a characterizing flavor in 2022, but the rule was never finalized. Massachusetts and California have implemented state-level bans. The welfare case for a ban depends entirely on the diversion structure: the FDA projected that a ban would cause 923,000 additional smokers to quit, but evidence from Massachusetts shows substantial diversion to non-menthol cigarettes and cross-border pur-

[7] A nice feature of diversion $D_{j\to k}$ is that it can be measured in terms of second-choices as well as small price changes. See Conlon and Mortimer (2021).

chasing in neighboring states (Gammon et al., 2022). A ban is equivalent to an infinite tax on one product variant, and if most of the diversion is to non-menthol cigarettes ($e_k \approx e_j$), the intensive margin externality reduction is near zero. If some diversion is to illicit menthol products ($e_0 \geq e_j$), the ban could increase total harm on that margin. The Canadian federal menthol ban in 2017 provides a cleaner test; Stoklosa et al. (2020) find limited illicit market growth, though some switching to non-menthol variants.

4 Policy Implications

Three practical lessons emerge from this analysis.

1. When the outside option is "dirty," tax rates should be lower than naive Pigouvian rates. This is most relevant for cannabis, where the illicit market poses greater risks than the legal market, but also applies to alcohol (if the outside option includes other intoxicating substances), SSBs (if untaxed sugary alternatives are close substitutes), and vehicles (if the outside option is keeping an old, high-emitting car). In the extreme, $D^e_{j\to 0}$ can be negative, implying that the optimal correction involves *subsidizing* the legal product.

2. The best simple tax depends on the covariance structure. When product-specific Ramsey taxes are administratively infeasible, the welfare cost of using a simple attribute tax depends on how well e_j proxies for $D^e_{j\to 0}$. If $\text{Cov}(e_j, D_{j\to 0}) > 0$ (products with higher externality content also have higher diversion to the outside good), an attribute tax does well. If this correlation is weak or negative, alternative instruments—such as minimum price floors, which effectively target the cheapest products with the highest $D_{j\to 0}$—can outperform attribute taxes. This helps explain the adoption of minimum unit pricing for alcohol in Scotland (Griffith et al., 2022) and the discussion of price floors in cannabis regulation.

3. Market power amplifies the cost of naive taxation. Firms with market power set markups proportional to inverse demand elasticity, which is typically correlated with product quality and cost but not with externality content. In every market considered here, the existing markup pattern works *against* the externality correction: premium products are overpriced (relative to the externality-correcting benchmark) and budget products are underpriced. This means that in concentrated markets, the welfare gap between a simple attribute tax and the optimal product-specific correction is larger than under perfect competition.

4. Constructing Estimates of $D_{j\to k}$ and $D^e_{j\to 0}$ is easier than it looks. For our distilled spirits example we used the estimated demand system from Conlon and Rao (2025) to recover $D_{j\to k}$. In many cases, the data will not be sufficient to estimate a full demand system. However, in our automobile example we relied instead of second-choice surveys to estimate $D_{j\to k}$ instead.

5 Conclusion

The conventional wisdom that externalities should be corrected by taxing the externality-generating attribute is an implication of the additivity principle *plus* the assumption that products within a market have independent demands. In any differentiated product market (and every sin-good market is one), the marginal external damage from product j depends on the full substitution structure, captured by the substitution-adjusted measure $D^e_{j\to 0}$. The three conditions under which a simple attribute tax suffices (clean outside option, homogeneous externality content, uniform diversion to outside good) fail to varying degrees in alcohol, cannabis, sugar-sweetened beverages, and vehicle emissions markets.

The practical lesson for tax design is that policymakers should evaluate proposed excise taxes not only by how well they target the externality-generating attribute, but by how well they account for *where consumers go* when prices rise. In some markets, notably cannabis, ignoring this substitution structure could make policy counterproductive.

References

Allcott, Hunt, Benjamin B. Lockwood, and Dmitry Taubinsky, "Regressive Sin Taxes, with an Application to the Optimal Soda Tax," *Quarterly Journal of Economics*, 2019, *134* (3), 1557–1626.

_ , _ , and _ , "Should We Tax Sugar-Sweetened Beverages? An Overview of Theory and Evidence," *Journal of Economic Perspectives*, 2019, *33* (3), 202–227.

Anderson, Michael L. and Maximilian Auffhammer, "Pounds That Kill: The External Costs of Vehicle Weight," *Review of Economic Studies*, 2014, *81* (2), 535–571.

Conlon, Christopher T. and Julie Holland Mortimer, "Empirical Properties of Diversion Ratios," *RAND Journal of Economics*, 2021, *52* (4), 693–726.

_ and Nirupama Rao, "The Cost of Curbing Externalities with Market Power: Alcohol Regulations and Tax Alternatives," 2025. Working Paper, NYU Stern.

_ , **Julie Holland Mortimer, and Paul Sarkis**, "Estimating Preferences and Substitution Patterns from Second-Choice Data Alone," 2025. Working Paper.

_ , **Nirupama Rao, and Yutong Wang**, "Who Pays Sin Taxes? Understanding the Overlapping Burdens of Corrective Taxes," *Review of Economics and Statistics*, 2024, *106* (6), 1719–1736.

Dubois, Pierre, Rachel Griffith, and Martin O'Connell, "How Well Targeted Are Soda Taxes?," *American Economic Review*, 2020, *110* (11), 3661–3704.

Gammon, Dorian G., Todd B. Rogers, Ellen M. Coats, Brett R. Loomis, and Matthew C. Farrelly, "National and State Patterns of Concept Cigarette Sales Following Implementation of a Statewide Menthol Cigarette Sales Restriction," *Nicotine & Tobacco Research*, 2022, *24* (8), 1212–1218.

Grieco, Paul L. E., Charles Murry, and Ali Yurukoglu, "The Evolution of Market Power in the U.S. Automobile Industry," *Quarterly Journal of Economics*, 2024, *139* (2), 1201–1253.

Griffith, Rachel, Martin O'Connell, and Kate Smith, "Tax Design in the Alcohol Market," *Journal of Public Economics*, 2019, *172*, 20–35.

_ , _ , **and** _ , "Price Floors and Externality Correction," *Economic Journal*, 2022, *132* (646), 2273–2289.

Grummon, Anna H., Benjamin B. Lockwood, Dmitry Taubinsky, and Hunt Allcott, "Designing Better Sugary Drink Taxes," *Science*, 2019, *365* (6457), 989–990.

Han, Joohun, Khondoker A. Mottaleb, John N. Ng'ombe, and Alvaro Durand-Morat, "Does Sin Tax on the Legal Market Facilitate the Illicit Market? An Ex-Ante Assessment on the US Cannabis Market," *Journal of the Agricultural and Applied Economics Association*, 2024, *3* (2), 330–345.

Hansen, Benjamin, Keaton Miller, Boyoung Seo, and Caroline Weber, "Taxing the Potency of Sin Goods: Evidence from Recreational Cannabis and Liquor Markets," *National Tax Journal*, 2020, *73* (2), 511–544.

Hollenbeck, Brett and Kosuke Uetake, "Taxation and Market Power in the Legal Marijuana Industry," *RAND Journal of Economics*, 2021, *52* (3), 559–595.

Ito, Koichiro and James M. Sallee, "The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel Economy Standards," *Review of Economics and Statistics*, 2018, *100* (2), 319–336.

Jacobsen, Mark R., Christopher R. Knittel, James M. Sallee, and Arthur A. van Benthem, "Sufficient Statistics for Imperfect Externality-Correcting Policies," Working Paper 22063, NBER 2016.

_ , James M. Sallee, Joseph S. Shapiro, and Arthur A. van Benthem, "Regulating Untaxable Externalities: Are Vehicle Air Pollution Standards Effective and Efficient?," *Quarterly Journal of Economics*, 2023, *138* (3), 1907–1976.

Kopczuk, Wojciech, "A Note on Optimal Taxation in the Presence of Externalities," *Economics Letters*, 2003, *80* (1), 81–86.

Miller, Keaton and Boyoung Seo, "The Effect of Cannabis Legalization on Substance Demand and Tax Revenues," *National Tax Journal*, 2021, *74* (1), 107–145.

O'Connell, Martin and Kate Smith, "Optimal Sin Taxation and Market Power," *American Economic Journal: Applied Economics*, 2024, *16* (4), 34–70.

Sandmo, Agnar, "Optimal Taxation in the Presence of Externalities," *Swedish Journal of Economics*, 1975, *77* (1), 86–98.

Seiler, Stephan, Anna Tuchman, and Song Yao, "The Impact of Soda Taxes: Pass-Through, Tax Avoidance, and Nutritional Effects," *Journal of Marketing Research*, 2021, *58* (1), 22–49.

Stoklosa, Michal, Jeffrey Drope, and Frank J. Chaloupka, "Effect of the Menthol Cigarette Ban in Canada on Illicit Cigarette Use and Sales of Roll-Your-Own Tobacco," *Tobacco Control*, 2020, *29* (Suppl 4), s247–s251.

White, Michelle J., "The "Arms Race" on American Roads: The Effect of Sport Utility Vehicles and Pickup Trucks on Traffic Safety," *Journal of Law and Economics*, 2004, *47* (2), 333–355.